Homework #1

Displays, Alpha Compositing, Image Processing,

Affine Transformations, Hierarchical Modeling

Assigned: Saturday, April 12th

Due: Thursday, April 24th
at the beginning of class

Directions: Please provide short written answers to the following questions on your own paper. Feel free to discuss the problems with classmates, but please ***answer the questions on your own and show your work.***

Please write your name on your assignment!

Problem 1: Short answer (10 points)

a) (6 Points) Much of our perception of 3D arises from the fact that we have two eyes, viewing a scene from two different viewpoints, so-called stereo vision. Given images recorded or rendered for two different viewpoints separated by the typical distance between human eyes, a 3D stereo display presents one image to one eye and the other image to the other eye; our eyes are then fooled into believing they are looking at an actual 3D scene. One approach to 3D displays is based on a standard color LCD display and a pair of LCD shutter glasses. The LCD display first shows a left-eye image, then a right-eye image, and so on, while the LCD shutter glasses synchronously let light reach the left eye, then the right, etc. An LCD shutter is essentially one giant LCD pixel (a crystal sandwiched between two polarizers) with no color filter, driven with a voltage to be either opaque or transmissive. The display and shutters are designed to give a reasonably bright picture when sitting naturally in front of the display. Assume that the LCD crystal at each display pixel is oriented the same way as every other pixel in the display (regardless of color filter).

- If you tilt your head sideways (i.e., tilting your head over to one of your shoulders, so that the imaginary line segment connecting your eyes is now aligned with the vertical direction), will the displayed images appear dimmer in one eye, both eyes, or neither eye? Justify your answer.

- Suppose you removed the LCD filter panel in front of the unpolarized backlight and looked at the even backlighting with your naked eye(s); you would see even, white light. Roughly how much dimmer would you expect that light to become after putting the panel back on and putting on the shutter glasses (which are turned on and shuttering), assuming the framebuffer is set to white at each pixel for both eyes? [Note that unpolarized light intensity is cut in half by linear polarization. Assume that the R,G,B sub-pixels each transmit 1/3 of the visible spectrum of the light. Perceived brightness is averaged over time.] Justify your answer.

b) (4 Points) Consider two vectors **u** and **v** which are of non-zero length and not parallel to each other. Which of the following is true and which is false:

1. $(\mathbf{u} \times \mathbf{v}) \times \mathbf{u} = \mathbf{u} \times (\mathbf{v} \times \mathbf{u})$

2. $\left[(\mathbf{u} \times \mathbf{v}) \times \mathbf{u}\right] \cdot \mathbf{u} = 0$

3. $\left[(\mathbf{u} \times \mathbf{v}) \times \mathbf{u}\right] \cdot \mathbf{v} = 0$

4. $\left\{\frac{\mathbf{u}}{\|\mathbf{u}\|} \times \left[\frac{(\mathbf{u} \times \mathbf{v}) \times \mathbf{u}}{\|(\mathbf{u} \times \mathbf{v}) \times \mathbf{u}\|}\right]\right\} \cdot \left\{\frac{\mathbf{u}}{\|\mathbf{u}\|} \times \left[\frac{(\mathbf{u} \times \mathbf{v}) \times \mathbf{u}}{\|(\mathbf{u} \times \mathbf{v}) \times \mathbf{u}\|}\right]\right\} = 1$

You do **not** need to justify your answer.

