CS 184/284A Spring 2025
Exam 2

April 28th, 2025
Time Limit: 110 minutes

Name: _____

SID Number: _____

Room: _____

Instructions:

- Please wait to look past this page until official start time.
- You must carefully read and sign the Honor Code below.
- Initial the top of every page, in case pages separate while scanning.
- This exam is closed book, except for two 8.5×11 pages of notes (double sided), handwritten.
- This exam is: 110 minutes long, 100 points, 16 pages (including this cover page) and 7 parts.
- Problem difficulty varies, so move on if feeling stuck on a problem.
- Answer each question in the space provided. Partial credit may be given on certain problems.
- ChatGPT helped with this exam, but all text was finalized by human and full responsibility is taken by your instructor.
- No corrections will be given during the exam. State any assumptions you require on page 18 (the Assumptions box).

Question	Points
1	20
2	10
3	20
4	8
5	10
6	12
7	20
Total:	100

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 signing 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 resources during this exam aside from my own page of notes.
- I will inform course staff immediately about any Honor Code violations I become aware of.

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.

Signature: _____

Date: _____

1. (Total : 20 points) Warmup - True / False

   Mark each statement True or False. (1 point each)

   (1a) T ◯ F ◯ **[Soln: F]** Importance sampling produces a biased estimate of the original continuous scene.

   (1b) T ◯ F ◯ **[Soln: T]** Path tracing stochastically estimates the integral in the rendering equation by simulating light paths through the scene.

   (1c) T ◯ F ◯ **[Soln: T]** Microfacet models assume a surface is composed of many tiny mirror-like facets with a distribution of orientations.

   (1d) T ◯ F ◯ **[Soln: T]** When imaging stars in the night sky with a regular camera, the circle of confusion increases as the sensor distance from the focal plane increases.

   (1e) T ◯ F ◯ **[Soln: T]** If you decrease your aperture from f/2.8 to f/5.6 and do not change your shutter speed, you must increase ISO by 4× to keep exposure content.

   Solution: True. Each stop halves light, so 2 stops down on aperture means that we must stop up ISO twice to maintain the same exposure.

   (1f) T ◯ F ◯ **[Soln: T]** The convolution theorem allows us to speed up convolutions with kernels of large spatial extent by performing multiplication in the frequency domain.

   Solution: True; convolutions with big kernels is often much more expensive than doing 2 FFTs, pointwise multiplication, and an inverse FFT.

   (1g) T ◯ F ◯ **[Soln: T]** In HDR imaging, multiple exposures are combined to capture scenes with high dynamic range.

   (1h) T ◯ F ◯ **[Soln: F]** Forward Euler can always be made stable if the time step is made small enough.

   (1i) T ◯ F ◯ **[Soln: F]** Forward kinematics lets the computer automatically compute the angle of each joint, whereas inverse kinematics allows the animator to have direct control over the angles of each joint.

   Solution: False. Forward kinematics gives full control over every angle, whereas inverse kinematics allows control over only the end position, and automatically calculates every angle.

   (1j) T ◯ F ◯ **[Soln: T]** Verlet integration is often used in particle-based simulations because of its numerical stability even though it does not conserve energy.

   (1k) T ◯ F ◯ **[Soln: T]** If two different lights have identical hue, saturation, and value, then they are metamers.

   Solution: True. Suppose A and B are metameric to some color C in HSV space. Then A and B are metameric by transitivity of metamerism.

   (1l) T ◯ F ◯ **[Soln: F]** The RGB color space is perceptually uniform.

   (1m) T ◯ F ◯ **[Soln: T]** Color constancy allows humans to perceive object colors as relatively constant under different lighting.

(1n) T ◯ F ◯ **[Soln: T]** The accommodation-vergence conflict in current headsets happens because your eyes must stay focused (accommodated) at a relatively distant screen while converging to fuse a nearby virtual object.

Solution: True

(1o) T ◯ F ◯ **[Soln: T]** In virtual reality (VR) the user is completely immersed and sees only light produced by the headset display, whereas in augmented reality (AR) the display simply overlays information on the real world.

Solution: True

(1p) T ◯ F ◯ **[Soln: T]** Foveated rendering lowers image resolution in the periphery based on eye-tracking, cutting GPU cost while preserving detail at the gaze point.

