

R·I·T

2017 Imaging Science Ph.D. Qualifying Examination

June 9, 2017

9:00AM to 12:00PM

IMPORTANT INSTRUCTIONS

You must complete two (2) of the three (3) questions given for each of the core graduate classes. The answer to each question should begin on a new piece of paper. While you are free to use as much paper as you would wish to answer each question, please only write on one side of each sheet of paper that you use AND STAY INSIDE THE BOX! Be sure to write your provided identification letter, the question number, and a sequential page number for each answer in the upper right-hand corner of each sheet of paper that you use. When you hand in your exam answers, be certain to write your name on the supplied 5" x 8" paper containing your provided identification letter and place this in the small envelope, and then place this envelope along with your answer sheets in the large envelope.

**ONLY HAND IN THE ANSWERS TO THE QUESTIONS THAT YOU WOULD
LIKE EVALUATED**

Identification Letter: _____

**THIS EXAM QUESTION SHEET MUST BE HANDED BACK TO THE
PROCTOR UPON COMPLETION OF THE EXAM PERIOD**

Please show all work and write your answers clearly to obtain full credit for your solutions. Clearly label the problems that you are working out on a given page to ensure that the grader will be able to easily discern which problem is being worked on each page.

Core 1: Fourier Methods in Imaging, Select TWO from 1-3

1. Consider three imaging systems that act on one-dimensional functions.

(a) [20%] The transfer function of the first system is:

$$H_1[\xi] = STEP[\xi] \cdot \exp[+i \cdot \pi]$$

Find an expression for the image that results from this system if the input is $f[x] = \delta[x - 1]$.

(b) [10%] Sketch the image from part (a) as real part, imaginary part, magnitude, and phase

(c) [10%] The transfer function of the second system is:

$$H_2[\xi] = STEP[\xi] \cdot \exp[+i \cdot \pi \cdot \alpha_0 \cdot \xi \cdot STEP[\xi]]$$

What are the dimensions of the parameter α_0 ?

(d) [20%] Evaluate and sketch the image produced by this system for the same input used in part (a), again if the value of the parameter is $\alpha_0 = +2$.

(e) [20%] The transfer function of the third system is:

$$H_2[\xi] = \exp[+i \cdot \pi \cdot \alpha_0 \cdot |\xi|]$$

Sketch the real part, imaginary part, magnitude, and phase of this transfer function, again for the case $\alpha_0 = +2$.

(f) [20%] Evaluate and sketch the image produced by this system for the same input used in part (a).

2. A 1-D real-valued input object $f[x]$ has been “blurred” by convolution with the quadratic-phase impulse response:

$$h[x] = \exp \left[+i\pi \left(\frac{x}{\alpha_0} \right)^2 \right]$$

where α_0 is a positive real number with dimensions of length.

- (a) [40%] Find an expression for the impulse response $w[x]$ of the inverse filter.
(b) [60%] Determine the output:

$$f[x] * w[x] = g[x]$$

if $w[x]$ is the function from part (a) and:

$$f[x] = 4 \cdot \exp \left[+i\pi \left(\frac{x-4}{\alpha_0} \right)^2 \right] + 2 \cdot \exp \left[+i\pi \left(\frac{x+2}{\alpha_0} \right)^2 \right] + \exp \left[+i\pi \left(\frac{x+3}{4\alpha_0} \right)^2 \right]$$

and explain the features of this result.

3. Consider a camera with a discrete sensor that is an $N \times N$ array of square pixels of width d_0 spaced by the the same pitch in the orthogonal directions: i.e., that $\Delta x = \Delta y \geq d_0$. This camera is used to image a band-limited object.
- (a) [10%] Specify the limiting frequency of the image on the sensor that ensures that there is no aliasing; this is the Nyquist frequency.
 - (b) [20%] Assuming that the input object is bandlimited as determined by part (a), find an expression for the 1-D profile of the normalized MTF due to the finite sensor size along the x -axis if the pixel dimension d_0 is smaller than the pixel pitch Δx . In other words, the fractional area of the discrete sensor array that is sensitive to light (i.e., the “fill factor” of the array) is less than 100%.
 - (c) [10%] Evaluate the numerical value of the MTF at the Nyquist sampling frequency.
 - (d) [10%] Graph the 1-D MTF evaluated in part (b) in the case where the pixel size is equal to the pixel pitch: $d_0 = \Delta x$; the fill factor of this array is 100%.
 - (e) [20%] Now consider the case of a sensor with 100% fill factor that may be translated relative to the image of a stationary object projected by the optics onto the sensor plane. The following procedure is used to collect the image: (1) the image is recorded with the sensor in its original position; (2) the sensor is translated along the x -direction by $\frac{\Delta x}{2}$ and another image recorded; (3) the sensor is translated along the y -direction by $\frac{\Delta x}{2}$ and another image recorded; and (4) the sensor is translated along the x -direction by $-\frac{\Delta x}{2}$ and a fourth image recorded. The final image is formed by combining these four images to create a $2N \times 2N$ array of pixels with a sample pitch of $(\Delta x)_1 = \frac{\Delta x}{2}$. Determine the Nyquist frequency of the new imaging system and the value of the MTF of the system at this new Nyquist frequency.
 - (f) [10%] Graph the 1-D MTF of the new system.
 - (g) [20%] Characterize the possible benefits and limitations of images of stationary objects created with the new system.

