CS 184/284A Spring 2020 **Name:** _____

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

April 23rd, 2020

Time Limit: 110 Minutes

- 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).
- Please read carefully, sign and submit the Honor Code described on the next page.
- This exam is occurring online. You must not communicate with anyone while you are taking this exam, unless there is an emergency.
- This exam is open book.
- This exam is 110 Minutes, and has a total of 110 points. There are 19 pages (including this cover page) and 7 parts.
- Problem difficulty varies throughout the exam, so move on if you find yourself stuck.
- (Improved since Online Exam 1) You will submit your answers for all questions in a single Gradescope "online assignment", titled GRADESCOPE EXAM 2. You will submit multiple-choice and short text answers directly through your browser. For select questions marked UPLOAD PHOTO, you will write your answers by hand on your own paper, take a photo and upload to the relevant questions in GRADESCOPE EXAM 2. You are allowed 15 minutes after the exam is over to complete photographing and uploading photos of your work. If necessary because of uploading problems, you can email your photos to vio@berkeley.edu with subject 'CS184 Exam2'.
- Communications: During the exam we will post any live corrections in a Piazza post, and you may submit questions to staff via an online form. Both can be found at Piazza post @504.

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 used no more than 110 minutes 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.

(1a) 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: _____

(1b) Self-Portrait Photo

UPLOAD PHOTO **Please take a photo of yourself, and upload to Gradescope Exam 2. (*You may decline this task without penalty, and if so, write words to that effect on a piece of paper, and upload a photo of that instead.*)**

2. (Total : 18 points) True / False

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

(2a) (1 point) **T** Supersampling reduces aliasing by blurring out and removing frequencies above the image resolution.

(2b) (1 point) **T** Extraordinary vertices during subdivision will tend to pull new vertices towards their location and smooth the nearby mesh more quickly than regular vertices.

(2c) (1 point) **T** The naive, full project 3-1 BVH intersection tests can be accelerated if all rays are known to go in the -z to +z direction.

(2d) (1 point) **F** Adding an EPS_D shift away from the intersection point of each ray does not result in a biased render result.

(2e) (1 point) **F** Irradiance enables path tracers to work with a quantity of light independent of distance.

(2f) (1 point) **F** Caustics can be easily captured with camera ray-based path tracing as long as light importance sampling is used at the end of the rays.

(2g) (1 point) **T** An object at the focal point of the camera will focus at infinity.

(2h) (1 point) **T** An ideal pinhole camera will not exhibit any circles of confusion.

(2i) (1 point) **F** An change in f-stop from f/4.0 to f/2.0 results in twice as much light entering the lens.

(2j) (1 point) **T** Long exposure times can help mitigate the effects of shot noise.

(2k) (1 point) **F** Inverse kinematics is the preferred animation approach when requiring precise, fine-tuned motion.

(2l) (1 point) **F** There exist no numerical integration methods that are both unconditionally stable and physically accurate.

(2m) (1 point) **T** Inaccuracies in the results of fluid simulation and other dense matter physics are partially due to the intractability of simulating global particle interactions.

(2n) (1 point) **T** Adaptive step sizes would be well suited for simulating sudden jerks and other high frequency forces.

(2o) (1 point) **F** For a color, perceived hue varies linearly with changes in R, G, and B contributions.

(2p) (1 point) **F** One can get around the problem of "negative" color contributions in a 3-color gamut by exploiting the opponent color theory.

(2q) (1 point) **T** Having four distinct cones would theoretically reduce the number of perceived metamers.

(2r) (1 point) **T** Substituting the Red channel in an LED display with a metamer will not change the resulting colors displayed.

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

(3a) (6 points) Transforms

Consider the following input image:

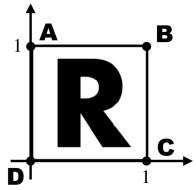

Draw the output image after applying the following transform:

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

Draw this grid with the output image on your own paper. Please label the coordinates of points A, B, C, and D as well.