Solution: True

(1q) T ◯ F ◯ **[Soln: F]** In ray tracing, shadows are determined only by checking rays from the light source to the camera.

Solution: False; shadow rays go from the surface point to the light, not to the camera.

(1r) T ◯ F ◯ **[Soln: T]** Radiance does not change with distance along a ray in a vacuum.

Solution: True; radiance is invariant along a ray in a vacuum.

(1s) T ◯ F ◯ **[Soln: F]** Bilinear filtering is sufficient by itself to remove aliasing from minified textures.

Solution: False; there are cases where mipmaps are required (bilinear alone is not sufficient).

(1t) T ◯ F ◯ **[Soln: T]** Uniform spatial partitions suffer in cases like the "teapot in a stadium".

Solution: True; grid cells will be insufficiently sized as many cells will have little to no data (just in the stadium). The uniform grid is mainly beneficial in circumstances where the density of objects is more uniform.

2. (Total : 10 points) Physical Simulation and Animation

(2a) [4 points] Below are the five main operations that are run each timestep in the mass-spring model from Homework Assignment 4. However, they may or may not be in the correct order. Write a number next to each step indicating the correct order to run them in, with 1 being the first step and 5 being the last step.

☐ Handle external collisions

☐ Apply external forces

☐ Constrain length of springs

☐ Use Verlet integration to update positions

☐ Handle self collisions

Solution:

4 Handle external collisions

1 Apply external forces

5 Constrain length of springs

2 Use Verlet integration to update positions

3 Handle self collisions

(2b) [6 points] Using unmodified explicit Euler's method, calculate the position over time of a particle with unit mass in a vector field (external force) defined by force$(x, y) = (-2, -y/2)$ starting at time $t = 0$ with position $(-2, 2)$ and velocity $(2, -2)$. Start at $t = 0$ and end at $t = 4$, with time-step $\Delta t = 2$. Draw the positions of the particle on each time step on the graph below.