Problem 2: Image processing (24 points)

In this problem, you will consider several convolution filtering operations and their behaviors. You do not need to worry about flipping filters before sliding them across images; i.e., assume filters are pre-flipped. In addition, assume that the y-axis points up, the x-axis points to the right, and the lower left corner of the image is at (0,0). *For each sub-problem, justify your answer.*

a) (2 points) The image you're editing is too dark, and you decide you need to amplify the value of each pixel by a factor of 4. Suggest a convolution filter that will quadruple the value at each pixel of the image without changing it in any other way. (Technically, after scaling pixel values, they could be out of range; assume that any needed clamping will be taken care of later, after filtering).

b) (4 points) While taking a photograph with your digital camera, you fail to hold the camera steady, and it translates from left to right while the shutter is open. You discover this later when you see that vertical edges, in particular, have been blurred a bit (an effect called "motion blur"). You decide to filter the image so that vertical edges are sharpened, but horizontal edges are unchanged. Suggest a single convolution filter that does this.

c) (4 points) After thinking a little more about the previous picture, you decide that motion blur is cool, and you want to apply it to another image. In this case, though, you want to simulate the effect of a camera translating diagonally along the $x = -y$ direction while the shutter is open. Suggest a convolution filter that would accomplish some diagonal blurring along that direction by averaging across m pixels.

d) (4 points) Describe a non-constant image that, when convolved with your diagonal blur filter from (c), would be unchanged by the filter. (You may ignore the boundaries.)

e) (10 points) Suppose you pad the boundary of an image in some way that allows you to compute output values for every pixel being filtered by a convolution filter. For an image of dimensions n x n and a filter of dimensions m x m, how many output pixels will be influenced by input pixels "hallucinated" beyond the boundary of the image? For simplicity, assume that m is odd. However, m and n may otherwise have arbitrary positive values.

Problem 3: Triangle coordinates (24 points)

Consider triangle Δ_{ABC} and a point Q depicted below:

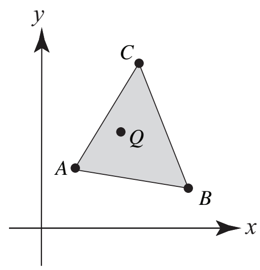

A, B, C, and Q lie in the x-y plane, so, **neglecting the homogeneous coordinate**, we can write out their **3D** coordinates as:

$$A = \begin{bmatrix} A_x \\ A_y \\ 0 \end{bmatrix} \quad B = \begin{bmatrix} B_x \\ B_y \\ 0 \end{bmatrix} \quad C = \begin{bmatrix} C_x \\ C_y \\ 0 \end{bmatrix} \quad Q = \begin{bmatrix} Q_x \\ Q_y \\ 0 \end{bmatrix}$$

The last coordinate is the z coordinate in this case, and we know that $A_z = B_z = C_z = Q_z = 0$. Note that Q is depicted as lying inside of Δ_{ABC}, but you should not assume that it does unless stated otherwise in a sub-problem. Assume that $A \neq B \neq C$, i.e., the triangle is not "degenerate." Further assume that if you curl the fingers of your right hand from A to B to C, your thumb will point in the direction of the positive z-axis.

(a) (3 points) Using cross and/or dot products, devise a test (with the help of equations) to determine if point Q lies inside of Δ_{ABC}. You should assume that the edges and vertices of the triangle are part of the interior of the triangle. [You do not need to expand any cross or dot products in your answer, but you may do so if it helps you.]

(b) (2 points) Using cross and/or dot products, compute the unit-length normal to Δ_{ABC}. Your solution should work for 3D points in general, i.e., not depend on the fact that these points lie in the x-y plane. We will use the right-hand rule for triangles, which means that as you curl the fingers of your right hand from A to B to C, your thumb will point in the direction of the normal. [You do not need to expand any cross or dot products in your answer, but you may do so if it helps you.]

(c) (3 points) Using cross and/or dot products, compute the area of the triangle, Area(Δ_{ABC}). This time you **do** need to expand all cross and/or dot products based on the elements of A, B, and C. Multiply out all terms; e.g., an expression like $(a+b)(c+d)$ should be expanded to $ac+ad+bc+bd$.

