CS 184/284A Spring 2021 **Name:** _____

Exam 2 (Online) **SID Number:** _____

April 22, 2021

Time Limit: 120 minutes

IMPORTANT: PLEASE READ CAREFULLY. NEW INSTRUCTIONS FOR 2021.

- Please carefully read, sign and submit the Honor Code described on the next page.
- You will need your own blank paper, a pen, a web browser, your class Gradescope account, and an internet-connected device capable of taking and uploading photos (a smartphone is ideal).
- You will have 2 hours to complete and submit your exam on Gradescope. You can choose to start and complete the exam anytime during the 24 hours of April 22nd 2021, California time (i.e. 12:01am to 11:59pm).
- During the 24 hour period, you must not communicate with anyone about the exam, including posting to or reading from the internet any information about the exam questions. This is a strict requirement.
- This exam is open book, and open internet, subject to the limitations described above.
- This exam is 120 minutes, and has a total of 110 points. There are 13 pages (including this cover page) and 7 parts. Problem difficulty varies throughout the exam, so move on if you find yourself stuck. You are encouraged to show your work and what you know for partial credit.
- For some questions, you will have to handwrite the solution and upload a picture/file to Gradescope. For these, please write out each answer on a separate page and upload a separate file for each question. Uploading time is built into the 2 hours, but if you encounter uploading difficulties you may email your photos to rishi.upadhyay@berkeley.edu with subject 'CS184 Exam'.
- For fairness to students taking the exam at different times, the teaching staff will not answer any questions or post any corrections about the exam during the 24 hour period.

Problem	Points	Score
1	0	
2	18	
3	16	
4	12	
5	24	
6	18	
7	22	
Total:	110	

1. Honor Code and Academic Integrity Certification

The Honor Code is the commitment and work of students individually and as a community to uphold honesty and integrity. This commitment is comprehensive. Specifically on exams, it means not giving or receiving any help, or using any resources that are not permitted. The Honor Code requires that students take an active role in seeing to it that others as well as themselves uphold the letter and spirit of the Honor Code.

By uploading my signature I certify that:

- All the work submitted in my name for this exam is my work alone.
- I have not given or received any help during this exam.
- I have not used any un-permitted resources during this exam (see cover page for permitted resources).
- I have timed myself and used no more than the allotted time to write my answers for this exam.
- If I become aware of any Honor Code violations related to this exam, I will inform the course staff immediately.

I understand that a false statement on this declaration is a violation of the Honor Code, and would be subject to the highest level of disciplinary sanction under the University Code of Conduct.

1.i. Certification

UPLOAD PHOTO **On a sheet of your own paper, please sign to provide your certification of this statement, write your name, SID number and the date. Take a photo of your page and upload to Gradescope Exam 2.**

Signature: _____

Name: _____

SID Number: _____

Date: _____

2. (Total : 18 points) True / False

Mark each statement true or false. (1 point each)

(2a) (1 point) ___ For any mesh, it is possible to generate a single 2D texture image (or a UV mapping) that induces no distortion.

(2b) (1 point) ___ If a vertex on the mesh is deformed by adding an offset to it, it's UV coordinate also changes.

(2c) (1 point) ___ In image processing, "magnification" means that you are sampling from a smaller region in the original image.

(2d) (1 point) ___ Intersecting a ray with a sphere involves solving a quadratic equation.

(2e) (1 point) ___ With the microfacet BRDF model, the surface will appear more shiny if the microfacet normal distribution function is more concentrated about the surface normal.

(2f) (1 point) ___ An anisotropic BRDF is effectively a three-dimensional function.

(2g) (1 point) ___ The field of view of a camera increases as the sensor is brought closer to the lens, all else being equal.

(2h) (1 point) ___ As long as it remains further than the focal point of the lens, a sensor that is brought closer to the lens will create photographs that are focused further away in the world.

(2i) (1 point) ___ In the JPEG image compression standard, we encode the spatial variation in lightness at higher resolution than spatial variaion in hue and saturation.

(2j) (1 point) ___ In a perfect image sensor with no device-related noise, a pixel that captures 10,000 photoelectrons will have an SNR of 100.

