CS 184/284A Spring 2022 **Name:** _____

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

April 21, 2021

Time Limit: 120 minutes

- 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	14	
4	15	
5	22	
6	16	
7	25	
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) ___ Sampling at more than twice the highest frequency in a signal will avoid aliasing.

(2b) (1 point) ___ In Cubic Hermite interpolation, the inputs are the positions and tangent vectors of the curve at the beginning and end of the curve segment.

(2c) (1 point) ___ The Blinn-Phong shading model is not a physically accurate model of surface reflection.

(2d) (1 point) ___ Importance sampling can actually increase variance in our Monte Carlo integration estimate, contrary to our goals, if we sample from a probability density function that is a poor approximation for the integrand.

(2e) (1 point) ___ Two lights of the same luminance can emit spectral power distributions with different total radiometric power.

(2f) (1 point) ___ In physically-based rendering, the global illumination solution contains the lighting contributions from all bounces of light from zero to infinity.

(2g) (1 point) ___ In electronic rolling shutter, visual artifacts are a result of the fact that different pixels are exposed to light for different durations of time.

(2h) (1 point) ___ In a dolly-zoom that maintains the size of an actor's face, the camera is continuously moved closer to the actor while the lens is continuously zoomed to a longer focal length.

(2i) (1 point) ___ We can derive the Thin Lens Equation from the Gauss Ray Diagram by applying similar triangles analysis and algebra.

(2j) (1 point) ___ In taking a photo where the SNR is shot limited, if I double the ISO and halve the shutter duration, SNR will go down by a factor of $\sqrt{2}$.

(2k) (1 point) ___ In JPEG compression, we save bits by reducing the resolution of the luma channel.

(2l) (1 point) ___ A 4D light field camera samples the 2D directional distribution of light arriving at every 2D position on the sensor.

(2m) (1 point) ___ In animation, parameter values between keyframes are usually computed using linear interpolation.

(2n) (1 point) ___ In particle simulation, the Forward Euler method is simple but fundamentally unstable.

(2o) (1 point) ___ Chromatic adaptation is the property of the human visual system of inferring the color of objects by inferring and factoring out the color of the light illuminating them.

(2p) (1 point) ___ Many automatic white balance (AWB) algorithms work by dividing R, G, B values, respectively, by R_W, G_W, B_W, where the latter are an estimate of the values associated with a white object in the scene.

(2q) (1 point) ___ The human brain infers the relative power of light at every wavelength in the visible spectrum, from the S, M, L values that it receives from the retina.

(2r) (1 point) ___ In the color matching experiment, if the test color is out of gamut, at least one primary light must be transferred to the test light side of the system in order to achieve a color match.

3. (Total : 14 points) Graphics Pipeline and Geometry

(3a) (6 points) Transforms

Consider the following input image:

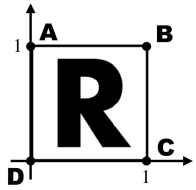

Draw the output image after applying the following transform:

$$\begin{bmatrix} -4 & 2 & 3 \\ -1 & -5 & 3 \\ 0 & 0 & 1 \end{bmatrix}$$

Draw this grid with the output image on your own paper. Please label the coordinates of points A, B, C, and D as well. Make sure you get the orientation of the R correct in your drawing, but the details of the font are not important.