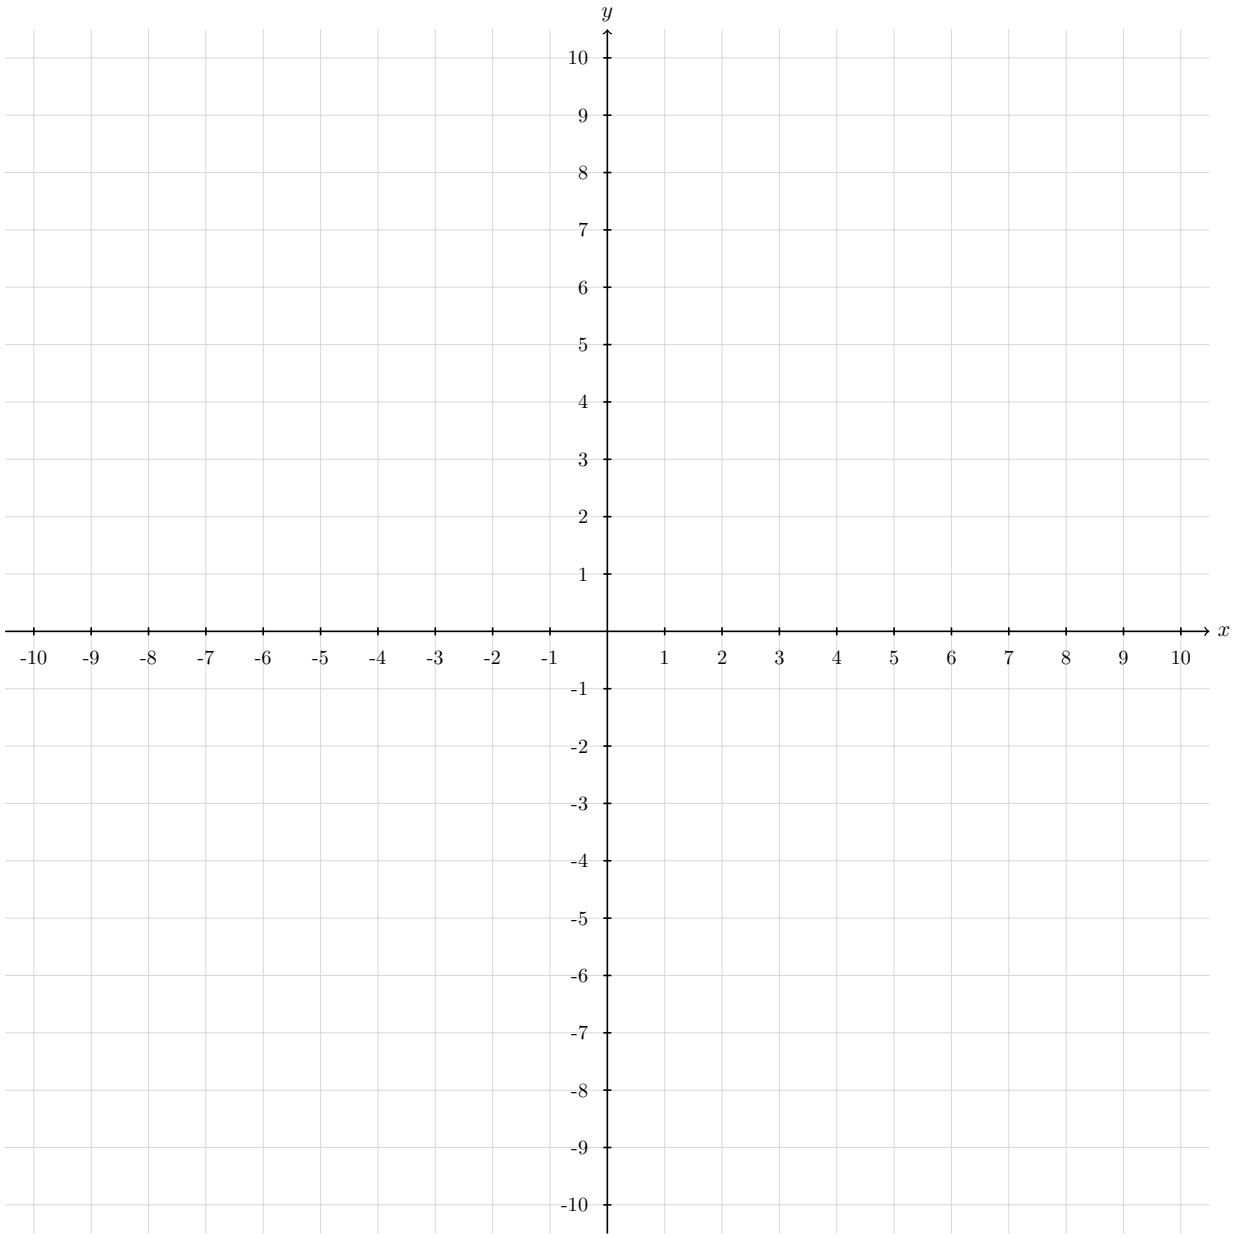

Solution:
$t = 0, p = (-2, 2), v = (2, -2), a = (-2, -1)$
$t = 2, p = (2, -2), v = (-2, -4), a = (-2, 1)$
$t = 4, p = (-2, 10), v = (-6, -2), a = (-2, 5)$

3. (Total : 20 points) Color Science

(3a) [4 points] Nami, a display engineer, is building a new kind of TV that has 4 color primaries (Red, Green, Blue, and X) instead of just the usual 3. Nami believes that the added primary X will expand the color gamut of the TV, and therefore allow it to show new colors it could not before. Zoro, the company boss, is unconvinced. After all, 3 primaries is apparently good enough for every other TV on the market.

You decide to resolve this dispute by running a formal color-matching test. Before you run the test, you think about how things could turn out. Which of the following hypothetical outcomes would show that Nami is right? **Select all that apply.**

- ◯ 0% Green + 30% Blue + 5% X = 35% Red.
- ◯ 30% Red + 60% Blue + 10% X = 10% Green.
- ◯ 5% Red + 5% Green + 5% X = 15% Blue.
- ◯ 10% Red + 5% Green + 20% Blue = 70% X.

Briefly explain your answer in the box below.