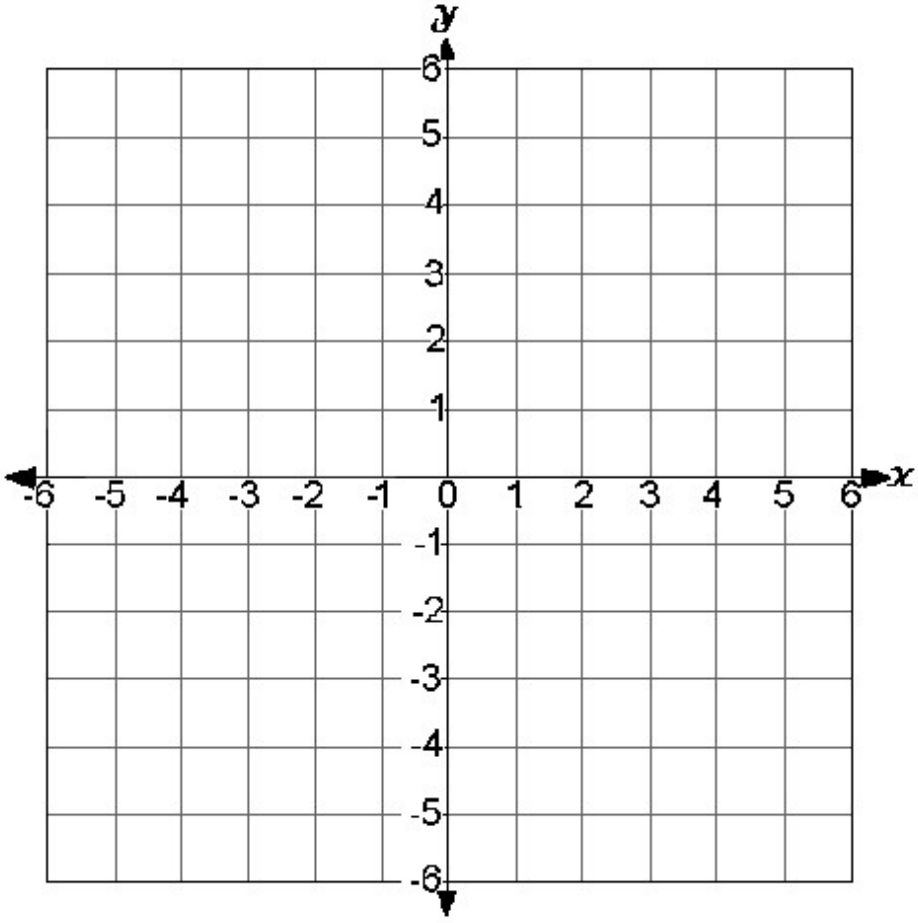

Solution:

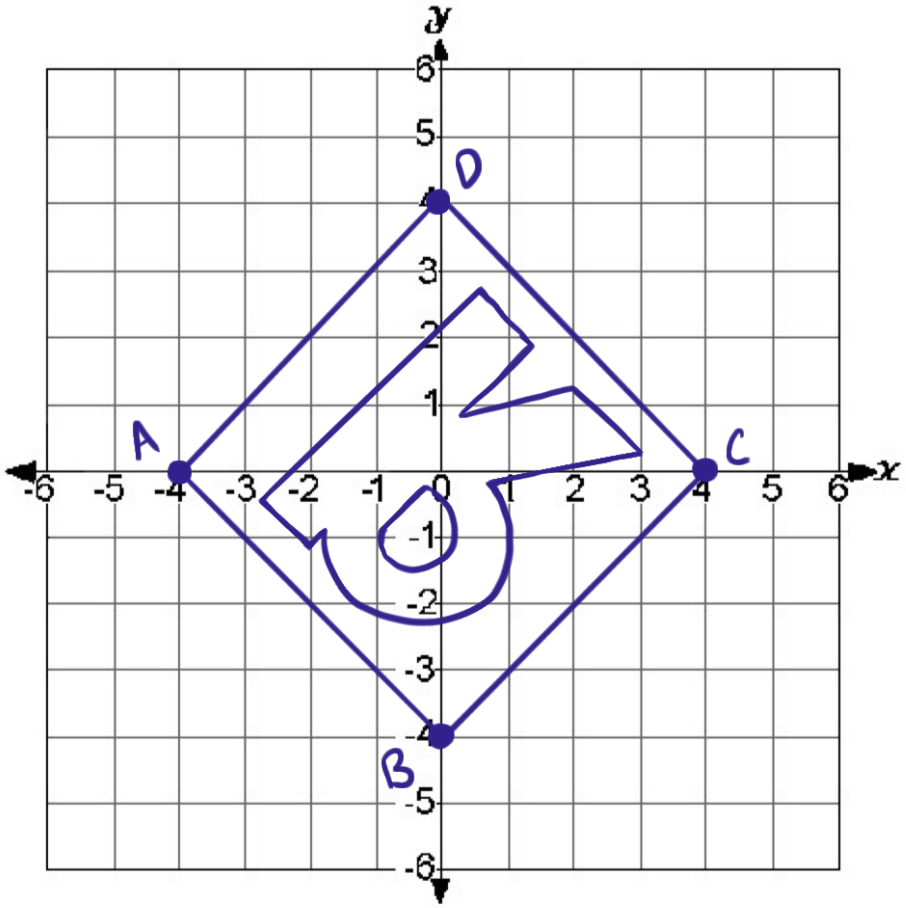

(3b) (6 points) Derive a quadratic polynomial $f(t)$ that satisfies $f(0) = \mathbf{p}_0$, $f'(1) = \mathbf{p}_2 - \mathbf{p}_1$, and $f(1) = \mathbf{p}_2$. Your answer should be in terms of t, $\mathbf{p}_0$, $\mathbf{p}_1$, and $\mathbf{p}_2$.

Answer: $f(t) =$_____

Derivation:

Solution: $f(t) = at^2 + bt + c$. The three constraints give the following matrix:

$$\begin{pmatrix} 0 & 0 & 1 \\ 2 & 1 & 0 \\ 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} \mathbf{p}_0 \\ (\mathbf{p}_2 - \mathbf{p}_1) \\ \mathbf{p}_2 \end{pmatrix}$$

Which means:

$$\begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} 1 & 1 & -1 \\ -2 & -1 & 2 \\ 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} \mathbf{p}_0 \\ (\mathbf{p}_2 - \mathbf{p}_1) \\ \mathbf{p}_2 \end{pmatrix}$$

Solving for a, b, c: $f(t) = (\mathbf{p}_0 + (\mathbf{p}_2 - \mathbf{p}_1) - \mathbf{p}_2)t^2 + (-2\mathbf{p}_0 - (\mathbf{p}_2 - \mathbf{p}_1) + 2\mathbf{p}_2)t + \mathbf{p}_0$

or, simplified: $f(t) = (\mathbf{p}_0 - \mathbf{p}_1)t^2 + (-2\mathbf{p}_0 - \mathbf{p}_1 + \mathbf{p}_2)t + \mathbf{p}_0$

Or, you could do this more simply using systems of equations.
We have 3 unknowns (well, one is pretty trivial...), and can come up with 3 equations to solve for them:
$c = \mathbf{p}_0$
$2a + b = (\mathbf{p}_2 - \mathbf{p}_1)$
$a + b + c = \mathbf{p}_2$
$\Rightarrow a + b = \mathbf{p}_2 - \mathbf{p}_0$

Subtracting the second equation by the third gives
$a = (\mathbf{p}_2 - \mathbf{p}_1) - (\mathbf{p}_2 - \mathbf{p}_0) = \mathbf{p}_0 - \mathbf{p}_1$
Now we just plug in for b, giving:
$b = \mathbf{p}_2 - \mathbf{p}_0 - a = \mathbf{p}_2 - \mathbf{p}_0 - (\mathbf{p}_0 - \mathbf{p}_1) = -2\mathbf{p}_0 + \mathbf{p}_1 + \mathbf{p}_2$
So we have $f(t) = (\mathbf{p}_0 - \mathbf{p}_1)t^2 + (-2\mathbf{p}_0 + \mathbf{p}_1 + \mathbf{p}_2)t + \mathbf{p}_0$

(3c) (2 points) Computational Complexity

Consider a piece of triangle mesh where all the vertices have valence 6 or less.

3c.i. (1 point) You first store the mesh as a simple list of triangles. Given a pointer to one triangle, what is the computational complexity of finding its three neighboring triangles? Assume there are n triangles.