For the remainder of the problem, we will safely ignore the z coordinate and work in **2D** only, but **now we will keep track of the affine coordinate** (w=1 for all points). We can now represent the triangle vertices and the point Q as:

$$A = \begin{bmatrix} A_x \\ A_y \\ 1 \end{bmatrix} \quad B = \begin{bmatrix} B_x \\ B_y \\ 1 \end{bmatrix} \quad C = \begin{bmatrix} C_x \\ C_y \\ 1 \end{bmatrix} \quad Q = \begin{bmatrix} Q_x \\ Q_y \\ 1 \end{bmatrix}$$

Again, to be clear, the last coordinate is now the affine w coordinate, which, for affine points, is always 1; i.e., $A_w = B_w = C_w = Q_w = 1$.

(d) (2 points) Suppose we create a 3x3 matrix [A B C], i.e., a matrix with columns filled by A, B, and C. Write out this matrix, explicitly filling in the elements of the matrix in terms of the elements of A, B, and C, and compute its determinant, det[A B C]. Again, multiply out all terms.

(e) (1 point) Based on your answers to (c) and (d), what is Area(Δ_{ABC}) in terms of det[A B C]?

(f) (1 point) In general, we can express Q as a weighted sum of A, B, and C; i.e., $Q = \alpha A + \beta B + \gamma C$, where α, β, and γ are scalars. In order for this to be a proper affine combination, what constraint is placed on α, β, and γ? Explain.

(g) (3 points) Now we will work on solving for α, β, and γ. Write out a matrix equation of the form $M\,\mathbf{p} = \mathbf{r}$:

$$\begin{bmatrix} m_{00} & m_{01} & m_{02} \\ m_{10} & m_{11} & m_{12} \\ m_{20} & m_{21} & m_{22} \end{bmatrix} \begin{bmatrix} p_0 \\ p_1 \\ p_2 \end{bmatrix} = \begin{bmatrix} r_0 \\ r_1 \\ r_2 \end{bmatrix}$$

where M is a 3x3 matrix, $\mathbf{p}$ is the column vector of unknowns, i.e., $\mathbf{p} = [\alpha\ \beta\ \gamma]^T$ and $\mathbf{r}$ is a column vector with three elements. I.e., explicitly write out the matrix equation, filling in the elements of M, $\mathbf{p}$, and $\mathbf{r}$, in terms of α, β, γ, and the elements of A, B, C, and Q. (Do not apply the matrix, just set up the equation.) Hint: you can expand $Q = \alpha A + \beta B + \gamma C$ explicitly in terms of the elements of A, B, C, and Q and the result should be equivalent to your matrix equation.

(h) (2 points) We can solve for $\mathbf{p}$ using Cramer's rule. In particular, for a matrix equation $M\,\mathbf{p} = \mathbf{r}$ as above, we can solve for $\mathbf{p}$ as ratios of determinants:

$$p_0 = \frac{\det\begin{bmatrix} r_0 & m_{01} & m_{02} \\ r_1 & m_{11} & m_{12} \\ r_2 & m_{21} & m_{22} \end{bmatrix}}{\det\begin{bmatrix} m_{00} & m_{01} & m_{02} \\ m_{10} & m_{11} & m_{12} \\ m_{20} & m_{21} & m_{22} \end{bmatrix}} \quad p_1 = \frac{\det\begin{bmatrix} m_{00} & r_0 & m_{02} \\ m_{10} & r_1 & m_{12} \\ m_{20} & r_2 & m_{22} \end{bmatrix}}{\det\begin{bmatrix} m_{00} & m_{01} & m_{02} \\ m_{10} & m_{11} & m_{12} \\ m_{20} & m_{21} & m_{22} \end{bmatrix}} \quad p_2 = \frac{\det\begin{bmatrix} m_{00} & m_{01} & r_0 \\ m_{10} & m_{11} & r_1 \\ m_{20} & m_{21} & r_2 \end{bmatrix}}{\det\begin{bmatrix} m_{00} & m_{01} & m_{02} \\ m_{10} & m_{11} & m_{12} \\ m_{20} & m_{21} & m_{22} \end{bmatrix}}$$

Note how the denominator is always the determinant of M, and the numerator is the determinant of a matrix that consists of M with one of the columns replaced with the elements of $\mathbf{r}$. Based on your answer to (g), re-write these same "Cramer's rule" equations using α, β, γ, and the elements of A, B, C, and Q.