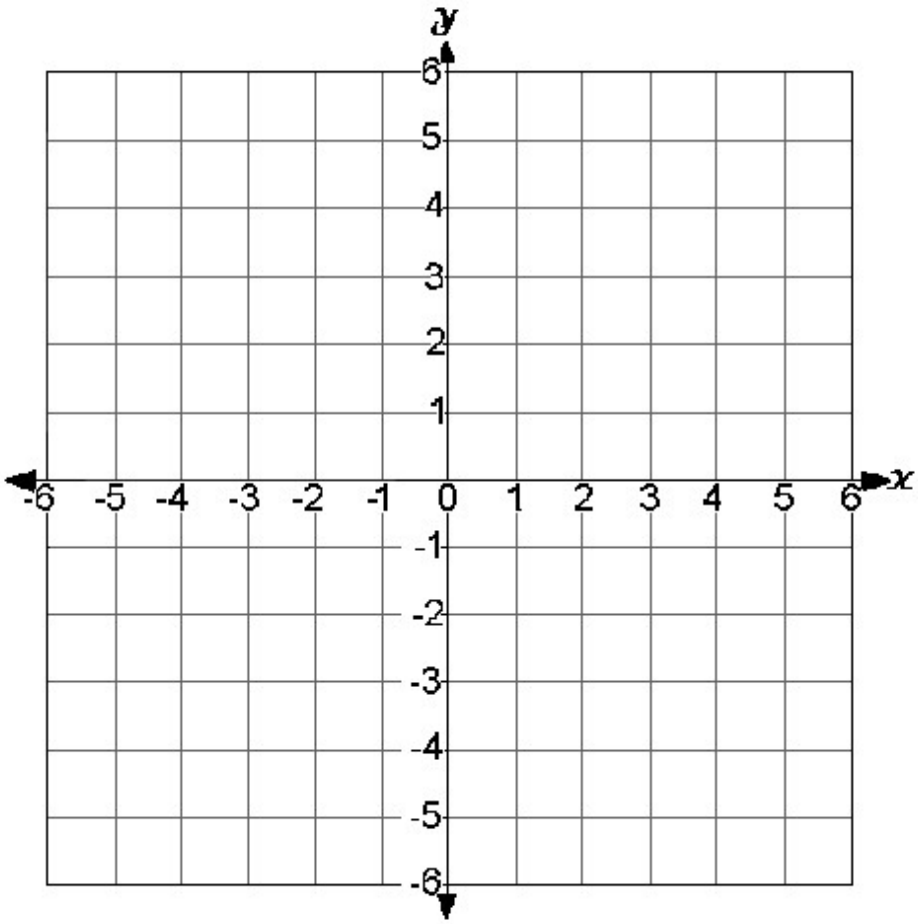

(3b) (2 points) Which of these statements are true? Select all that apply.

In mipmapped texture mapping:

(i) many copies of the texture image are stored at different resolutions.

(ii) the mipmap level is computed at every rasterization sample

(iii) a portion of the rendered scene where a texture is more minified (i.e. the texture appears on a smaller area of screen pixels) will sample from a higher resolution mipmap level

(iv) Mipmap texture mapping will cause less overblurring compared to supersampled rendering without mipmaps.

(3c) (6 points) Consider the following initial (left) and final (right) half-edge meshes.