Answer: _____

Solution: $O(n)$ for the list of triangles

3c.ii. (1 point) Next, you decide to store the mesh in a Half Edge data structure. Now what is the computational complexity of finding the three neighboring triangles?

Answer: _____

Solution: $O(1)$ for the Half Edge datastructure.

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

(4a) (2 points) A sphere of radius 10 cm is at a distance of 10 meters from you. What solid angle does it subtend from your perspective? Note that we can assume that the radius of the subtended object is much smaller than the distance to the object.

Answer: _____

Solution:
It subtends a solid angle of $\pi(0.1)^2/10^2 = \pi/10000$.

(4b) (4 points) Importance Sampling Bug

You add importance sampling to a portion of your Monte Carlo renderer. The noise in the resulting image goes down for the same number of samples, but the image gets brighter. You're sure the implementation was correct before adding importance sampling. Write one or two sentences about what might have happened.

Solution: The fact that noise went down is good, and indicates that the importance sampling function was a good guess for the underlying integrand. However, the fact that the image got brighter indicates that the normalization was incorrect – one explanation is that you forgot to divide each sample's value by the probability that you drew that sample.

(4c) (4 points) Path Tracing

What would happen if we were not allowed to use Russian Roulette in the path tracing algorithm?

Solution: Either (1) we would have an infinite recursion bug, or (2) we might have a fixed maximum number of bounces, say N, which would be biased and omit lighting contributions involving more than N bounces of light from light sources to the camera. The second answer is enough to get full points on this subpart.

(4d) (2 points) Radiometry vs Photometry

You shine two light sources on a surface, one at a time, and measure the irradiance and illuminance. The irradiance from the second light is twice that of the first, but the illuminance is the same. What does this mean, and what is the explanation?

Solution:
The irradiance being double means that the second light is putting twice as much energy onto the surface per unit time. The illuminance being the same means that the perceived brightness of the surface by a human observer is the same. One explanation, not unique, is that the second light emits a significant amount of energy at wavelengths outside the visible spectrum.

5. (Total : 24 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 1/30 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: _____

Solution:
Aperture size: f/2 blurs the background the most.
Shutter speed: $1/(30 * 16) = 1/480$ seconds. 1/500 seconds also accepted.
Opening the aperture by 1 f-stop doubles the exposure, and there are 4 f-stops from f/8 to f/2 (f/5.6, f/4, f/2.8, f/2). This means that opening the aperture to f/2 gives us $2 * 2 * 2 * 2 = 16$ times the original exposure. To compensate, we decrease the shutter speed by a factor of 1/16 to keep exposure the same.

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: _____

Solution:
Aperture size: f/32 maximizes depth of field, since it's the smallest possible aperture opening.
Shutter speed: $16 * 1/30 = 8/15$ seconds. 1/2 seconds also accepted.
Closing the aperture by 1 f-stop halves the exposure, and there are 4 f-stops from f/8 to f/32 (f/11, f/16, f/22, f/32). This means that decreasing the aperture to f/32 gives us 1/16 times the original exposure. To compensate, we increase the shutter speed by a factor of 16 to keep exposure the same.

(5b) (6 points) Sensor Noise

You continue to work on your auto-exposure control algorithm for the camera. You are concerned that for shots with small aperture, the shutter duration can be too long.

5b.i. What is the image quality problem that can occur if the shutter duration is too long?

Answer: _____

Solution: Motion blur.

5b.ii. In order to decrease the shutter duration while maintaining the same final exposure, you consider increasing the ISO (gain) from ISO 100 to ISO 1600. What is the image quality problem that can occur if you make this change?

Answer: _____

Solution: More noise.

5b.iii. The SNR (signal-to-noise ratio) will change if you increase the SNR from ISO 100 to ISO 1600 as described. Assuming that the SNR is limited by shot noise, which may not be true as described in class, what is the change in SNR? In your answer, use the variable prev_snr to denote the SNR when the ISO was 100.

Answer: _____

Solution: $prev_snr/4$

5b.iv. You make the change in SNR from ISO 100 to ISO 1600 as described, and observe that the resulting SNR is worse than predicted – it is not limited by shot noise. What is a likely source of the problem?

Answer: _____

Solution: Read noise.