Core 2: Optics, Select TWO from 4-6

Each answer must be accompanied by supporting illustrations. Be sure to label the axes and all important characteristics in your sketch. Illegible work will not be graded. Partial credit will be awarded if it demonstrates a level of competency.

Please use a straight edge ruler and a single-purpose calculator when needed. If you not have these, please tell the proctor, who will distribute them.

4. (Diffraction) Two point sources of wavelength λ and equal strength are located in the object plane at the distances $x_0 = \pm a$ from the optical axis. Assume the sources are mutually coherent (i.e., light from the sources can interfere).
- (a) Using Huygen's principle, determine a complex algebraic expression (in terms of x, z, x_0 , and any other pertinent variables) for the net electric field at the point x in a plane located a distance z from the object. Simplify the expression by use of a series expansion, assuming $x \pm a \ll z$.
 - (b) A lens of focal length \mathbf{f} and semidiameter R is placed a distance $z = 2\mathbf{f}$ from the point sources. Write an expression for the electric field at the immediate output face of the lens.
 - (c) Using the Huygens-Fresnel integral for the field in the x' plane, a distance Δz from the initial field $E[x]$:

$$E[x', \Delta z] = \frac{\exp[+ik \cdot \Delta z]}{i\lambda \cdot \Delta z} \int_{-\infty}^{+\infty} E[x] \exp\left[+ik \cdot \frac{(x - x')^2}{2 \cdot \Delta z}\right] dx$$

Determine a simplified expression for the complex electric field if $\Delta z = 2\mathbf{f}$. (You must evaluate the integral to receive credit for this problem.)

- (d) On the diagrams that you have drawn for the above problems, sketch the rays that provide a geometrical optics solution. Based on this sketch, write a simple expression for the electric field in the geometrical optics approximation. Prove that this expression agrees with your solution in part (c) in the limit as $R \rightarrow \infty$.

5. Geometrical Optics: An optical relay is a periodic series of lenses that transfers rays from one end of a device to another.
- (a) Using ray tracing, prove that a series of identical lenses of focal length \mathbf{f} , separated by a distance $d = \mathbf{f}$, will transfer rays without magnification if the object is placed a distance $d_0 = \mathbf{f}$ from the first lens. (Draw the initial object large enough to avoid drawing errors – roughly 2 cm should be sufficient).
 - (b) How many lenses are required to obtain an erect image of magnification $M = +1$? Support your answer with your ray diagram. (An image will appear a distance $d_i = \mathbf{f}$ after the last lens in the sequence.)
 - (c) Determine the magnification of this relay system if the object is instead placed a distance $d_o = 2\mathbf{f}$ from the initial lens in the sequence. (Hint: this may be solved using similar triangles. You may wish to make a ray tracing drawing to check your solution.)
 - (d) Verify your result in part (b) using ABCD matrices. That is, prove that the image height is equal to the object height if $d_o = d_i = \mathbf{f}$. (Hint: you may wish multiply the matrices in an organized fashion, e.g., pair-wise).

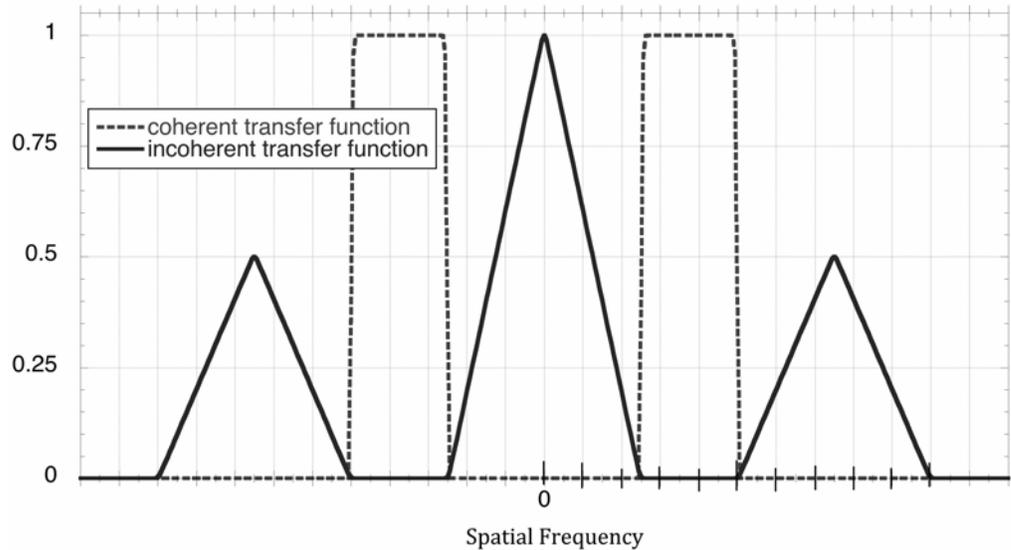
6. Modulation transfer function: The coherent and incoherent transfer functions for a one-dimensional lens of focal length $f = 158$ mm with a central obscuration and outer width R_{out} are shown in the plot below for an image distance of $d_i = 2f$. The OTF is a real-valued function. The incoherent cutoff frequency is $500 \frac{\text{cycles}}{\text{mm}}$ and the wavelength is $\lambda_0 = 633$ nm.