Solution: Only the first three are correct. First, re-arrange the equations to solve for X. In each case, at least one of the RGB primaries needs to be negative in order to make a match, which is impossible. These situations would show that X is outside the gamut of RGB, and thus is a new color to add to the TV. In contrast, the fourth option instead shows that X can already be made with RGB.

(3b) [4 points] You run the experiment and find that Nami is right. However, you discover later that your test subject, Chopper, is color-blind and has only two out of the three cone types. Chopper's result indicated he saw that X is outside of the gamut of RGB. Would this still be true for a color-normal observer, i.e. someone with all three cone types? Briefly explain why or why not.

Solution: Yes, X would still be outside of the RGB gamut. There was no setting of Red, Green, or Blue that could satisfy both of Chopper's cone types at once. Those two cones would still be unsatisfied regardless if the missing third one was. Therefore, X is still outside of the gamut of Red, Green, and Blue.

(3c) (Note: this question is unrelated to the previous parts.)

It is discovered that a pod of exceptionally intelligent dolphins have engineered their own smartphone, called the FishPhone 5. A key feature of the FishPhone is that it can take a photo of a color, and then display a flawless metamer of that color on its screen. Although dolphins only have **one** cone type, you are studying this technology to see if it is also possible for humans.

We have the following definitions:

- The FishPhone's single light detector has detection efficiency $r(\lambda)$.
- The FishPhone's single output primary has power spectra $p(\lambda)$.
- The detector is taking a photo of a light spectra, $n(\lambda)$.
- The dolphin cone's spectral sensitivity function is $s(\lambda)$.

3c.i. [1 point] The FishPhone takes a picture of the spectra n. Write the equation for x, the response of the light detector.

Solution:

$$x = \int r(\lambda)n(\lambda)d\lambda$$

3c.ii. [2 points] Next, the FishPhone applies some function f onto x to control how intensely to power its output primary. Write the equation for $y(\lambda)$, the total output power spectra. You can assume $f(x)$ is given.

Solution:

$$y(\lambda) = p(\lambda)f(x)$$

3c.iii. [3 points] Next, the dolphin user looks at the color displayed on the FishPhone. Let z_{actual} be the response of the dolphin's single cone type to output light y. Also, let z_{goal} be the cone response that the FishPhone is trying to achieve. Write an equation for each z. You can assume x and y are given if needed.

Solution:

$$z_{\text{actual}} = \int s(\lambda)y(\lambda)d\lambda$$

$$z_{\text{goal}} = \int s(\lambda)n(\lambda)d\lambda$$

3c.iv. [4 points] Recall that we want $z_{\text{actual}} = z_{\text{goal}}$. Specify a choice of f and r which achieves this. Designing underwater screens is difficult, so assume that p is given and cannot be changed. Your solution must not depend on n.

Hint: $f(x) = ax$ for some scalar a.

Solution:

$$r(\lambda) = s(\lambda)$$

$$f(x) = x \div \int s(\lambda)p(\lambda)d\lambda$$

3c.v. [2 points] You want to bring the FishPhone's capabilities to human smartphones. Describe (in words) a choice of analogous r_1, r_2, r_3 for a human smartphone that would enable perfect metamer display for humans. Ignore any potential issues with the display's output gamut.

Solution: Set r_1, r_2, r_3 to be the human photoreceptor sensitivities for L M and S cones.

4. (Total : 8 points) Virtual Reality

(4a) [2 points] In a buggy VR app, the left and right eye cameras are rendering from the same virtual camera position. What visual error would this cause?

Solution: No depth perception, since the scene appears flat or 2D, breaking immersion.

(4b) [2 points] In an AR headset, a virtual coffee mug should appear anchored to a real table. If the system has higher latency than it should and the user rotates their head quickly towards the right, what happens to the perceived alignment of the coffee mug relative to the table?

Solution: The virtual coffee mug appears to momentarily drift incorrectly toward the right, lagging behind the correct, left-ward moving position it should appear in the visual field to appear anchored at the same position on the real table.

(4c) [4 points] (Virtual Reality) In VR what is the **vergence-accommodation conflict**? Explain why this occurs in typical head-mounted displays and describe one hardware or software technique that has been proposed to address it.

