

Computer Graphics - CSE2215 - 2020/2021

Resit Examination

1. Given a 512×512 grayscale image, if you apply a 3×3 Box filter (averages nine pixels) will the result be the same as applying first a 1×3 average filter followed by a 3×1 average filter? Both smaller filters (1×3 and 3×1) compute the average of three pixels. Explain your answer (only answering yes/no does not lead to any points). You can also ignore any border treatment for your explanation. You can use the following 3×3 image patch to help in your explanation. For example, the top left corner of this patch has gray value a , and so on.

a	b	c
d	e	f
g	h	i

Solution: Applying first the horizontal 1×3 filter the new values of the middle column are:

$$b' \leftarrow (a + b + c)/3$$

$$e' \leftarrow (d + e + f)/3$$


$$h' \leftarrow (g + h + i)/3$$

Then, applying the vertical 3×1 filter on the new middle column composed by b' , e' and h' :

$$e'' \leftarrow (b' + e' + h')/3 = (a + b + c + d + e + f + g + h + i)/9$$

which is the same as applying the 3×3 box filter directly over e .

2. Looking at this depiction, why are two vanishing points enough to create the perspective foreshortening? Given a camera projection matrix, how would you determine the locations of the vanishing points? What happens if the camera is tilted slightly downwards?



116

Taken from w3.joshuanava.biz

Solution: Only two points are needed because the vertical lines are parallel on the image, thus the third vanishing point is projected to infinity.
 To determine the locations take the directions of the sides of the building in 3D (with $w = 0$) and project using the projection matrix.
 If the camera is tilted a third vanishing point would appear at the top (currently, it is at infinity).

3. Given $P(\alpha) = \alpha * T1 + 2 * \alpha * T2 + (1 - 3 * \alpha) * T3$, can you provide constraints on α to ensure that the point P is inside the triangle $T1, T2, T3$ (assuming it is non-degenerate, meaning the points form an actual triangle)? Explain your choice. You can, if needed, rely on the test for being inside of a triangle, which we use for a general point on the supporting plane of a triangle.

Solution: $\alpha * T1 + 2 * \alpha * T2 + (1 - 3 * \alpha) * T3 = T3 + \alpha(T1 - T3) + 2 * \alpha(T2 - T1)$

Thus $\alpha \geq 0$ and $\alpha + 2 * \alpha \leq 1$, hence $\alpha \leq 1/3$.

We can also do it the other way around:

Know that a point lies inside of the triangle, when we have

$P = T3 + b(T2 - T3) + a(T1 - T3)$ with $a, b > 0$ and $a + b \leq 1$.

$T3 + b(T2 - T3) + a(T1 - T3) = (1 - a - b)T3 + aT1 + bT2$

We see $\alpha = a, 2 * \alpha = b, (1 - 3\alpha) = (1 - a - b)$, thus

$1 \geq \alpha \geq 0$ because of the constraint on a

$\alpha \geq 0, \alpha \leq 0.5$ because of the constraint on a

$\alpha \geq 0, \alpha \leq 0.5$ because of the constraint on a

$a + b < 1$ implies $(1 - 3\alpha) \geq 0$ thus $\alpha \leq 1/3$.

Hence, $0 \leq \alpha \leq 1/3$.

4. When points lie on the same line, they are called collinear. Which of the following points r is collinear to points $p = (2.0, 6.0, 2.0)$ and $q = (3.5, 8.5, 1.0)$? All points are expressed in homogeneous coordinates.


A. $r = (-0.25, 0.25, 1.0)$

B. $r = (-1.0, 1.0, 1.0)$

C. $r = (-1.0, -1.0, 1.0)$

D. $r = (2.0, 4.0, 1.0)$

5. Given a robot as illustrated below, your goal is to apply the necessary transformations to move the hands from the original (rest) pose to the target pose.



The parts expressed in their original coordinate frame have the following properties:

- all green parts (arm, forearm and hand) and the body are aligned with the vertical (y) axis
- the shoulders are aligned with the horizontal (x) axis
- all starting points are at the origin, where the starting point of a body part is the first vertex encountered when following the depicted connectivity from the red dot to this body part.