- (a) The aperture is comprised of two rectangle functions. Using the plot below and the equation for the incoherent cutoff frequency, determine the values of the inner and outer sizes R_{in} and R_{out} of the lens aperture function in units of mm.
- (b) An object having an intensity profile given by

$$I_o[x] = 1 + \cos\left[2\pi\frac{x_o}{\Lambda}\right]$$

is placed in the object plane, where $\Lambda = 2.667 \mu\text{m}$. Determine the intensity profile of the image $I_i[x_i]$, using the incoherent transfer function presented below.

- (c) Compare the obscured lens in the previous parts to an unobscured lens having the same outer width R_{out} and the same incoherent cutoff frequency. That is, plot the incoherent transfer function of the unobscured aperture. Determine the intensity profile of the image $I_i[x_i]$ for the unobscured lens assuming the same object given in (b).
- (d) Prove that the image contrast obtained for the obscured lens is twice as large as that for the unobscured lens.



Core 3: Vision, Select TWO from 7-9

7. The human visual system has been optimized over a long period of evolution. Today, the ‘typical eye’ has an acuity of 20/20 under ‘optimal conditions,’ has a field of view extending over 180° horizontally, can detect and differentiate among a huge number of different spectral mixtures while being relatively insensitive to changes in the spectral content of the illuminant, and over a range of illumination that spans at least 10 orders of magnitude, without saturating or suffering from noticeable noise.

But physical measurements of the optics of the human eye indicate (i) significant chromatic aberration when broadband illumination is used, (ii) significant spherical aberration even when illumination is limited to a narrow range of wavelengths, and (iii) severe loss of image quality even a few degrees away from the optical axis.

- (a) [20%] Describe how it is possible for humans to have such high perceived visual image quality when we know from physical measurements that severe chromatic and spherical aberrations are inherent in the relatively simple, two-element optical design of the eye.
- (b) [20%] Describe how it is possible for humans to have 20/20 acuity and such a wide field of view given what we know about how acuity falls off so dramatically off axis. [Include in your answer a definition of “20/20 acuity” and how it is measured.]
- (c) [20%] How does the visual system achieve such a high *between-scene* dynamic range? In other words, how is it possible to perceive detail in bright sunlit scenes at one time, then under moonlit and cloudy conditions later?
- (d) [20%] How does the visual system achieve such a high *within-scene* dynamic range? In other words, how is it possible to simultaneously perceive detail in bright, sunlit regions and in dark shadows?
- (e) [20%] We know from physiological studies that ‘color’ information is carried in the human visual system in only three neural channels, yet most observers can differentiate among millions of different spectral combinations. On the other hand, it is also true that there are sets of spectral ‘metameric pairs’ that appear identical, despite the fact that their spectra are completely different. Discuss this apparent dichotomy, including in your discussion some reason(s) that humans evolved to have such a small number of chromatic channels.

8. When trying to understand the performance of the human visual system (HVS), vision scientists measure both the modulation-transfer function (MTF) and contrast-sensitivity function (CSF) of observers.

- (a) [30%] Describe the MTF. Include in your description a sketch of the MTF for a typical observer under ideal conditions, with both axes labeled and scaled. Discuss why the MTF isn't "perfect," i.e., explain why the curve varies from an ideal system that passes all spatial frequencies perfectly without affecting the output.
- (b) [30%] Describe the CSF. Include in your description a sketch of the CSF for a typical observer under ideal conditions, with both axes labeled and scaled. Discuss why the CSF isn't "perfect," i.e., explain why the curve varies from an ideal system that passes all spatial frequencies perfectly without affecting the output.

NASA and DARPA are collaborating on a humanoid robot for use in extraterrestrial exploration

- (c) [20%] Among the sensors now available in robots are *binocular-vision sensors*. If you were to have an analogous metric to the MTF of the HVS, what would you measure in the robot's "visual system?"
- (d) [20%] If you were to have a measure parallel to the CSF of the HVS, what would you measure in the robot's "visual system?"



Jeff Pelz cannot be serious including this image of the illegitimate offspring of Gort and Robbie the Robot in this test.

9. This is Figure 6 from Daly's classic paper, *Visible differences predictor – an algorithm for the assessment of image fidelity*, which introduced an explicit 2D model for observer CSF. In referring to the figure, Daly says, “Shown in Figure 6 is a typical CSF for the 2-D frequency plane, showing the bandpass nature and oblique effect.”