(5c) (10 points) Assuming a camera with an ideal thin lens with 50 mm focal length, with an f/4.0 aperture, and which is focused very far away (i.e. at optical infinity). Now consider a tiny firefly at a distance of 1 meter.

5c.i. UPLOAD PHOTO **Draw a ray diagram representing this problem that includes the following items and labels them: tiny firefly, lens, focal point, sensor, rays from the firefly passing through the lens onto the sensor, circle of confusion, distance from lens to firefly z_o, distance from lens to image of firefly z_i, and focal length f. Your drawing should be neat and clear, but does not need to be drawn to scale.**

Solution:

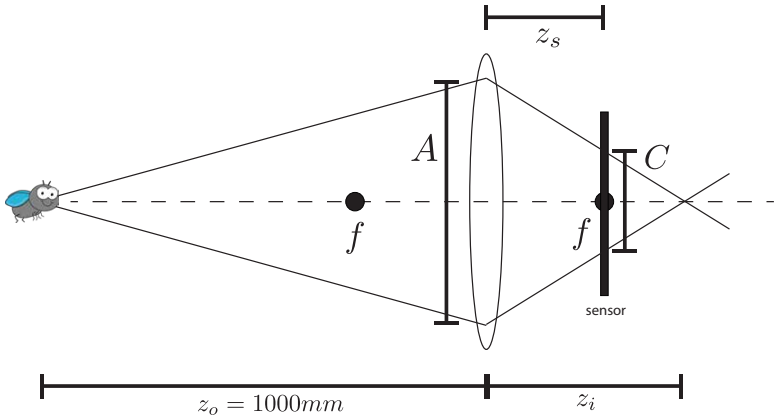

5c.ii. UPLOAD PHOTO **Derive the size of the circle of confusion (in mm) that the firefly makes on the sensor surface.**

Solution:

$$f = 50\text{mm} \tag{1}$$

$$z_o = 1000mm \tag{2}$$

$$\frac{1}{z_i} = \frac{1}{f} - \frac{1}{z_o} \quad \text{thin lens equation} \tag{3}$$

$$\frac{1}{z_s} = \frac{1}{f} - \frac{1}{\infty} = \frac{1}{f} \tag{4}$$

$$C = A \cdot \frac{z_i - z_s}{z_i} \quad \text{from diagram by similar triangles} \tag{5}$$

$$= \frac{f}{4} \cdot \frac{z_i - f}{z_i} \quad \text{by definition of F} - \text{number} \tag{6}$$

$$= \frac{f}{4}\left(1 - \frac{f}{z_i}\right) \tag{7}$$

$$= \frac{f}{4}\left(1 - \left(1 - \frac{f}{z_o}\right)\right) \tag{8}$$

$$= \frac{f^2}{4z_o} \tag{9}$$

$$= \frac{50^2}{4 \cdot 1000}\text{mm} \tag{10}$$

$$= 0.625\text{mm} \tag{11}$$

(5d) (4 points) Image Processing by Convolution

You process an image as follows. First, you convolve each row of pixels by a 1D filter kernel equal to $[1, 2, 1]$. Second, you convolve each column of pixel with the same convolution kernel. The result is equal to convolving by what 2D filter kernel?

Solution:

$$\begin{pmatrix}1\\2\\1\end{pmatrix} * \begin{pmatrix}1 & 2 & 1\end{pmatrix} = \begin{pmatrix}1 & 2 & 1\\2 & 4 & 2\\1 & 2 & 1\end{pmatrix}$$

Explanation:

Try it out with a 3x3 matrix! Convolve the rows of this matrix with the filter horizontally:

$$\begin{pmatrix}a & b & c\\d & e & f\\g & h & i\end{pmatrix}$$

to get this:

$$\begin{pmatrix}a+2b+c\\d+2e+f\\g+2e+h\end{pmatrix}$$

and then convolve that vertically to get: $(a+2b+c)+(2d+4e+2f)+(g+2e+h)$

Looking at this result, we can see that this is equivalent to convolving the entire 3x3 matrix by this 3x3 kernel, which is our answer.

$$\begin{pmatrix}1 & 2 & 1\\2 & 4 & 2\\1 & 2 & 1\end{pmatrix}$$

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

(6a) (10 points) Euler Steps

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

$$(x, y) = (3, 0) \qquad (\dot{x}, \dot{y}) = (-3, 3)$$

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 of this spring is $k = 1$, and its rest length is 3.