(2k) (1 point) ___ For a time step of 0.1 seconds, it's reasonable to model the forces applied to a particle in any system as a constant.

(2l) (1 point) ___ Verlet integration dissipates energy in the system by recomputing the velocity at the constrained position.

(2m) (1 point) ___ The average of any two 3x3 rotation matrices results in a valid rotation matrix.

(2n) (1 point) ___ Blend shapes can be learned from data, if the data is well-processed so all samples have the exact same number of vertices that are in correspondence with each other.

(2o) (1 point) ___ The watercolor illusion reveals that human color perception exhibits a "flood-fill" component.

(2p) (1 point) ___ Maxwell's color matching experiment showed that human color perception is three dimensional, because three primary colors form a basis for matching the appearance of any color.

(2q) (1 point) ___ In human color vision, there is no spectral power distribution that can stimulate only the M cone cell.

(2r) (1 point) ___ A display with red, green, blue and yellow pixels could reproduce any color in the human gamut.

3. (Total : 16 points) Graphics Pipeline and Geometry

(3a) (6 points) Transforms and coordinate frames.
Ann and Bali are taking a photo of a cat:

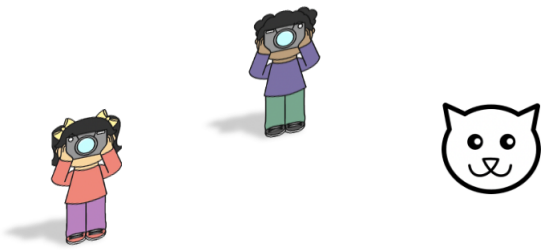

In the camera coordinate frame of Ann, the cat's location is X_A.
The world-to-camera transform of Ann's camera is $T_A \in \mathbb{R}^{4\times4}$.
The world-to-camera transform of Bali's camera is $T_B \in \mathbb{R}^{4\times4}$.
What is the location of the cat, X_B, in Bali's camera coordinate frame?
Write its expression.

(3b) (2 points) Which of these statements are true? Homogenous coordinates can:

(i) express sum of two translation vectors as a matrix multiplication.
(ii) avoid division operation in perspective projection
(iii) express both points and vectors
(iv) makes affine transformation commutative

(3c) Take this scene below. The direction of the incoming light is $l = (-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 0)$. We sampled 3 points in this scene, x_1, x_2, x_3, and computed their surface normal, which are:

$$n_1 = (-\frac{1}{\sqrt{2}}, \frac{1}{\sqrt{2}}, 0)$$
$$n_2 = (1, 0, 0)$$
$$n_3 = (0, 1, 0).$$

3c.i. (6 points) Compute the shading at the 3 sampled points using Lambertian shading. Assume that the incoming radiance is 1 everywhere in this scene and the diffuse coefficient is $k_d = 1$.