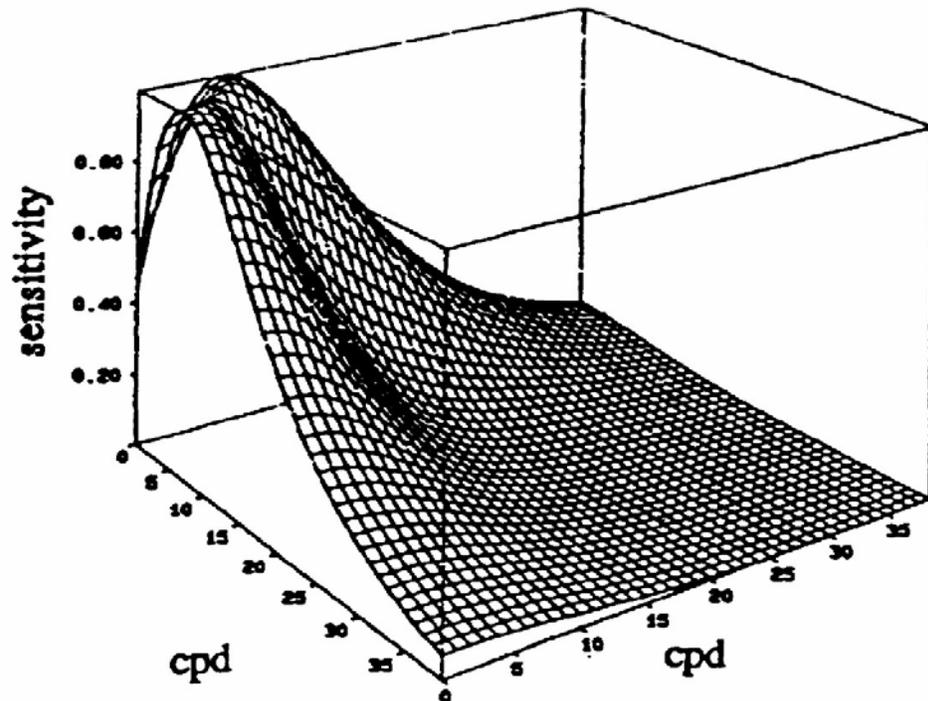


Figure 6

- (a) [20%] Describe the “bandpass nature” Daly refers to, including the spatial frequency of peak response and f_{\max} along the horizontal and vertical axes.
- (b) [30%] Explain the reason for this bandpass nature; specifically, the reason(s) for the loss of response at high and low frequencies.
- (c) [30%] What is the ‘oblique effect?’
- (d) [20%] Discuss to what extent man-made optical systems share either of these characteristics.

R·I·T

2017 Imaging Science Ph.D. Qualifying Examination

June 9, 2017

1:00PM to 4:00PM

IMPORTANT INSTRUCTIONS

You must complete two (2) of the three (3) questions given for each of the core graduate classes, except everyone must answer all parts of the Graduate Laboratory question. The answer to each question should begin on a new piece of paper. While you are free to use as much paper as you would wish to answer each question, please only write on one side of each sheet of paper that you use AND STAY INSIDE THE BOX! Be sure to write the identification letter provided to you this morning, the question number, and the page number for each answer in the upper right-hand corner of each sheet of paper that you use.

ONLY HAND IN THE ANSWERS TO THE QUESTIONS THAT YOU WOULD LIKE EVALUATED

Identification Letter:

THIS EXAM QUESTION SHEET MUST BE HANDED BACK TO THE PROCTOR UPON COMPLETION OF THE EXAM PERIOD

Core 4: Radiometry, Select TWO from 10-12

Several constants that you may or may not need in your solutions are provided here for your convenience:

| | |
|--|--|
| Planck's Constant | $h = 6.626 \times 10^{-34} \text{ J s}$ |
| Speed of light in vacuum | $c = 2.99792458 \times 10^8 \text{ m s}^{-1} \approx 3 \times 10^8 \text{ m s}^{-1}$ |
| Boltzmann Constant | $k = 1.381 \times 10^{-23} \text{ J K}^{-1}$ |
| Approximate Solar Temperature | $T_{\text{solar}} = 5800 \text{ K}$ |
| Approximate Room Temperature | $T_{\text{room}} = 312 \text{ K}$ |
| Earth-Sun Distance | $D_{\text{earth-sun}} = 1.4960 \times 10^{11} \text{ m}$ |
| Radius of the Sun | $R_{\text{sun}} = 6.955 \times 10^8 \text{ m}$ |
| Stefan-Boltzmann constant | $\sigma = 5.67 \times 10^{-8} \text{ W m}^{-2} \text{ K}^{-4}$ |
| Absolute zero on Kelvin scale converted to Celsius | $-273.15 \text{ }^\circ\text{C}$ |

10. You are studying the effect of rising ocean temperatures on the health of a stony coral species, *Cyphastrea serailia*, in the Pacific Ocean. You have built a submersible craft and designed a camera system that can be submerged in a waterproof case, along with two identical isotropic light sources as shown in Figure 1. You are studying a coral colony of this species that has attached itself to the side surface of an old shipwreck as shown. You are taking your measurements at night-time using only the two isotropic light sources shown as the illumination of the coral. The camera is affixed to a rod attached to the submersible and is separated from the corals by a distance d_c , while the two light sources are at the distance d_s from the surface of the coral. The light sources each have power Φ_{src} and are separated from each other by the distance w . The camera is mounted at the midpoint of your vessel halfway between the two light sources. In your study area, assume that the extinction of the water column $\beta_w(\lambda)$ is constant throughout the water column. The optics of your camera detector have a transmittance of τ_{cam} , the camera has f-number $f\#_{\text{cam}}$, with detector elements that are ℓ_{pixel} on each side and noise equivalent power NEP_{cam} . Assume that the camera is responsive to a narrow wavelength band around the wavelength λ_{cam} . Also assume that the coral surface is Lambertian.