(i) (3 points) Assume Q is inside of Δ_{ABC}. In this case, all of the determinants in (h) are positive. Based on your answer to (e), re-write your answer to (h) in terms of areas of triangles.

(j) (2 points) Suppose $Q = B$. What should α, β, and γ be? Justify your answer in terms of areas of triangles.

(k) (2 points) Suppose Q is halfway between A and C. What should α, β, and γ be? Justify your answer in terms of areas of triangles.

Problem 4. 3D affine transformations (26 points)

The basic scaling matrix discussed in lecture scales only with respect to the x, y, and/or z axes. Using the basic translation, scaling, and rotation matrices, specify how to build a transformation matrix that scales along any ray in 3D space. This new transformation is determined by the ray origin **p** = (p_x, p_y, p_z) and direction vector **v** = (v_x, v_y, v_z), and the amount of scaling s_v. For clarity, a diagram has been provided, showing a box being scaled with respect to a given ray. Your answer should work for *any* ray, not just the case shown in the picture.

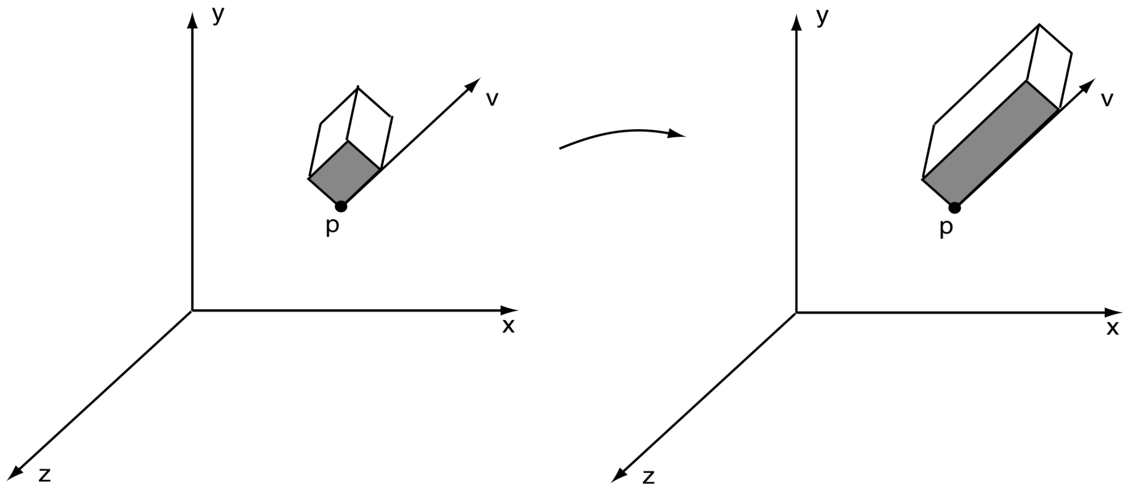

You can use any of the following standard matrices (from lecture) as building blocks: canonical axis rotations $R_x(\alpha)$, $R_y(\beta)$, $R_z(\gamma)$, scales $S(s_x,s_y,s_z)$, and translations $T(t_x,t_y,t_z)$. You don't need to write out the entries of the 4x4 matrices. It is sufficient to use the symbols given above, supplied with the appropriate arguments. All scale factors are strictly positive. You must compute the angles of any rotations required. Note that you may require inverse trigonometric functions, and you should assume that $\cos^{-1}(x)$ outputs a range of $[0..\pi]$, and that $\sin^{-1}(x)$ and $\tan^{-1}(x)$ each outputs a range of $[-\pi/2..\pi/2]$.