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.**

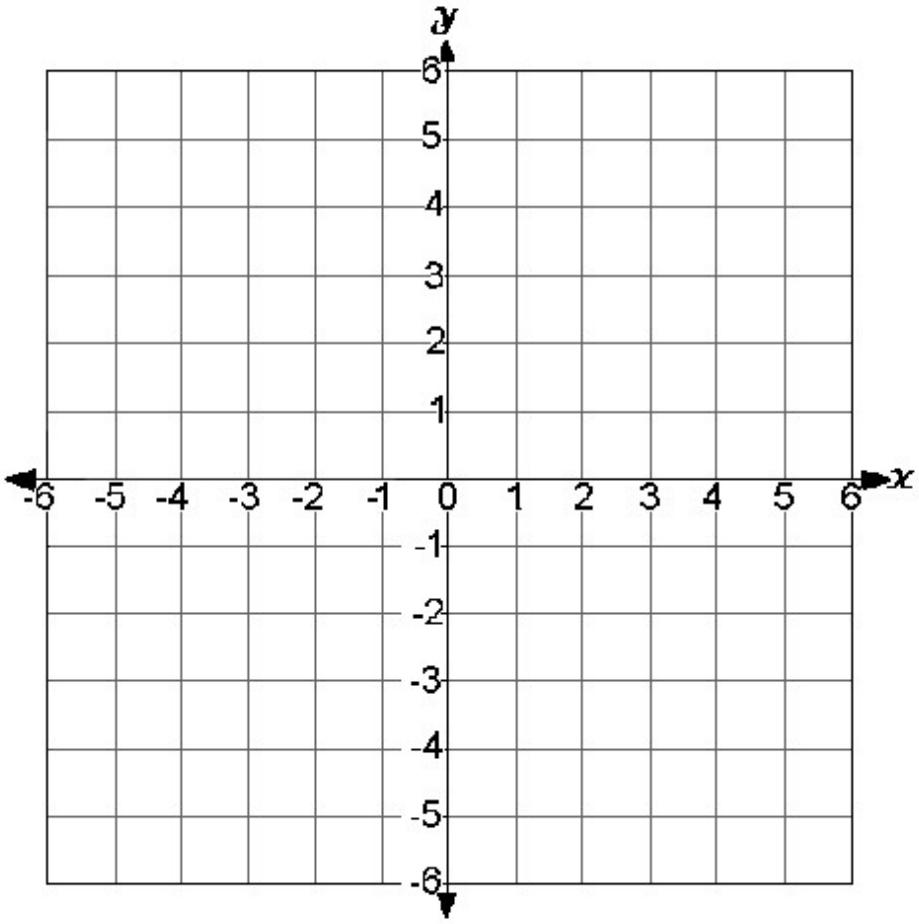

Solution:

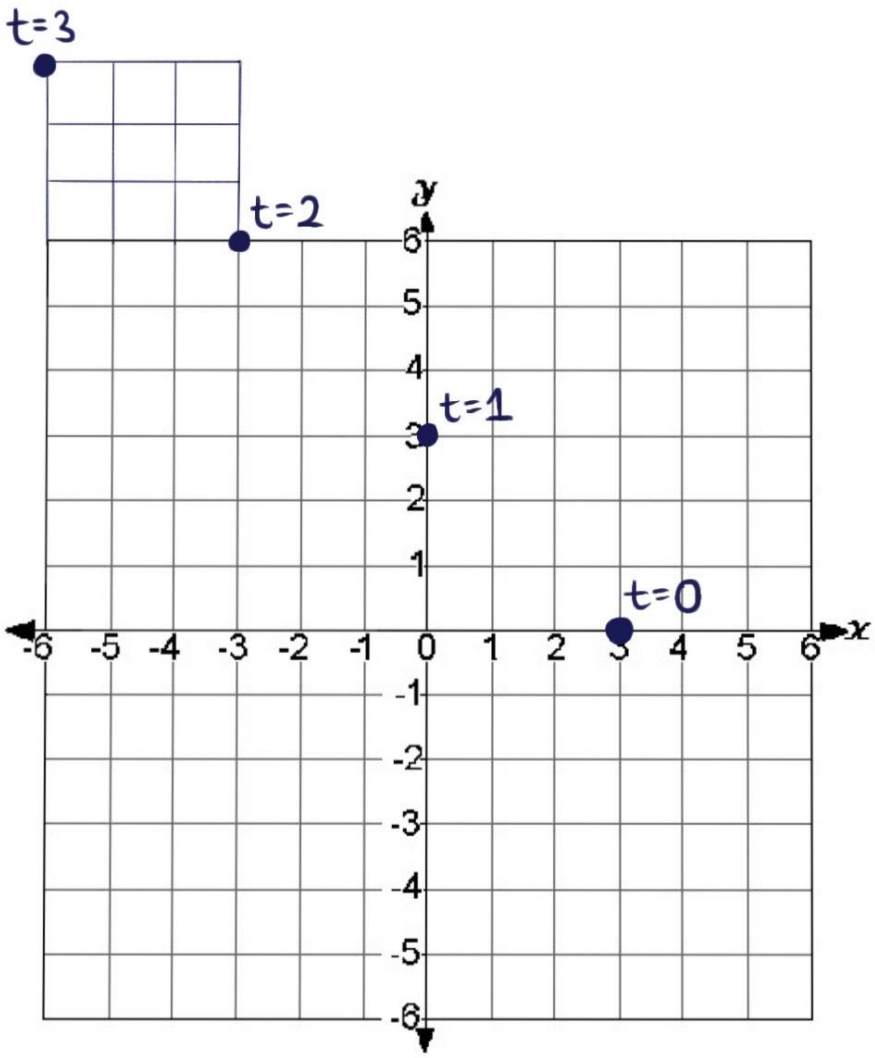

(6b) (6 points) Simulation Stability

Jack wrote a mass-spring simulation program using explicit Euler's method. The program works well in scene A, but blows up after a few iterations in scene B. If Jack's implementation is correct, say whether each of the following modifications could make his program more stable in scene B, and a sentence on why?

(i) Decrease simulation step size.

(ii) Switch to implicit Euler implementation.

(iii) Add some damping to the system.

(iv) Increase stiffness of the springs.

Solution: The answers are :

(i) Yes, because the error per time-step is reduced.

(ii) Yes, because implicit Euler can be unconditionally stable.

(iii) Yes, because damping reduces the energy at each time step.

(iv) No, because higher stiffness of the springs means larger forces at each time step and will increase the rate of divergence.

(6c) (2 points) You are given the output of a particle simulation. In other words, you are given a sequence of particle positions and velocity vectors over time. What method would you use to fill in the positions at intermediate time points, for each particle, to generate a smooth output rendering?

A) Hermite spline interpolation

B) Euler integration

C) Rasterization

D) Mesh traversal

E) Ask your computer nicely

Solution:
A) Hermite spline interpolation.

7. (Total : 22 points) Color

For this section 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) Color Vision Deficiency and Metamerism

A person has color vision deficiency, where the retina contains only two types of cone cells, the S and the M types, and NOT the L type. For this person, two lights with different spectral power distributions (SPDs) appear as identical colors. As we did in class, we can represent the two SPDs by vectors s_1 and s_2. For these two SPDs to appear identical to this person, it must be the case that the following two equations are satisfied:

$$\begin{aligned} r_S \cdot s_1 &= r_S \cdot s_2 \\ r_M \cdot s_1 &= r_M \cdot s_2 \end{aligned} \tag{12}$$

7a.i. (1 point) It turns out that s_1 and s_2 do NOT appear identical for a person with regular, trichromatic color vision. For a person with normal eyesight, write down an inequality using the variable r_L that must be satisfied at the same time as the two equations above.

Solution: $r_L \cdot s_1 \neq r_L \cdot s_2$.

7a.ii. (3 points) You set a goal of modifying either s_1 or s_2 so that the resulting pair appear identical to both people described above. Luckily, you have a third light source with SPD vector s_3, that appears black to the person with color vision deficiency, and NOT black to the person with normal color vision. As a first step, write down three equations or inequalities involving s_3 that must be true.