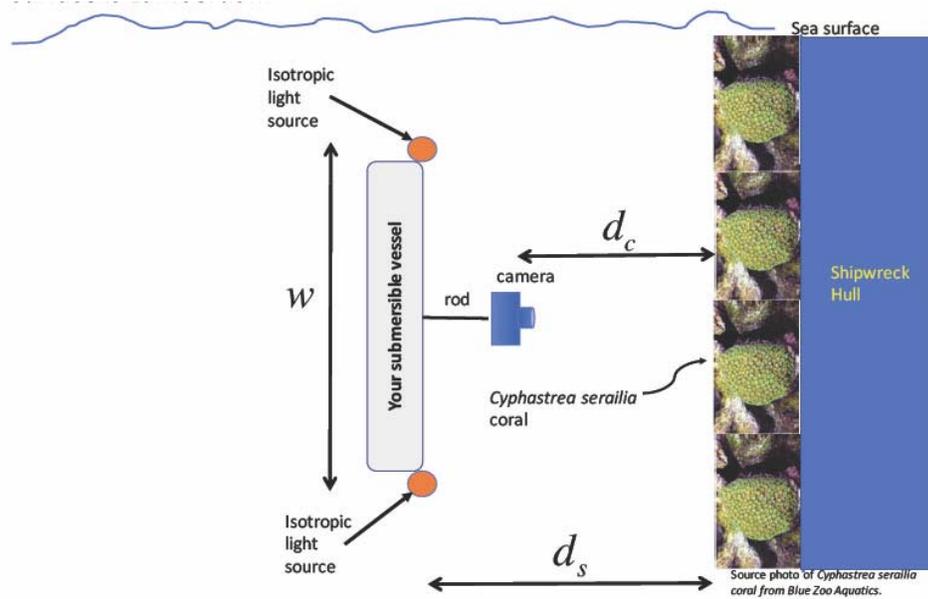


Figure 1. Camera system setup for coral monitoring.

- (a) [50%] Find a mathematical expression that could be used to determine whether a change in reflectance of the surface could be observed by your camera system in terms of the provided variables.
- (b) [30%] You were out just yesterday to take measurements, and the waters were at their seasonal average for this time of year, however, a warming trend that is expected over the next two weeks will expose the coral at your site to warmer waters. You are familiar with a recent study in which the effect of warming waters on coral reflectance was studied as a function of time (Figure 2 below). You examine the reflectance curves and need to decide whether your camera system could detect the changes observed in this study. For the listed values of the parameters:

| |
|---|
| $\Phi_{\text{src}} = 800 \text{ W}$ |
| $\lambda_{\text{cam}} = 600 \text{ nm}$ |
| $\tau_{\text{cam}} = 0.95$ |
| $NEP_{\text{cam}} = 1.25 \times 10^{-11} \text{ W}$ |
| $\ell_{\text{pixel}} = 10 \mu\text{m}$ |
| $f\#_{\text{cam}} = 1.4$ |
| $\beta_w(\lambda_{\text{cam}}) = 0.16 \text{ m}^{-1}$ |
| $w = 2 \text{ m}$ |
| $d_c = 0.75 \text{ m}$ |
| $d_s = 1 \text{ m}$ |

determine whether your camera would be able to detect a change in the reflectance after 14 days like that observed in the study using the results that you found in part a. In Figure 2 below, you will use the data in the right-hand column which

shows the impact of heating on a variant of this coral that is found offshore (as is the case on the shipwreck described in this problem).