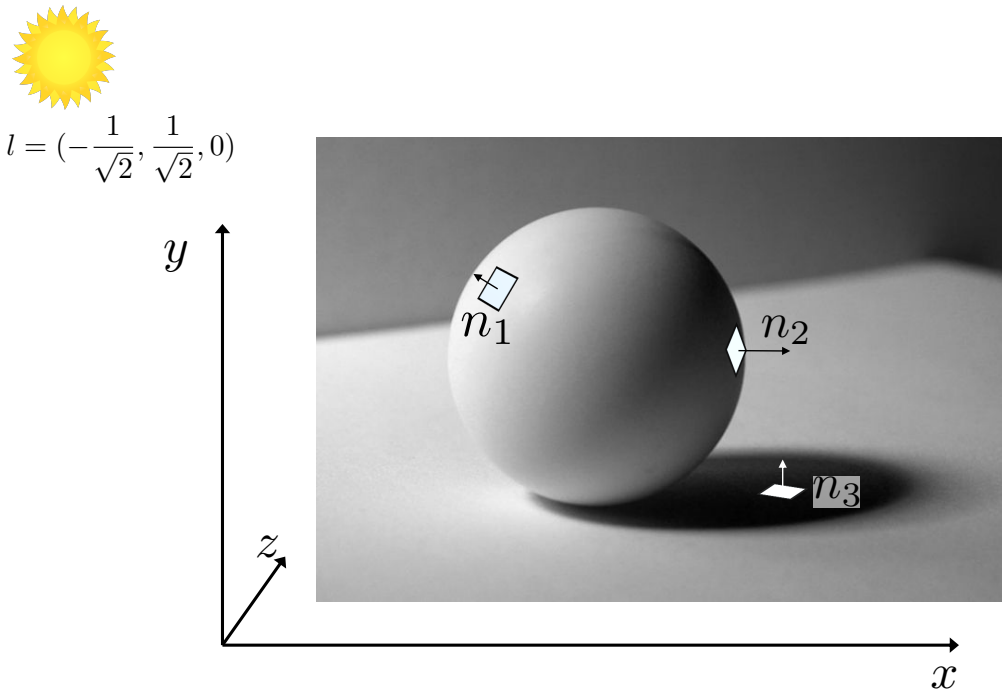

3c.ii. (2 points) Does Lambertian shading give values that agree with this image? If not why?

4. (Total : 12 points) Radiometry, Monte Carlo Integration, Physically-Based Rendering

(4a) (2 points) Solid Angle

A square of side length 2 meters is at a distance of 100 meters from you. What solid angle does it subtend from your perspective? Note that you can assume that the width of the subtended object is much smaller than the distance to the object.

Answer: _____

(4b) (4 points) An isotropic point light at $(0, 3, 4)$ with radiant flux (power) of 100 watts illuminates a surface location at the origin with surface normal vector $(0, 0, 1)$. What is the irradiance at this point?

(4c) Path Tracing

You model and path trace the scene shown in the photo above. Consider altering the path tracing algorithm so that it stops recursion after just 1 or 2 bounces of light. In the questions below you can assume that a lot of samples are drawn so that Monte Carlo image noise is low. In this photo, assume that the light is a diffuse area light source on the ceiling (there is a black rectangular card to block direct light going into the camera's lens.)

4c.i. (2 points) Notice the soft edges on the shadows cast on the floor. Will the rendered image with only 1 bounce of light show these soft shadow edges? Explain your answer briefly.

4c.ii. (2 points) Notice the "color bleeding" of the red wall onto the floor on the left, and from the green wall onto the floor on the right. Will the rendered image with only 1 bounce of light show this color bleeding? How about the rendered image with only 2 bounces of light? Explain your answer briefly.

4c.iii. (2 points) Briefly describe the main visual difference between the rendered image with only 2 bounces of light as compared to the image rendered by regular path tracing.

5. (Total : 24 points) Imaging: Cameras, Lenses, Sensors, Image Processing

(5a) Light Field Camera Design

Ms. B has built a plenoptic-sytle light field camera, with design parameters A, F, a, f and d_{sensor} (see diagram below). She is wondering how to set the gap between the microlens array and the sensor, d_{sensor}. As shown on the illustration below, the microlens array is placed at the focal plane of the main lens (with aperture diameter A and focal length F). The microlenses have an aperture of size a and focal length of f (not shown on the diagram), with the sensor at a distance of d_{sensor} from the microlenses. Note that the illustration is not to scale, and in general the microlenses are very small compared to the main lens. [Trivia: the microlens sensor parameter values below match those found in a Lytro light field camera.]