Solution: Vergence-accommodation conflict arises when the eyes converge (rotate inward or outward) to look at different virtual depths but must accommodate (focus) at a fixed display distance (the screen in the headset). In natural vision, vergence and accommodation are coupled, but VR breaks this coupling, which can cause eye strain, discomfort, or blurred vision.

One technique to address this:

- **Varifocal displays:** Dynamically adjust the focal distance of the display based on where the user is looking (e.g., using eye tracking).
- **Light field displays:** Emit rays in different directions to simulate correct focal cues at multiple depths.

5. (Total : 10 points) Short Answer: Rasterization, Geometry, Monte Carlo

(5a) [2 points] (Rasterization) After the `LookAt` transform (recall that this transforms from world space coordinates to camera space coordinates) is applied in the graphics pipeline, but before the projection matrix, describe:

(a) Where is the camera located?

(b) What direction does the $-z$ axis point?

Solution: (a) The camera is at the origin. (b) The $-z$ axis points along the viewing direction (into the scene).

(5b) [2 points] (Geometry) In a HalfEdge mesh data structure, each `HalfEdge` points to the *opposite* `HalfEdge` (sharing the same edge but opposite direction) and one other `HalfEdge` element. Define / describe that other `HalfEdge` element.

Solution: The **next** HalfEdge in the current face loop (proceeding around the face).

(5c) [2 points] (Barycentric) When using barycentric coordinates to interpolate the color across a triangle, where the three corners are pure red, green and blue, what are the barycentric coordinates of the point in the triangle that is gray?

Solution: $\left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$.

(5d) [2 points] (Monte Carlo) For an unbiased Monte Carlo estimator $\hat{I}_N$, how does its variance scale with the number of independent samples N?

Solution: $\text{Var}[\hat{I}_N] \propto \frac{1}{N}$.

(5e) [2 points] (Texture mapping)

Suppose we are using MIP-mapping on a texture as implemented in Homework 1. Suppose also that we are rasterizing a single pixel on the screen whose footprint projects to a large $N \times N$-pixel area on the texture image. What level MIP-map do we use? Assume trilinear filtering.

Solution: $\log_2(N)$

6. (Total : 12 points) Cameras, Lenses, Imaging

(6a) Consider the following diagram, which you may remember from discussion section. This diagram is not drawn to scale!

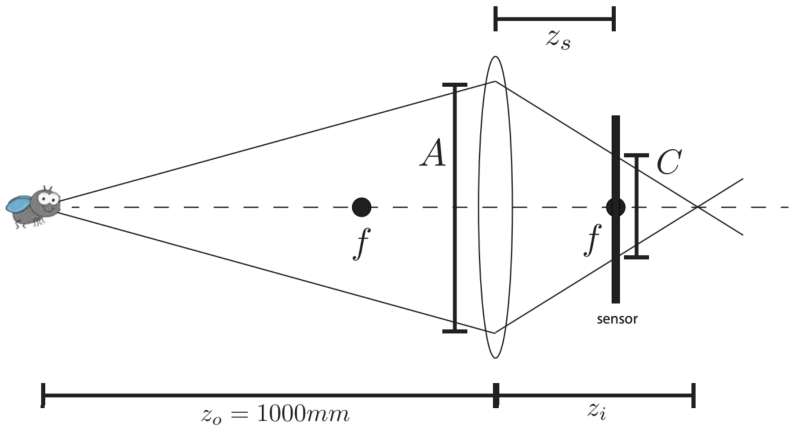

In this variant of that problem, assume that:

- $A = 25$mm
- $C = 2.5$mm

6a.i. [6 points] Calculate focal length f of this lens. Show your thinking and work for full or partial credit.