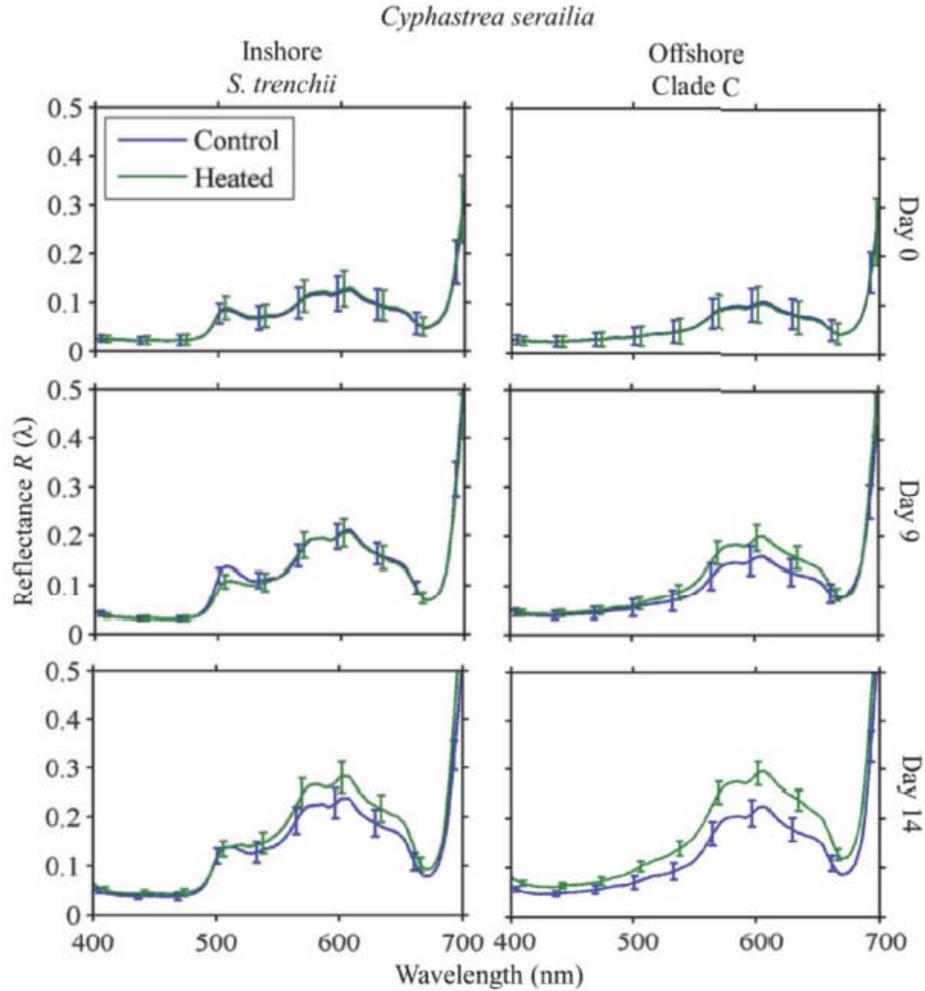


Figure 2. Figure from (Russell, Dierssen, et al., 2016) showing the effect on reflectance for two variants of the coral *Cyphastrea serailia* (in this problem, use the data in the right-hand column in your analysis).

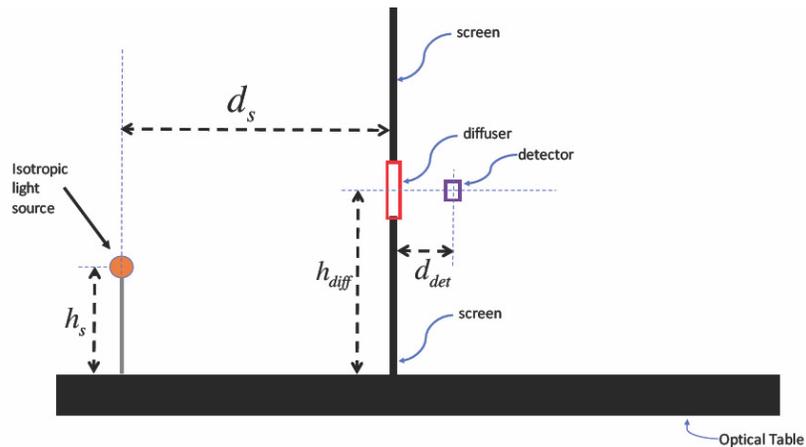
- (c) [20%] You later modify your camera setup on the submersible craft so that the rod holding the camera is a different length, but after the modifications you take measurements with the submersible craft being the same distance from the coral. That is, the distance of the light sources from the coral remains the same value of d_s , but the distance of the camera from the coral surface is now a new distance d_n (rather than d_c). find an expression relating the minimum detectable change in reflectance found in part a to the minimum detectable change in reflectance with the new rod length.

11. You are setting up an experiment on an optical table to test a new detector as shown in side view in the Figure below. An isotropic light source with radiant intensity I_s is mounted at a height h_s above the table at a distance d_s from a screen with an embedded Lambertian diffuser with transmittance τ_{diff} . Assume that the light is monochromatic with wavelength λ_s . The diffuser is circular with radius R_{diff} and is at a height h_{diff} above the optical table. Close to the diffuser on the other side of the screen is the new detector, which is aligned along the center axis of the diffuser at a distance d_{det} from the diffuser as shown. The diameter of the detector is D_{det} , and the detector has responsivity $\mathfrak{R}(\lambda)$.