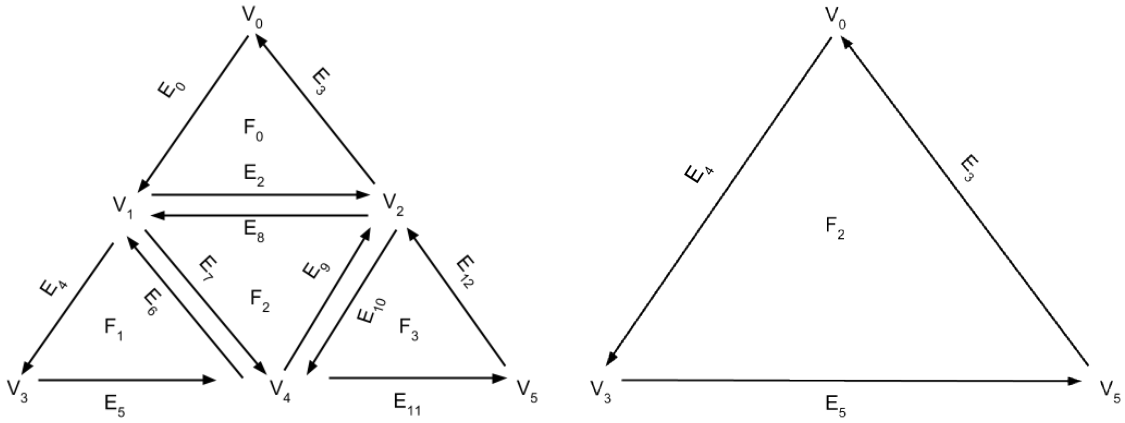

For this problem, please use `halfedge()`, `vertex()`, `face()`, `twin()`, `next()` as the calls to both get and set those items. For example, given a halfedge pointer $\texttt{E}_2$, you could set its twin to the halfedge belonging to some vertex $\texttt{V}_1$ by writing:
$\texttt{E}_2 \rightarrow \texttt{twin}() = \texttt{V}_1 \rightarrow \texttt{halfedge}()$ (this example is not based on the above diagrams)

You are trying to modify the mesh on the left to look like the one on the right. You are initially only given a pointer to the halfedge $\texttt{E}_3$. How can you correctly assign $\texttt{E}_3$'s new vertex, using the fewest number of calls? (Recall that a halfedge's vertex is at the head of the halfedge – e.g. in the mesh on the left, $\texttt{E}_3 \rightarrow \texttt{vertex}()$ is $\texttt{V}_2$.) Assume you cannot access any elements that are not shown. Write your solution in one line. Explain your work for partial credit.

4. (Total : 15 points) Radiometry, Monte Carlo Integration, Physically-Based Rendering

(4a) (3 points) Assume ω_1 and ω_2 are the solid angles subtended by two spheres C_1 and C_2, respectively, from a given point P. Mark the following statements true or false.

___ If $\omega_1 > \omega_2$ then C_1 must be larger than C_2.

___ If the radius of C_1 is greater than C_2, and if C_1 is closer to P, then it must be true that $\omega_1 > \omega_2$.

___ If C_1 and C_2 are the same size and their centers are the same distance from P, then ω_1 and ω_2 must be equal.

(4b) (6 points) A point light at position $(6, 0, 8)$ (in meters) with radiant flux (power) of 100 watts directs all its light uniformly into the hemisphere directly below it. Some of this light falls on a flat, tilted surface passing through the origin, with surface normal vector $(1, 1, 1)$. What is the irradiance at the origin? Show your work and use the correct units.

(4c) (6 points) In path tracing, we discussed two types of importance sampling: sampling over the surface of light sources, or sampling the BRDF in the directions where it is large.

Your friend is considering a point on the surface of a material at which there is a diffuse BRDF and light source that subtends a very small solid angle. To obtain a more accurate image, would it be better for your friend to importance sample the lights or the BRDF, and why?

5. (Total : 22 points) Imaging: Cameras, Lenses, Sensors, Image Processing

(5a) (4 points) Exposure

You are designing an auto-exposure control algorithm for a camera. Assume that your algorithm can change only two variables: the aperture size and the shutter duration. You can choose an aperture size of f/2, f/4, f/8, f/16 or f/32. You can choose any shutter duration from 1/4000 seconds to 4 seconds.

Throughout this question, assume that the correct exposure level for each shot is equivalent to an aperture of f/8 with a shutter duration of 1/60 seconds.

5a.i. 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 without changing the exposure. What lens and shutter settings do you choose?

Aperture size: _____

Shutter speed: _____

5a.ii. The scene detection algorithm says the next shot is a landscape, so aesthetically your algorithm should maximize the depth of field without changing the exposure. What lens and shutter settings do you choose?

Aperture size: _____

Shutter speed: _____

(5b) Comparative Depth of Field in Large and Small Cameras

Your friend comments that her large, standalone camera takes photos with shallow depth of field that she likes. And in contrast, that her smartphone (in regular mode without rendered effects like "portrait mode" that simulate shallow depth of field) takes photos with a large depth of field, where almost everything looks like it is in focus at the same time.