Solution:

$$\begin{aligned} r_S \cdot s_3 &= 0 \\ r_M \cdot s_3 &= 0 \\ r_L \cdot s_3 &> 0 \end{aligned}$$

7a.iii. (8 points) You realize that you can accomplish your goal by adding some s_3 into one of the other two lights, creating an additive linear combination as we saw in class. Using this clue, specify how to modify either s_1 or s_2, so that the resulting pair of SPDs appear identical to both people described above. Show mathematically that the resulting pair of lights appear identical to both people described above.

Solution: Without loss of generality, assume $r_L \cdot s_1 < r_L \cdot s_2$, such that $r_L \cdot (s_2 - s_1) > 0$. We want to find scalar constant k, such that $s_1' = s_1 + ks_3$ and s_2 appear identical.
For s_1' and s_2 appear identical to both the color vision deficient person and the

normal color vision person, all three of the following equations must be true:

$$\begin{aligned} r_S \cdot s_1' &= r_S \cdot s_2 \\ r_M \cdot s_1' &= r_M \cdot s_2 \\ r_L \cdot s_1' &= r_L \cdot s_2. \end{aligned} \tag{13}$$

Expanding the third equation and solving for k, we find that $s_1' = s_1 + \frac{r_L \cdot (s_2 - s_1)}{r_L \cdot s_3} s_3$. Finally, we verify mathematically that s_1' and s_2 appear identical to both people as follows:

$$\begin{aligned} r_S \cdot s_1' &= r_S \cdot \left(s_1 + \frac{r_L \cdot (s_2 - s_1)}{r_L \cdot s_3} s_3 \right) \\ &= r_S \cdot s_1 + \frac{r_L \cdot (s_2 - s_1)}{r_L \cdot s_3} r_S \cdot s_3 \\ &= r_S \cdot s_1 \\ &= r_S \cdot s_2 \\ r_M \cdot s_1' &= r_M \cdot \left(s_1 + \frac{r_L \cdot (s_2 - s_1)}{r_L \cdot s_3} s_3 \right) \\ &= r_M \cdot s_1 + \frac{r_L \cdot (s_2 - s_1)}{r_L \cdot s_3} r_M \cdot s_3 \\ &= r_M \cdot s_1 \\ &= r_M \cdot s_2 \\ r_L \cdot s_1' &= r_L \cdot \left(s_1 + \frac{r_L \cdot (s_2 - s_1)}{r_L \cdot s_3} s_3 \right) \\ &= r_L \cdot s_1 + \frac{r_L \cdot (s_2 - s_1)}{r_L \cdot s_3} r_L \cdot s_3 \\ &= r_L \cdot s_2 \end{aligned}$$

7a.iv. (2 points) Describe a real spectral power distribution that satisfies the conditions for s_3.

Solution: One answer is a monochromatic light around 700nm. This appears as a red color, and stimulates the L cone but not the M or S cones.

(7b) Color Reproduction Using Color Matching Functions.

7b.i. (4 points) Someone runs a color matching experiment for you on a display with red, green and blue pixels. A test subject looks at monochromatic light of each wavelength in turn, and determines the red, green and blue pixel intensities required to achieve a color match. The results of the experiment are given to you as color matching functions, represented as vectors $\overline{r}$, $\overline{g}$, $\overline{b}$. For each wavelength, the color matching function gives the intensity of the corresponding pixel color required to achieve a match. Now, you are given a target spectral power distribution, represented as a vector s. Write down expressions for the intensities (scalar) R, G and B for the red, green and blue pixels, required to reproduce the color of s on the display.

Solution:

$$R = s \cdot \overline{r}$$
$$G = s \cdot \overline{g}$$
$$B = s \cdot \overline{b} \tag{14}$$

7b.ii. (2 points) What does it mean if the R, G or B values obtained in the previous part are negative?

Solution: It means that the color of s is outside the color gamut of the display and cannot be reproduced exactly.

(7c) (2 points) Perceptual Color Spaces

In the CIELAB color space, a non-linear gamma function is used to warp the value scale of L^*, a^* and b^*. In a sentence, what benefit does this non-linear warping provide?

Solution: It provides a more perceptually uniform color space.