- (a) [40%] Find an expression for the photon flux into the detector in terms of one or more of the given variables.
- (b) [20%] If the detector integrates the incoming light over a period t_{int} , what is the shot noise of the detector in electrons in terms of one or more of the variables provided so far?
- (c) [20%] Evaluate numerically the shot noise in “electrons” found in part (b) for the following numerical values of the parameters:

| |
|--|
| $I_s = 11.14 \frac{\text{W}}{\text{sr}}$ |
| $h_s = 0.1 \text{ m}$ |
| $d_s = 2.5 \text{ m}$ |
| $\tau_{\text{diff}} = 0.85$ |
| $\lambda_s = 1.13 \mu\text{m}$ |
| $R_{\text{diff}} = 0.23 \text{ m}$ |
| $h_{\text{diff}} = 0.45 \text{ m}$ |
| $d_{\text{det}} = 0.1 \text{ m}$ |
| $D_{\text{det}} = 14.4 \mu\text{m}$ |
| $\mathfrak{R}(\lambda) = 0.37 \frac{\text{A}}{\text{W}}$ |
| $t_{\text{int}} = 2.25 \times 10^{-3} \text{ s}$ |

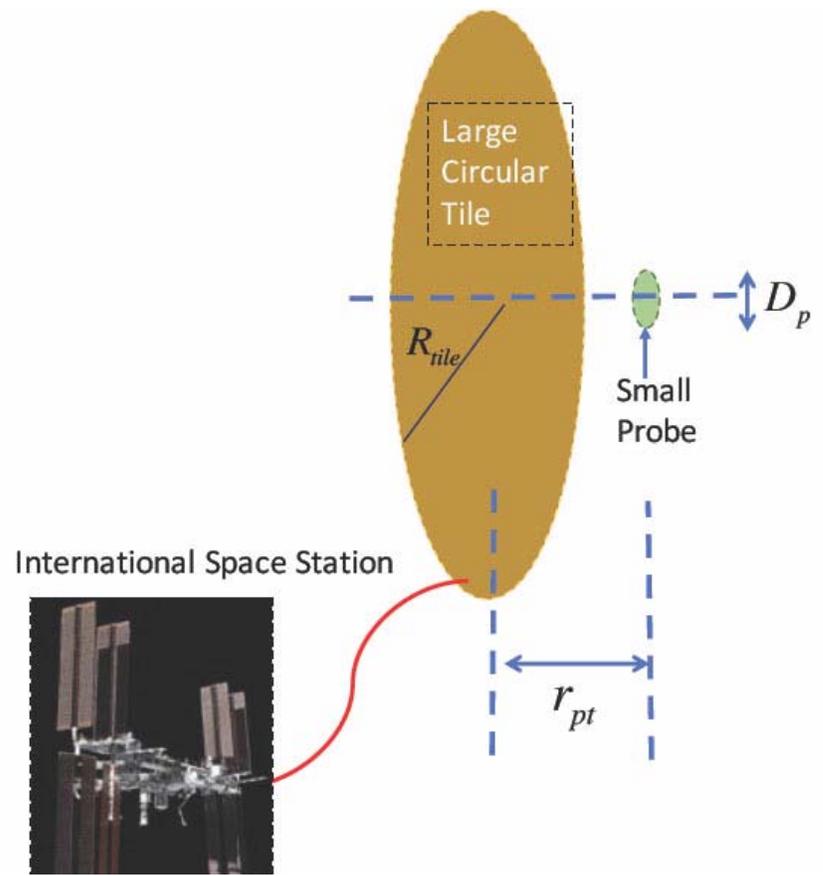
- (d) [20%] If the other sources of noise (e.g., read noise) of the system total 825 electrons, what is the total noise of the system? Is the system shot noise limited?



12. Imagine that you are an astronaut on the International Space Station (ISS), and you are taking a space-walk with a thermal probe. You are examining a large circular tile attached externally to the ISS. Through the link to the ISS, the circular tile is heated and maintains a constant temperature T_{tile} . The radius of the large circular tile is R_{tile} ; the tile has emissivity ϵ_{tile} . In close proximity to the tile, you position a very small circular probe facing the tile and aligned along the same central axis as the tile. The small probe has diameter D_p and emissivity ϵ_p . The distance between the probe and the large circular tile is r_{pt} . The tile is insulated on the side that is not facing the probe so that you can assume that there are no energy losses from the back side of the tile. In this problem, you can assume that the sun is hidden behind the Earth and that there are no other sources of radiation influencing the probe. Also, assume that the tile and the probe can each be treated as a greybody. In some parts of the problem below you are asked to evaluate the mathematical expression. In other parts, you will evaluate your expression numerically using one or more of the following values:

| |
|------------------------------------|
| $T_{\text{tile}} = 335 \text{ K}$ |
| $R_{\text{tile}} = 125 \text{ cm}$ |
| $\epsilon_{\text{tile}} = 0.63$ |
| $D_p = 5 \text{ cm}$ |
| $\epsilon_p = 0.92$ |
| $r_{\text{pt}} = 4 \text{ cm}$ |

- (a) [30%] Find an expression for the flux onto the small probe from the large circular tile.
- (b) [10%] Numerically evaluate your result for the flux expression that you calculated in part (a) using one or more of the values provided above.
- (c) [30%] If the uncertainty in the temperature of the tile is 1.3%, and the uncertainty in the emissivity of the tile is 2.1%, and you can assume that other sources of uncertainty are negligible, find an expression for the uncertainty in the flux expression derived in part (a) and evaluate this expression for the uncertainties in the temperature and emissivity of the tile.
- (d) [20%] The probe reaches an equilibrium temperature after some period of time. Find a mathematical expression for the temperature of the small probe.
- (e) [10%] Evaluate the equilibrium temperature of the probe found in part (d) using one or more of the values given in the table above.