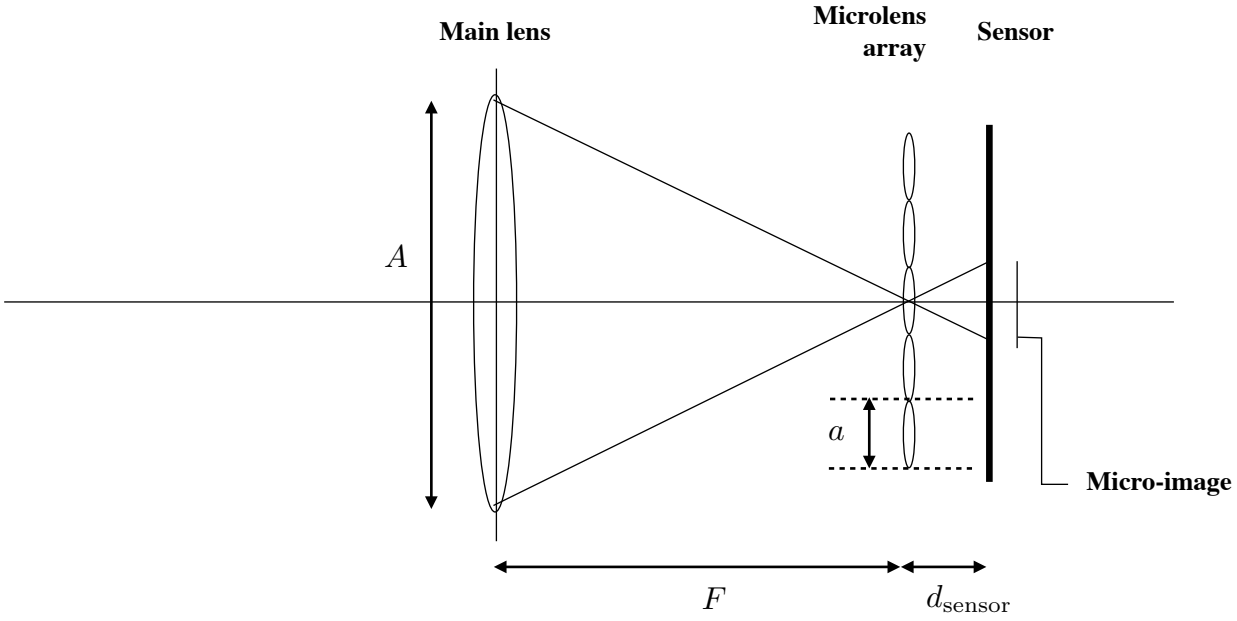

5a.i. (3 points) Recall that in this type of light field camera, each microlens acts like a tiny camera, and creates a micro-image of the main lens aperture onto the sensor. Ms. B realizes that to make as sharp an image as possible, she should set d_{sensor} such that the tiny cameras are focused on the main lens. Given this focusing goal, derive an expression for d_{sensor} in terms of the other camera parameters. Make sure to explain your work for partial credit.

5a.ii. (1 point) In reality, $F \gg f$. In Ms. B's camera, $F = 50$ mm, and $f = 0.04$ mm (40 microns). Given these numbers, what is the value of d_{sensor} (you can give it to three significant figures of precision)?

5a.iii. (4 points) Ms. B is worried about how accurately d_{sensor} needs to be set in the camera she is building. If the real gap in her camera differs by an error Δd from the target gap d_{sensor}, this will misfocus the micro-images appearing under each microlens. The misfocused micro-image of the lens aperture is blurred by a circle of confusion with a diameter C. As a step towards determining how accurately d_{sensor} needs to be set, derive an expression for this diameter, C, in terms of the camera design parameters and Δd. Make sure to explain your work for partial credit.

5a.iv. (2 points) Ms. B decides that this misfocus must have a circle of confusion diameter

that is no larger than the width of one sensor pixel. In her camera, the width of one sensor pixel $\Delta p = 0.00014$ mm (1.4 microns), and $a = 0.02$ mm (20 microns). Given these numbers, continuing the previous part, solve for the allowable error Δd in the gap between the microlens array and the sensor (you can give it to three significant figures of precision)? Sbow your work for partial credit.

(5b) Exposure

You are designing an auto-exposure control algorithm for a camera. Assume that your algorithm can change only three variables: the aperture size, shutter duration and ISO. You can choose an aperture size of exactly one of the following values: f/2, f/4, f/8, f/16 or f/32. You can choose any shutter duration that is longer than 1/4000 seconds, which is the physical limit of the mechanical shutter. You can choose an ISO of exactly one of the following values: 100, 200, 400, 800 or 1600.

5b.i. (4 points) Consider a shot in extremely bright light, where the correct exposure is equivalent to an aperture of f/8 with a shutter duration of 1/1000 seconds at ISO 100. The scene detection algorithm says the next shot is a portrait of a person, so aesthetically your algorithm should blur the background as much as possible. What camera settings should your auto-exposure control algorithm choose? Briefly justify your answer.