For this problem, assume that the large camera has a sensor size of 25x25mm, and the smartphone has a sensor size of 5x5mm. Assume that both sensors have a resolution of 4000x4000 pixels. Assume that the lenses on the large camera and the smartphone have focal lengths of 25mm and 5mm, respectively, and both have an F-number of F2. Note that these specifications mean that both cameras have the same angular field of view.

5b.i. (6 points) Assume that both the large camera and smartphone are used to take a picture focused on a person's face at a distance of 0.5 meter. In this problem, you will calculate the diameter of the circle of confusion of light coming from the background (i.e. at optical infinity), in sensor pixels, for both the large camera and smartphone. Show your work. You may wish to draw a diagram.

It is recommended that you compute the following values on your way to the final answer.

In the large camera:

- Distance (mm) between sensor and lens
- Diameter (mm) of the circle of confusion on the sensor, for light coming from optical infinity
- That diameter in sensor pixels

In the smartphone camera:

- Distance (mm) between sensor and lens
- Diameter (mm) of the circle of confusion on the sensor, for light coming from optical infinity
- That diameter in sensor pixels

Finally, what is the ratio of these diameters of the circle of confusion for the large camera compared to the smartphone?

5b.ii. (4 points) You point out to your friend that this difference in depth of field and background blur will become less apparent when focused on a subject further away. As an example, consider taking a full-body portrait of a person who is 2 meters tall, such that the resulting photo shows the person as half the full height of the photo. Again, show your work for partial credit.

First, calculate the distance at which the person must be from the camera (lens) to take this photo. (Recall that both cameras have the same angular field of view, so this distance is the same for both cameras.)

_____ meters

Next, assume that you focus both cameras on the person at that distance, and repeat your calculations for the diameter of the circle confusion of light coming from the background (i.e. at optical infinity) in sensor pixels, for both cameras.
Once again, here are some values that might be useful to compute along the way for both cameras.

- Diameter in sensor pixels for large camera
- Diameter in sensor pixels for smartphone
- Ratio of the diameters

Is the ratio less than for the previous question where the camera was focused much closer?

5b.iii. (8 points) Your friend says, let's consider the following thought experiment. Let's say I take a photo A with my large camera. If I took a photo B with exactly the same composition with my smartphone, we know that the depth of field would be larger and the background blur will be less. But if I could magically shrink the whole world by a factor of 5 (the factor by which the smartphone sensor width is smaller than the large camera sensor), I could take a photo C with the smartphone that has the exact same composition as A, and it would have exactly the same depth of field and background blur as A. (Be careful – to achieve the same composition in C as in A, you will have to adjust the distance of the camera to the scene.) Is your friend correct? If not, describe a counterexample scene, and show that the background blur would be different. If your friend is correct, prove it.

6. (Total : 16 points) Animation: Kinematics, Motion Capture, Physical Simulation

(6a) (2 points) Kinematics

Consider the one segment arm of length r attached to a piston as shown in the diagram below. For each x (piston displacement) and θ (arm rotation angle), the end of the arm is at some point p in the 2D plane.

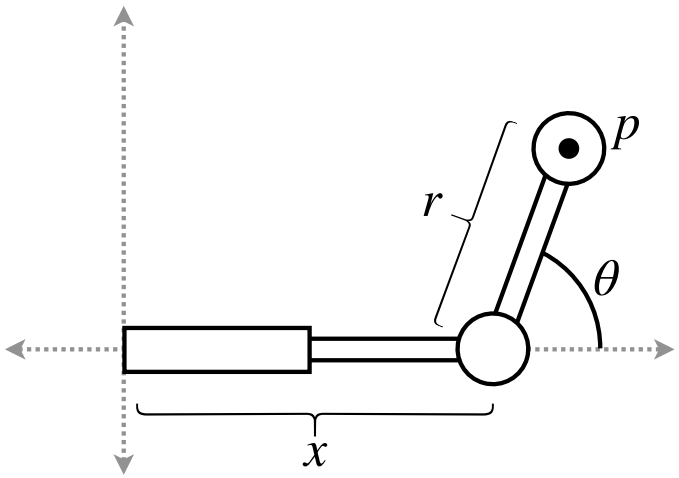

Write down a formula for the position $p = (a, b)$ in the plane as a function of the controllable parameters x and θ.