There are many possible solutions to this problem. To constrain the space of answers, and to give you a solution hint: you must cause the **v** direction to align with the y-axis at some stage of your solution.

Show your work, using words and drawings as needed to support your answer.

Problem 5. Hierarchical modeling (16 points)

Suppose you want to create the hierarchical model shown below. The model is comprised of five parts, labeled **A**, **B**, **C**, **D,** and **E**, and each part is drawn as one of five primitives given below (they are already scaled to the correct sizes). The following transformations are available to you:

- R(θ) – rotate by θ degrees (counter clockwise)
- T(a, b) – translate by $[a \;\; b]^T$

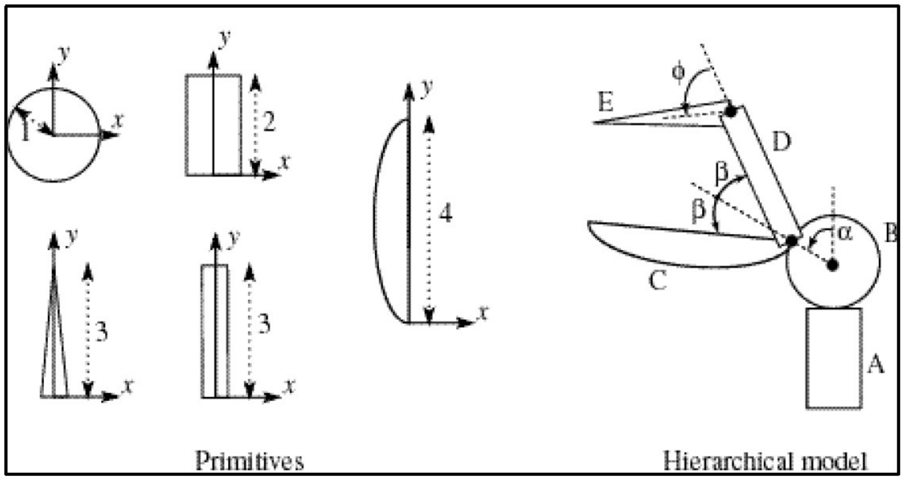

(a) (10 points) Construct a tree to specify the hierarchical model where the nodes of the tree should be labeled **A**, **B**, **C**, **D**, and **E**, with **A** as the root. Along each of the edges of the tree, write expressions for the transformations that are applied along that edge, using the notation given above (you do not need to write out any 3x3 (2D affine) matrices). Insert numerical values (i.e., for primitive sizes) where available. Remember that the order of transformations is important! Show your work wherever the transformations are not "obvious." Your tree should contain a bunch of boxes (or circles) each containing one part number letter; these boxes should be connected by line segments, each labeled with a corresponding transformation that connects child to parent. The tree must have one or more branches in it. If two parts are connected physically, then they should be connected in the tree, as long as you don't form a cycle by connecting them.

(b) (2 points) Write out the full transformation expression to be applied to the part labeled **E**.

(c) (2 points) Suppose we now want to compute the location of the tip of the part labeled **C**, where the tip is the point on **C** farthest from part **B**. One way to do this is to imagine adding one more primitive **F** that consists of an infinitesimal point at the origin (in its local coordinates, as was the case for each of the primitives in the drawing above). We can then add this primitive to the hierarchical model, attached to **C**, with a suitable transformation. What is the transformation expression to be applied to **F**?

(d) (2 points) To compute the location of the point at the tip of part **C**, we would just apply the transformation expression to the location of the point **F** in its local coordinates, which is (in 2D affine coordinates) $[0 \; 0 \; 1]^T$. Suppose $\alpha = 270^{\circ}$ and $\phi = \beta = 0$. Evaluate the location of the tip of **C** using the matrices in part (c), writing out the matrices in full to compute the result.