Aperture size: _____

Shutter speed: _____

ISO: _____

Explanation:

5b.ii. (4 points) The next shot is in dim light in the evening, where the correct exposure is equivalent to an aperture of f/8 with a shutter duration of 1/8 second at ISO 100. The scene detection algorithm says this shot is a landscape, so aesthetically your algorithm should maximize the depth of field without changing the exposure.
However, image blur due to hand-shake is a concern. Recall that a common rule of thumb in photography is that the shutter duration should be at least as fast as the inverse of the focal length of the lens (35 mm equivalent) to avoid motion blur due to hand-shake. The camera here has a 30 mm focal length lens, so you should make sure to choose a shutter speed less than or equal to 1/30 second.
What lens and shutter settings do you choose? Justify your answer.

Aperture size: _____

Shutter speed: _____

ISO: _____

Explanation:

(5c) (6 points) Exposure and Noise

We have a camera with a 16MP sensor where the pixels have 100% quantum efficiency and zero sensor-related noise (i.e. the only noise present in the signal is due to Poisson shot noise). We take several photos of a patch of a gray wall that is uniformly illuminated. The pixels do not saturate in any of the photos. The camera has a 50 mm lens and a sensor with area 10x10mm containing 16 million pixels. For the photos below, we use the aperture and shutter duration described, and we downsample the resulting photographs to the resolution listed. If the individual pixels in the first photo have an SNR as stated below, calculate the SNR for the individual pixels in the other two photographs. Show your work.

Photo 1: f/2.0 lens aperture, 1/100 second shutter duration, and downsampled to a 4 million pixel image.

SNR = 100

(i) Photo 2: f/1.0 lens aperture, 1/50 second shutter duration, and downsampled to a 1 million pixel image.

SNR = _____

(ii) Photo 3: f/4.0 lens aperture, 1/200 second shutter duration, and not downsampled (i.e. a 16 million pixel image).

SNR = _____

6. (Total : 18 points) Animation: Kinematics, Motion Capture, Physical Simulation

(6a) Forward Kinematics

Shown below is a two bone kinematic chain.

On the left are the bone coordinate systems for each bone. Transformations T_1 and T_2 take each bone to their rest pose, shown on the right.

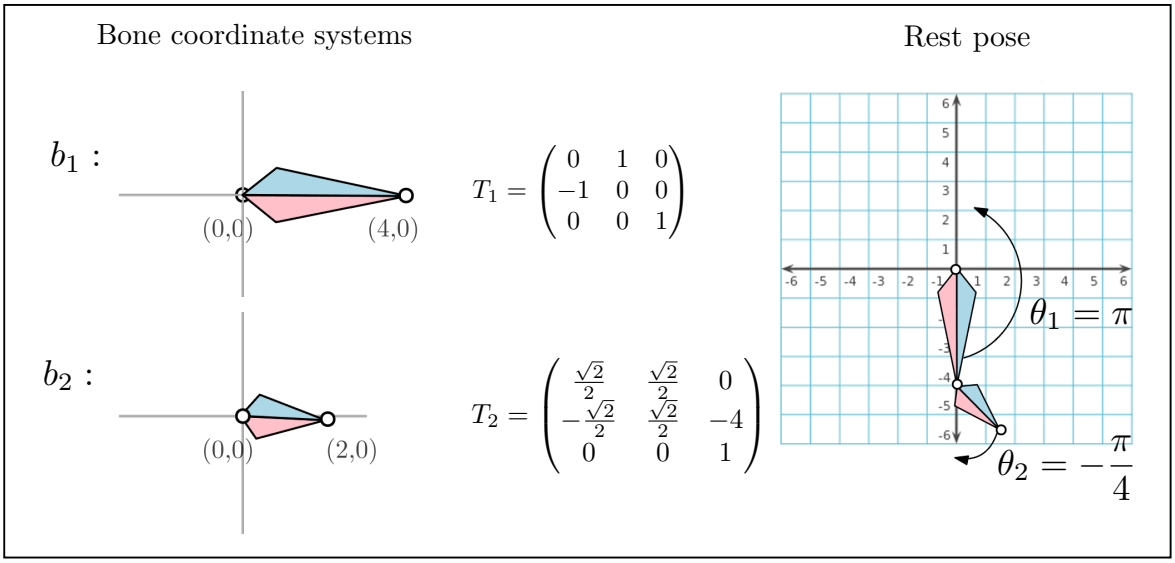

An artist tells you that they want to apply a relative rotation to each bone: $\theta_1 = \pi$ and $\theta_2 = -\pi/4$.