Solution:
The image sensor is at the focal point, so $z_s = f$. Using similar triangles, notice that $\frac{A}{z_i} = \frac{C}{z_i - z_s} = \frac{C}{z_i - f}$. Also recall that using the Thin Lens Equation, $z_i = \frac{z_o f}{z_o - f}$.

$$z_i - f = \frac{C z_i}{A} \tag{1}$$

$$f = z_i \left(1 - \frac{C}{A}\right) \tag{2}$$

$$f = \frac{z_o f}{z_o - f} \left(1 - \frac{C}{A}\right) \tag{3}$$

$$z_o f - f^2 = z_o f \left(1 - \frac{C}{A}\right) \tag{4}$$

$$f \left(f + z_o \left(1 - \frac{C}{A}\right) - z_o\right) = 0 \tag{5}$$

Because f is not zero, we solve for the non-zero f value as $f = z_o \frac{C}{A} = 1000 \left(\frac{2.5}{25}\right) = 100$mm. Equivalently, and maybe even more efficiently, you could plug in the calculation of $z_i = \frac{f}{1 - \frac{C}{A}}$ into the Thin Lens Equation, then solve for f.

6a.ii. [3 points] There is another sensor plane, call it z'_s, where the circle of confusion would be exactly the same size as shown here for z_s. What is that depth z'_s? You can assume f is given. Show your work and you can round your answer to the nearest tenth of a millimeter.

Solution:

$$z'_s = z_i + (z_i - z_s) = 2\frac{z_o f}{z_o - f} - f$$

Plugging in the value of f as found in the previous part, $\frac{2 \cdot 1000 \cdot 100}{1000-100} - 100 = \frac{2000}{10-1} - 100 \approx 122.22$mm.

6a.iii. [3 points] What are the main visual differences between the photographs obtained at the two depths?

Solution: The image in the original camera setting will have smaller maginfication than the image in the second camera setting. So the fly will look larger at the second depth, as well as more blurry.
To be specific, the magnification in the first image is $M_1 = \frac{f}{z_o} = 0.1$. The magnification in the second image is $M_2 = \frac{z'_s}{z_o} \approx 0.122$, so the second image is 82%.

7. (Total : 20 points) Global Illumination and Realistic Materials

(7a) [4 points] **You Would Not Believe Your Eyes** Look at the images below, comparing the same scene lit by a large and small light source. Assuming uniform hemisphere sampling is used in evaluating the direct lighting integral at each point in the scene, why might smaller light sources create "fireflies," which are the very noisy pixels in the scene?

Large light sources:

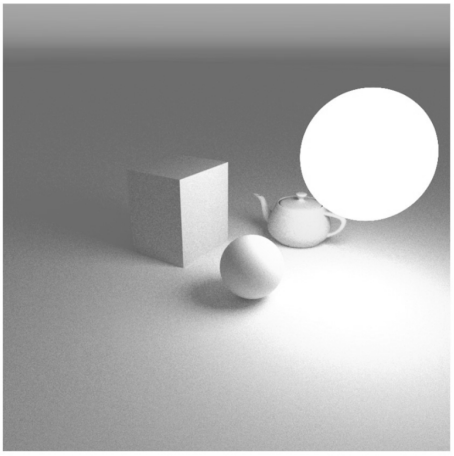

Small light sources:

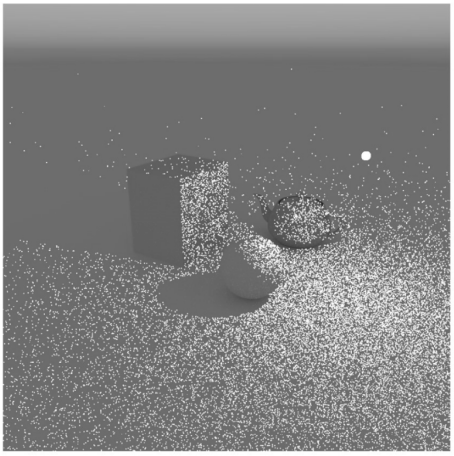

Solution: In uniform hemisphere sampling, rays are distributed evenly over the hemisphere. If a light source is small, the probability of a randomly sampled ray hitting it is very low. When a ray does happen to hit the small light, it carries a disproportionately large amount of energy (because its contribution is scaled by the inverse probability of hitting the light). This creates very bright, isolated pixels known as "fireflies." They occur because a few rare samples contribute a lot of light, while most samples contribute little or none.