(6b) (4 points) Now let's consider inverse kinematics of the one segment arm in the previous question. We are given the desired position p and we must solve for x and θ. Mark the following statements true or false.

___ There is always a unique solution.

___ There may be an infinite number of solutions.

___ The set of solutions, if it exists, is continuous in parameter space.

___ There are either zero, one or two solutions.

(6c) (10 points) Euler Steps

Suppose we have a particle of mass $m = 1.0$ with the following initial conditions at time $t = 0$:

$$(x, y) = (0, 2) \qquad (\dot{x}, \dot{y}) = (4, 1)$$

This particle is subject to a force due to an undamped spring attached to the origin at one end and the particle at the other. The spring's position at the origin is fixed, whereas the particle is allowed to move around. The spring constant is $k = 1$, and its rest length is 2.

We wish to find the trajectory of this particle using explicit (non-modified) Euler integration with a timestep of $\Delta t = 1$. **Draw the first three steps of this particle (not including its position at $t = 0$) and write out its final computed position.**

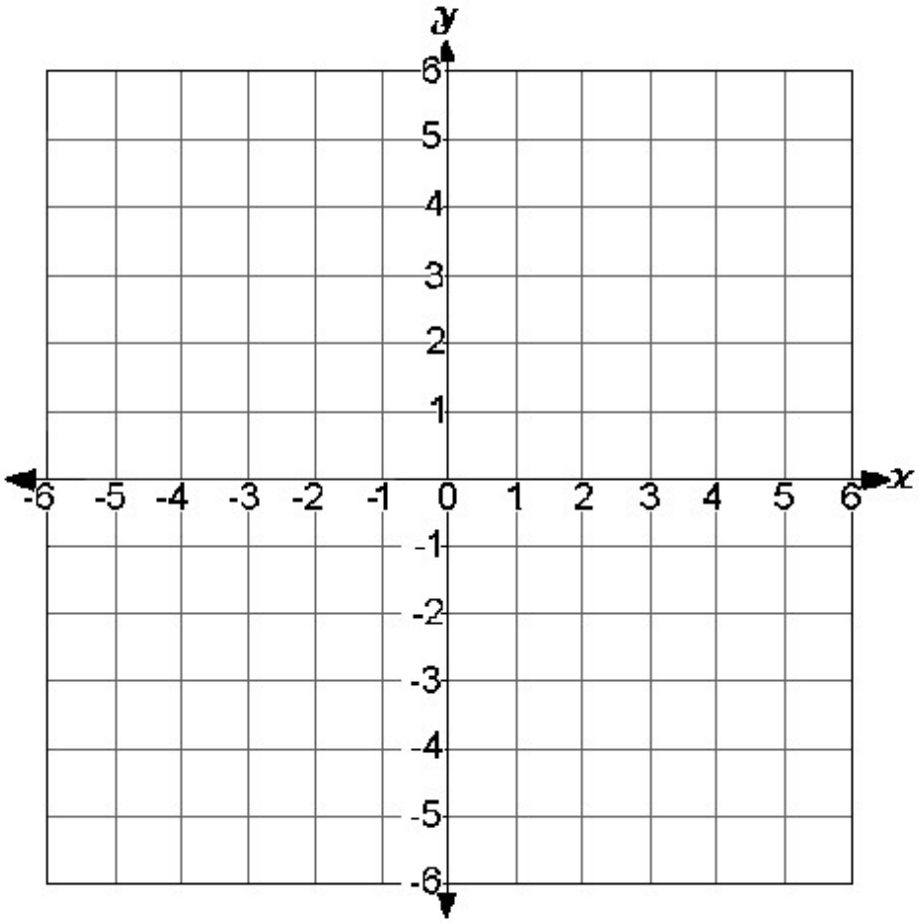

7. (Total : 25 points) Color

(7a) Color Reproduction for Aliens

An alien species magically transports to earth. They are almost exactly like us humans in every way, except that their color perception is different. It turns out that instead of S, M and L cone cells, they have S, P and L and cone cells. The spectral response of the P cone cell is approximately midway between that for S and M cones (similar to one of the cone cells found in pigeons).