6a.i. (4 points) Draw the configuration of the bones after applying the relative rotation below:

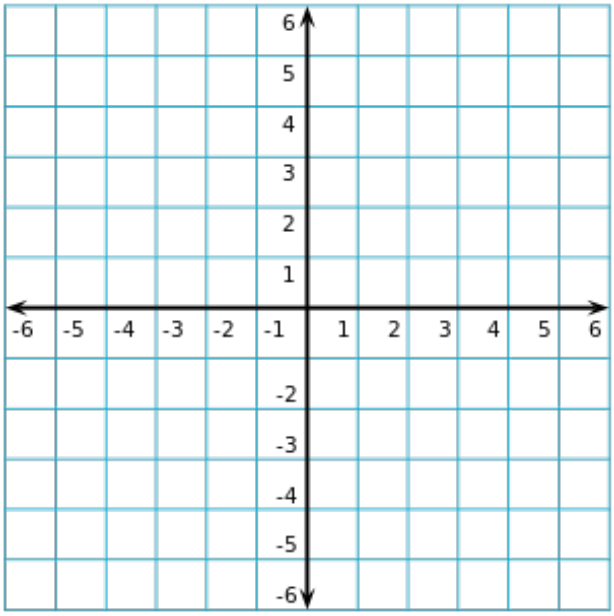

6a.ii. (8 points) Write down the matrix expression for the absolute transformation applied to each bone from rest pose in order to get the configuration above. Let $\hat{T}_1$ and $\hat{T}_2$ be the absolute transformations. You can use R_θ to express a 2D rotation matrix with rotation θ, which you can specify. You can write your solutions in terms of these

variables along with T_1 and T_2.

(6b) (2 points) Great news! We were able to obtain 3D scans for 100 different Dachshunds, all nicely processed so that each dog is a mesh with 1000 vertices, with the same set of faces with vertex-to-vertex correspondences (i.e. vertex #8 is always the vertex on the left eye across all dogs). Here are some samples:

If we compute the shape blend shapes of Dachshunds from this data using PCA, what effects would the first principal component (eigenvector corresponding to the largest eigenvalue) most likely capture?

(i) the fur color
(ii) the leg lengths
(iii) the shoulder widths
(iv) the torso length

(6c) (4 points) Suppose we have a stable spring system hung from the origin (0,0) with the particle, x, at $(-5, 0)$ with velocity $\dot{x} = (0, 0)$, mass $m = 1$ and no internal or external force applied to it.

Suddenly, another particle hits particle x. At the moment of contact, the state of the particle is $x = (-5, 0)$, with velocity $\dot{x} = (0, -2)$.

What is the internal damping force at this moment, with a stiffness factor $k_d = 1$? Briefly explain your answer.

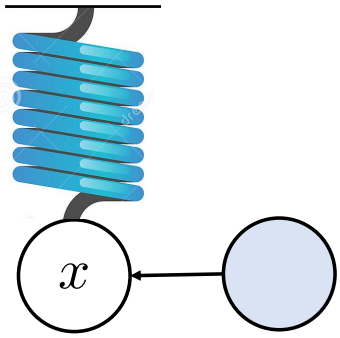

7. (Total : 22 points) Color

For this part of the exam, assume that r_S, r_M and r_L are vectors representing the spectral response of the S, M and L cone cells of the human retina as described in lecture.

(7a) Metamerism

Consider two spectral power distributions, $s_1(\lambda)$ and $s_2(\lambda)$ that are metamers. Under the same viewing conditions, s_1 and s_2 appear the same color to a human observer.

7a.i. (3 points) s_1 and s_2 need not carry the same total power, integrated over all visible wavelengths. Explain why, and give an example.