(7b) Microfacet Reflection Analysis: Circular vs. Linear Specular Footprints

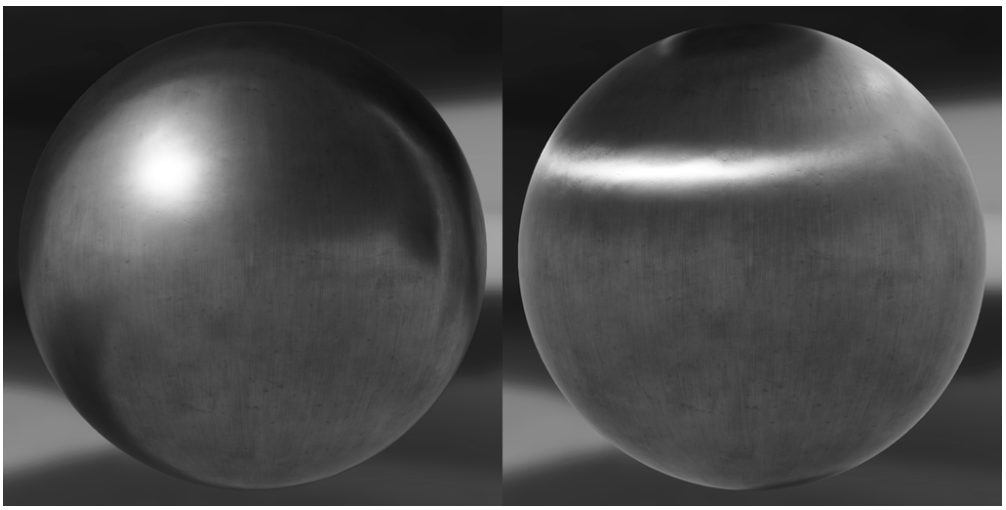

The image above (credit: https://google.github.io/filament) shows examples of an isotropic BRDF (left) vs anisotrophic BRDF (right). Notice the round vs linear specular highlights that are visually common with these materials.

Isotropic and Anisotropic BRDFs

In this problem, we'll analyze the appearance of these two types of BRDFs on a flat surface. You can assume the microfacet BRDF model discussed in lecture. The isotropic and anisotropic BRDFs have microfacet normal distribution functions (NDFs) defined as follows:

$$D_{\text{iso}}(N_x, N_y, N_z) = \frac{H(N_z)}{2\pi\,\sigma^2}\ \exp\left[-\frac{N_x^2 + N_y^2}{2\sigma^2}\right],$$

$$D_{\text{aniso}}(N_x, N_y, N_z) = \frac{H(N_z)}{2\pi\,\sigma_x\sigma_y}\ \exp\left[-\frac{N_x^2}{2\sigma_x^2} - \frac{N_y^2}{2\sigma_y^2}\right].$$

Here, $\mathbf{N} = (N_x, N_y, N_z)$ is a unit-length normal vector with $N_x^2 + N_y^2 + N_z^2 = 1$ and $N_z > 0$. $H(N_z)$ is the Heaviside step function, and ensures that the density is evaluated only on the upper hemisphere:

$$H(N_z) \ = \ \begin{cases} 1, & N_z > 0, \\ 0, & N_z \le 0. \end{cases}$$

For the isotropic case, $0 < \sigma < 1$ is the roughness parameter. For the anisotropic case, $0 < \sigma_x < 1$ and $0 < \sigma_y < 1$ control roughness in the x and y directions respectively, with

$$\sigma_y \gg \sigma_x$$

to model a combed surface with elongated microfacets along the y-axis.

Viewing a Flat Glossy Panel Lit By a Point Light

A large glossy panel lies in $z = 0$. A point light shines from the position $p_l = (1, 0, 1)$. A camera is positioned from the opposite direction at $p_c = (-1, 0, 1)$ and is pointed at the origin. We'll mathematically analyze the shape of the specular highlight that will result from each of the BRDFs.

7b.i. [2 points] First, for each point $(x, y, 0)$ on the panel, calculate normalized (unit-length) light and view vectors, l and v, pointing from $(x, y, 0)$ to the light and the camera, respectively.

7b.ii. [6 points] Recall the definition of the half-angle vector, $h = \text{bisector}(l, v) = \frac{l+v}{\|l+v\|}$. (In BRDF theory, recall that a microfacet reflects light from the light to the camera if h is parallel to the microfacet normal vector). Calculate the half-angle vector h.