For the target pose, the following properties are given:

- the red point is located at q in R^2
- the blue parts are not rotated with regards to the rest pose
- the left arm (depicted on the right - see figure) is rotated 45° clockwise with regards to the left shoulder
- the left forearm is rotated 60° counter-clockwise with regards to the left arm in the target pose
- the left hand is aligned with the positive x-axis
- the right arm, forearm and hand are symmetrical with regards to the left side
- the position of the arm, forearm and hand can be inferred from the lengths given in the figure

Given the following two transformation matrices:

- T_p : translates by p , where p is a vector in R^2
- R_θ : rotates by angle θ counter-clockwise

Define the sequence of matrices to multiply the points of the hands by in order to draw it at the target position. You should describe two sequences, one for the left hand and one for the right hand. You do not need to provide a single matrix, a product of matrices is fine.

Solution:

left hand:


$$T(q)T(0, d_{body})T(d_{shoulder}, 0)R(45)T(0, -d_{arm})R(60)T(0, -d_{forearm})R(-15)$$

right hand:

$$T(q)T(0, d_{body})T(-d_{shoulder}, 0)R(-45)T(0, -d_{arm})R(-60)T(0, -d_{forearm})R(15)$$

6. We take a photo of the leaning Pisa tower by placing our pinhole camera on the ground at a distance of 10m from the base of the tower (point p). Consider that the camera centre is at the origin, the camera is oriented along the positive z-axis, the image plane is at $z = 2$, and the ground is the plane $y = 0$.

Point r' is the projection of the top of the leaning tower (point r) and has height equal 1m on the image plane (the exercise uses an unconventionally large camera to ease calculations). Knowing that if the tower stood straight its top would be at point q at a distance of 10m from the ground, what is the distance from the top of the leaning tower to the ground?



- A. 5
B. 6

C. 7

D. 8

7. Given a point light source in R^3 at position $l = (0, 4, \sqrt{3})$ and a point $p = (0, 1, 0)$ on a diffuse surface (no specular component). The normal of the surface at point p is $n = (0, 0, 1)$. The light source has the following ambient, diffuse and specular RGB intensities: $(0, 0, 0)$, $(0.8, 0.8, 0.5)$, $(1, 1, 1)$, respectively. Using the Phong Illumination Model, what should the diffuse component of the surface (k_d) be in order for an observer at position $e = (2, 2, 7)$ to perceive the colour at point p as $(0.3, 0.4, 0.1)$?
- A. $k_d = (0.3, 0.4, 0.1)$
B. $k_d = (0.75, 1.0, 0.4)$
C. $k_d = (0.35, 0.5, 0.2)$
D. $k_d = (0.15\sqrt{3}, 0.2\sqrt{3}, 0.05\sqrt{3})$
8. Given a point light source at position $L = (-2, 3, 0)$, and a point $m = (1, 0, 0)$ on a mirror surface with normal $n = (0, 1, 0)$. If we rotate the mirror around point m , in what direction should its normal point, such that the light ray is reflected from m towards the point $Q = (4, 0, 0)$? Explain your steps to get to the answer (no points is attributed to just writing the final solution vector).

Solution:

normalize light direction $l = (-1/\sqrt{2}, 1/\sqrt{2}, 0)$

normalize desired reflected direction towards Q: $r = (1, 0, 0)$

add both and normalize again

normal direction $= (l + r)/\|l + r\| = (1 - 1/\sqrt{2}, 1/\sqrt{2}, 0)/(\sqrt{2 - \sqrt{2}})$

9. Assume a cylinder $x^2 + y^2 = 1$ made out of a mirror material and a light ray $r: r(t) = (2, -1, -2) + (-t, t, t)$. In what direction is the reflected light ray leaving? (Hint: Try $t = 1$ for the intersection) Calculate the result and explain your solution by commenting on it.

Solution:

ray : $(2, -1, -2) + a * (-1, 1, 1)$

For $t = 1$, we have $(1, 0, -1)$, which is on the cylinder.

$x^2 + y^2 = 1$ so the normal there is also $(1, 0, 0)$


The normalized incoming light direction is $(-1, 1, 1)/\sqrt{3}$

Hence, the reflected ray is: $(-1, 1, 1)/\sqrt{3} - 2 * (-1/\sqrt{3}, 0, 0)$

Thus: $(1/\sqrt{3}, 1/\sqrt{3}, 1/\sqrt{3})$

10. Assume two axis-aligned squares in the plane as shown in the image below. Square A has vertices (a, b, c, d) with respective RGB colours $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$, $(0, 1, 0)$. Square B has vertices (e, f, g, h) with respective colours $(1, 1, 0)$, $(0, 0, 1)$, $(0, 1, 0)$, $(0, 0, 1)$.