7a.ii. (3 points) s_1 and s_2 will also appear the same to a color-blind human observer who has only two types of cone cells instead of three. Prove why this is true.

(7b) Color Reproduction Across Displays

In this multi-part question we will consider the problem of reproducing color across displays. Specifically, consider two color displays, where we can set the color by directly choosing the intensity of the R, G and B pixels. Unfortunately, the color displays are from different manufacturers, and colors do NOT look the same across the two displays when we set the same R, G and B intensities for both.

You measure the emission spectra of each of the R, G, and B pixels on each display. Assume that s_{R1}, s_{G1}, s_{B1} are vectors representing the emission spectra for the first display, and that s_{R2}, s_{G2}, s_{B2} are vectors representing the different emission spectra for the second display.

Using these emission spectra, we will derive equations for setting the color on the second display to match the appearance on the first display, by solving for the necessary scalar intensity values R_2, G_2, B_2 for the second display, as a function of the values R_1, G_1, B_1 used on the first display.

7b.i. (2 points) Spectral Power Distributions Emitted by Displays

First, write down expressions for the spectral power distributions emitted by each display, as a function of R_1, G_1, B_1 and R_2, G_2, B_2. You may use any of the variables defined above. Explain your work for partial credit.

7b.ii. (5 points) Tristimulus Response to Colors Presented on Each Display

Now, derive expressions for the tristimulus responses (i.e. the L, M and S) responses of a human observer to the colors presented on each display, as a function of R_1, G_1, B_1 and R_2, G_2, B_2. Put your final answer in the form of a single matrix multiplied by column vectors $\begin{bmatrix} R_1 & G_1 & B_1 \end{bmatrix}^T$ for the first display, and similarly for the second display. You may use any of the variables defined above in your answer. Show your work for partial credit.

$$\begin{bmatrix} S_1 \\ M_1 \\ L_1 \end{bmatrix} =$$

$$\begin{bmatrix} S_2 \\ M_2 \\ L_2 \end{bmatrix} =$$

7b.iii. (5 points) Color Reproduction Across Displays

Now, solve for intensities $\begin{bmatrix} R_2 & G_2 & B_2 \end{bmatrix}^T$ for the second display, as a function of $\begin{bmatrix} R_1 & G_1 & B_1 \end{bmatrix}^T$. You can leave your answer as matrix multiplications but write out the matrices fully. Show your work for partial credit.

$$\begin{bmatrix} R_2 \\ G_2 \\ B_2 \end{bmatrix} =$$

7b.iv. (4 points) Gamut Problems
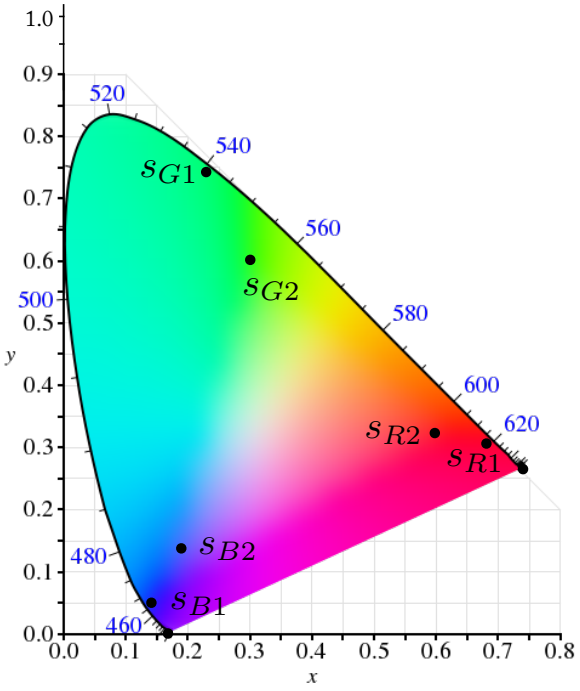
The solution from the previous part works well for many colors, but there is a problem. For some input R_1, G_1, B_1 values, the calculated R_2, G_2, B_2 values contain a negative number. The following chromaticity diagram shows the xy locations of the emission spectra of the R, G, and B pixels for the two displays. Describe or draw the region of the chromaticity diagram containing colors that will exhibit the described problem.