7b.iii. [4 points] If we further assume that $|x|$ and $|y|$ are small (where the specular highlight is relevant), we can approximate and simplify the formula for the half-angle vector to:

$$h(x, y) \approx \frac{(-x, -y, 1)}{\sqrt{x^2 + y^2 + 1}}.$$

Now, using this approximation and the two NDF formulas above, write down expressions for the reflected radiance in the direction of the camera: $L_{\text{iso}}(x, y) \propto D_{\text{iso}}\big(h(x, y)\big)$ and $L_{\text{aniso}}(x, y) \propto D_{\text{aniso}}\big(h(x, y)\big)$.

7b.iv. [4 points] Will the specular highlight viewed from the camera appear round and linear for the isotropic and anisotropic BRDFs? Justify your answer briefly.

Solution:

7b.i. [2 points] **Light and view directions**

For a surface point $P(x, y, 0)$

$$L(x, y) = \frac{\ell - P}{\|\ell - P\|} = \frac{(1 - x, -y, 1)}{A}, \qquad A = \sqrt{(1 - x)^2 + y^2 + 1},$$

$$V(x, y) = \frac{e - P}{\|e - P\|} = \frac{(-1 - x, -y, 1)}{B}, \quad B = \sqrt{(1 + x)^2 + y^2 + 1}.$$

7b.ii. [4 points] **Half-vector**

Add the two normalised directions and renormalise:

$$H(x, y) = \frac{L + V}{\|L + V\|} = \frac{\big(\Sigma_x, \Sigma_y, \Sigma_z\big)}{\sqrt{\Sigma_x^2 + \Sigma_y^2 + \Sigma_z^2}},$$

where

$$\Sigma_x = \frac{1 - x}{A} - \frac{1 + x}{B}, \quad \Sigma_y = -y\Big(\frac{1}{A} + \frac{1}{B}\Big), \quad \Sigma_z = \frac{1}{A} + \frac{1}{B}.$$

Not part of the solution, but relevant for transition to the next part, for small $|x|, |y|$ (region where highlights matter) the two denominators A and B are almost

equal, so the direction of H is very close to

$$\boxed{H(x,y) \;\approx\; \frac{(-x,\,-y,\,1)}{\sqrt{x^2+y^2+1}}}.$$

7b.iii. [4 points] **Reflected radiance**

Insert $\mathbf{N} = H(x,y) = (N_x, N_y, N_z)$ in the NDFs.

—

Isotropic BRDF

$$\begin{aligned} L_{\text{iso}}(x,y) &\propto \frac{1}{2\pi\sigma^2}\,\exp\Big[-\tfrac{N_x^2+N_y^2}{2\sigma^2}\Big] \\ &= \frac{1}{2\pi\sigma^2}\,\exp\Big[-\frac{x^2+y^2}{2\sigma^2\big(x^2+y^2+1\big)}\Big]. \end{aligned}$$

—

Anisotropic BRDF

$$\begin{aligned} L_{\text{aniso}}(x,y) &\propto \frac{1}{2\pi\sigma_x\sigma_y}\,\exp\Big[-\tfrac{N_x^2}{2\sigma_x^2}-\tfrac{N_y^2}{2\sigma_y^2}\Big] \\ &= \frac{1}{2\pi\sigma_x\sigma_y}\,\exp\Big[-\frac{x^2}{2\sigma_x^2(x^2+y^2+1)}-\frac{y^2}{2\sigma_y^2(x^2+y^2+1)}\Big]. \end{aligned}$$

7b.iv. [4 points] **Footprint shapes**

- **Isotropic highlight.** L_{iso} depends only on $r = \sqrt{x^2+y^2}$, so contours of equal brightness are *circles* centred at $(0,0)$. The specular spot therefore appears round.
- **Anisotropic highlight.** Because $\sigma_y \gg \sigma_x$, the exponent penalises x much more strongly than y. Brightness falls off rapidly in x but slowly in y, giving a long, narrow *elliptical streak*; when the anisotropy is extreme it approaches a thin *line* parallel to the y-axis, matching the photograph on the right.

Thus an isotropic NDF produces a circular specular footprint, whereas an anisotropic (combed) NDF yields an elongated, almost linear highlight, exactly as observed in practice.