One thing that is discovered is that our RGB smartphone displays don't produce convincing colors for aliens. Unlike for humans, for aliens the display colors don't appear the same as the corresponding colors in the real world.

Assume that r_S, r_P, r_M and r_L are vectors representing the spectral response of the S, P, M and L cone cells of the human and alien retinas as described in lecture. Also, assume that s_R, s_G and s_B are the spectral emission functions for the R, G and B pixels in the smartphone displays we are testing.

You are a color display engineer, and you take it upon yourself to reprogram some smartphones with modified color processing, with the goal of producing optimal color for our new alien friends. Given an SPD s of a color in the real world that we want to reproduce for an alien, in the next few subparts of this problem, you will derive new intensities R_A, G_A and B_A for the R, G, and B pixels in the smartphones, to accomplish the goal.

7a.i. (2 points) Spectral Power Distribution Emitted by Reprogrammed Smartphones

First, write down an expression for the SPD emitted by the reprogrammed smartphones, as a function of R_A, G_A and B_A. You may use any of the other variables defined above.

$s_{\text{display}} =$

7a.ii. (4 points) Tristimulus Alien Response

Next, write down expressions for the tristimulus response of the alien visual system to the real SPD s and the displayed SPD s_{display}.

$$\begin{bmatrix} S \\ P \\ L \end{bmatrix}_{\text{real}} =$$

$$\begin{bmatrix} S \\ P \\ L \end{bmatrix}_{\text{display}} =$$

for the s_R, s_G and s_B pixels in the smartphone.

7a.iii. (6 points) Alien Color Reproduction Formula
Finally, for the display to reproduce the color of s for the aliens, the tristimulus values in the previous part must be equal. Set them equal and solve for the intensities R_A, G_A and B_A as a function of the input s.

$$\begin{bmatrix} R_A \\ G_A \\ B_A \end{bmatrix} =$$

7a.iv. (2 points) The reprogrammed smartphones are very popular amongst the aliens, but nobody succeeds in reprogramming the smartphones to achieve exact color reproduction for both humans and aliens at the same time. Give a brief technical explanation of why this cannot be achieved.

7a.v. (3 points) It turns out that there are aliens with color vision deficiency (CVD), similar to humans, where they are effectively missing their L cone cells (they have only S and P cones). In an engineering brainstorming session at your company, someone proposes that you may be able to program an RGB TV's color processing so that it provides exact color appearance for an audience of CVD aliens as well as CVD humans who were missing M cone cells. Do you agree, and if so, briefly how would you achieve this? (It is an engineering meeting, so there are no marketing people to complain that there may not be a large market for CVD aliens and CVD humans who want to watch a movie together.)

7a.vi. (2 points) Continuing the previous part, another engineer suggests that you would have to have different color processing for an audience of the CVD aliens who have S and P cones only, and CVD humans who have S and M cones only (missing L cone cells). Briefly explain why you agree or disagree.

7a.vii. (2 points) Continuing the previous part, another engineer claims that you could never achieve exact color reproduction for an audience of the CVD aliens and CVD humans who are missing S cone cells. Briefly explain why you agree or disagree.

7a.viii. (4 points) The aliens settle into life on earth. You go back to the engineering drawing board to come up with a color display that will reproduce color appearance accurately for all aliens and humans looking at the display at the same time. What is the key change that you need to make in the hardware design of the display to accomplish this goal? And what is the resulting color processing formula you must use to calculate the intensities of the pixels in the new display?