Large circular tile and small probe

Core 5: Image Processing/Computer Vision, Select 2 from 13-15

13. Inferring Video Tags: You have decided to compete on the *YouTube-8M dataset* to win the \$100,000 prize. This dataset consists of 7 million YouTube videos (450,000 hours) and the goal is to infer multiple labels/tags for each individual video. For example, a particular video might have the tags ‘Video Game,’ ‘PC Game,’ and ‘Call of Duty.’ Each video has an average of 3.4 labels.

Google has extracted “Convolutional Neural Network” (CNN) features for you using the Inception-V3 CNN model, so for a particular video with T frames, you will have a collection of $V = [f_1, f_2, \dots, f_T]$ CNN feature vectors, where each $\mathbf{f}_t \in \mathbb{R}^d$. In a similar manner, they have also extracted audio features using a VGG-like CNN for audio data. Both of these are extracted at 1-second intervals.

- (a) [30%] Explain why transfer learning using pre-trained CNNs is extremely effective, and what we mean by an image embedding.
- (b) [40%] You have decided to use bag-of-visual-words to turn the T frame-level features for a video into a single vector. Give pseudo-code for training the bag-of-visual-words model and then pseudo-code for assembling the bag-of-visual-words feature vectors for a video V .
- (c) [20%] Explain why a typical multi-class classifier (e.g., *softmax*) is inappropriate for this problem. What kind of approach should we use for this problem?
- (d) [10%] Describe two ways for incorporating audio features into the system.

14. Segmentation: “Shot detection” is detecting where a video transitions or “cuts” from one shot to another, i.e., the positions in a video in which one scene is replaced with different visual content. This occurs frequently in broadcast video, as in these cases:
- Cutting from one person talking to another person in a show.
 - Cutting from a show to advertising.
 - Cutting from an on-site news reporter to an anchor at a desk.

Software-based shot detection is used heavily in video post-production. It is also used for automatic video indexing, content-based video retrieval, and video summarization.

You have been hired to do shot detection in broadcast video. Specifically, given a sequence of frames I_1, \dots, I_n representing a video, your task is to break this sequence into subsequences of similar frames. You should assume that $n = 10^5$, with each frame being a 1920×1080 image (1080p video).

- [30%] Describe a simple and computationally efficient algorithm for computing the possible shots, and explain the pros and cons of your approach.
- [10%] What is the major problem with using k -means for shot detection?
- [30%] One potential algorithm you could use to break the video into multiple shots is “mean-shift.” Explain how mean-shift works.
- [10%] What are two of the biggest challenges with applying mean-shift to this problem? Explain.
- [10%] After completing the shot-detection project, you have been asked to do a semantic segmentation for every frame of the video. What is the input and output of a semantic segmentation algorithm? In other words, what do these algorithms do?
- [10%] What is the difference between semantic segmentation and instance segmentation?

15. Creating Image Panoramas: In this problem, we will go through the steps for automatically stitching an image A to an image B to create a panorama by using a “homography.”
- (a) [20%] A homography matrix \mathbf{H} maps from image 1 to image 2, and you have pixel coordinates (x_1, y_1) on image 1 in non-homogeneous coordinates. Give the equations for computing the corresponding pixel location (x_2, y_2) in image 2 in non-homogeneous coordinates.
 - (b) [10%] What are the dimensions of a homography matrix? How many point correspondences are needed to compute the matrix for two images? Explain where the number comes from.
 - (c) [10%] Why are homogeneous coordinates used with a homography matrix? Explain.
 - (d) [10%] Describe the three properties that interest points should have.
 - (e) [10%] Briefly explain why RANSAC is used when fitting a homography matrix and give the steps.
 - (f) [20%] Give the high-level steps for automatically stitching an image A to an image B to create an image mosaic or panorama using a homography.
 - (g) [20%] When stitching two images together, there may be obvious artifacts that make the panorama look mediocre due to changes in lighting conditions (brightness, color, etc.) between the two images. Describe a technique to overcome this problem.

Core 6: Graduate Laboratory: Answer the question

16. In the Grad Lab project this year, the group was developing a system to detect automobile emissions. The class considered two detectors: a quad thermopile from Dexter, and a 2d array.
- (a) [20%] What are the pros and cons of the original Dexter thermopile point detector?
 - (b) [20%] What are the pros and cons for the 2d array?
 - (c) [20%] What was attempted to get the 2d array to work? Comment on what you learned in attempting each. Be specific.
 - (d) [10%] What are the primary sources of noise that have to be considered in the design? Which is the largest source of noise?
 - (e) [20%] Explain, using simple drawings to help, how the different bandpass filters allow the detection of a specific gas. How would a chopper wheel aid in the detection?
 - (f) [10%] If you could go back to the beginning of the project, what would you do differently in the project design/build/test?