The colour of any point inside one of the squares is defined by bilinear interpolation of the colours of its vertices. Inside the overlapping region, the final colour is defined by adding the interpolated colours of the two squares at point p . Given the following corners: $a = (0, 0)$, $c = (1, 1)$, $e = (0.5, 0.5)$, $g = (1.5, 1.5)$, $p = (0.75, 0.75)$, compute the resulting colour at p . Explain how you arrived at your solution (just providing the colour of p does not count as a valid explanation).



Solution: The vertices have the following weights:

$$w_a = w_g = 0.25 * 0.25 = 1/16$$

$$w_b = w_h = 0.75 * 0.25 = 3/16$$

$$w_c = w_e = 0.75 * 0.75 = 9/16$$

$$w_d = w_f = 0.25 * 0.75 = 3/16$$

The colour of p for square A is:

$$p_a = w_a * (1, 0, 0) + w_b * (0, 1, 0) + w_c * (0, 0, 1) + w_d * (0, 1, 0) = 1/16 * (1, 0, 0) + 3/16 * (0, 1, 0) + 9/16 * (0, 0, 1) + 3/16 * (0, 1, 0) = (1/16, 6/16, 9/16)$$


The colour of p for square B is:

$$p_b = w_e * (1, 1, 0) + w_f * (0, 0, 1) + w_g * (0, 1, 0) + w_h * (0, 0, 1) = 9/16 * (1, 1, 0) + 3/16 * (0, 0, 1) + 1/16 * (0, 1, 0) + 3/16 * (0, 0, 1) = (9/16, 10/16, 6/16)$$

The final colour is:

$$p_a + p_b = (1/16, 6/16, 9/16) + (9/16, 10/16, 6/16) = (10/16, 1, 15/16)$$

11. Given a screen of 400×200 pixels, we want to draw a red rectangle with a blue border. Our rasterization algorithm uses a simple overwrite method when primitives write to the same pixel. Hence, we will first draw a blue rectangle then a smaller red one, giving the result in the image below. The minimum and maximum coordinates of the red rectangle are $(-0.4, -0.75, 0, 1)$ and $(0.6, 0.25, 0, 1)$, respectively, after the rectangle was projected with the camera matrix. Please indicate the minimum and maximum coordinates of the blue rectangle after projection, each in the form $(x, y, 0, 1)$ with $x, y \in R$, if we want a border size of 4 pixels along the edges (see figure below)?



(the figure is for illustration only, it is not supposed to reflect the given coordinates precisely)

Solution:

1 unit horizontally = $2/400 = 1/200 = 0.005$

1 unit vertically = $2/200 = 1/100 = 0.01$

4 pixels horizontally = 0.02


4 pixels vertically = 0.04

corners of blue rectangle:


min (-0.42, -0.79)

max (0.62, 0.29)

12. Given the simplified dart board below formed by two concentric circles with alternating colours. The colour transitions occur precisely at the x and y axis.



To define the underlying geometry we approximate the outer circle by triangles. We start with four triangles where two edges of each triangle are aligned with the x and y axis, and one vertex is at the centre of the circle. We then subdivide each triangle into two by splitting in half the edge opposite to the centre vertex, obtaining 8 triangles. This subdivision scheme is further repeated many times until the outer circle is well approximated by the triangles, as illustrated below. Given this geometric approximation of the circle by triangles, which of the following textures could be used to achieve our board image above? You should pick the texture with minimum amount of texels that is able to texturize the board, and assume that you can specify new texture coordinates for each triangle at each vertex.




A. 1x2 texture:




B. 1x4 texture:



C. 2x2 texture:



D. 4x4 texture:



13. Shadow mapping suffers from a resolution mismatch issue leading to blocky artifacts. We have seen that for color textures mipmapping can avoid some of the texture-related problems. Discuss if this method is useful for shadow mapping. Would you have a suggestion yourself on how to introduce filtering of Shadow-map shadows (short answer)?

Solution:

Usually not useful. Filtering depths leads to "blockers" that are floating in the air. Especially if depth is non-uniform.

Solutions could be PCF (compare neighboring texels, interpolate result), or one could compare the filtered depth using a falloff that is first always lit and then slowly becomes darker the more things are in shadow - it is a poor man's version of PCF.