Homework #1

Displays, Image Processing,

Affine Transformations, Hierarchical modeling

Assigned: Saturday, April 9th

Due: Wednesday, April 20th
at the beginning of class

Directions: Please provide short written answers to the following questions, using this page as a cover sheet. Feel free to discuss the problems with classmates, but please ***answer the questions on your own***.

Name:____

Problem 1: Short answers (17 points)

Provide short answers to each of the following questions:

a) (3 points) What is double-buffering and what is its purpose?

b) (3 points) How do you normalize a convolution filter, and what is the purpose of doing so?

c) (3 points) Suppose, for each pixel of a grayscale image, you only average pixels in a 5x5 neighborhood that have values within +/-5% of the pixel being filtered. Could this be written as a convolution filter? Explain.

d) (4 points) How would you compute the unit-length normal to a triangle in 3D with vertices A, B, and C, specified according to the right-hand rule (where curling the fingers of your right hand from A to B to C will leave your thumb pointing along the normal direction)? What happens if A, B, and C are co-linear?

e) (4 Points) Consider a pair of three dimensional vectors, *u* and *v*, which are of non-zero length and not parallel to each other. Which of the following is True, False, or Nonsense (i.e., involving an operation that cannot be performed):

$$(v \times u) \times u = u \times (u \times v)$$

$$(v \bullet u) \times u = u \times (u \bullet v)$$

$$(v \times u) \bullet u = u \bullet (u \times v)$$

$$\frac{u}{\|u\|} \bullet v = u \bullet \frac{v}{\|v\|}$$

You do **not** need to justify your answer for part (e).

Problem 2: Color LCD displays (12 points)

LCD displays operate on the principle of polarizing light at the entrance to each crystal, twisting the polarization by a voltage controlled amount, then passing the resulting light through a final polarizer. The end result is that the light coming out of the LCD has a particular intensity and (fixed) polarization. The color of the light reaching the viewer is controlled with a color filter (red, green, or blue) over each crystal.

Light coming out of the LCD has “linear polarization,” characterized by an angle θ, with a corresponding polarization vector:

$$\mathbf{p} = \begin{bmatrix} \cos\theta \\ \sin\theta \end{bmatrix}$$

A polarizing filter also has an orientation β with polarization vector:

$$\mathbf{f} = \begin{bmatrix} \cos\beta \\ \sin\beta \end{bmatrix}$$

The filter only passes the component of the input polarization aligned with filter’s polarization. Specifically, it passes a fraction α of the light:

$$\alpha = \left|\mathbf{f} \bullet \mathbf{p}\right|^2 = \left|\mathbf{f}^T \mathbf{p}\right|^2 = \left|\cos\beta\cos\theta + \sin\beta\sin\theta\right|^2$$

Now, suppose, through a quirk of manufacturing, all the blue cells for an LCD display have polarization rotated 90 degrees with respect to all the green and red cells. Let's say that the blue polarization is horizontal ($\theta = 0^{\circ}$), and the green and red polarizations are vertical ($\theta = 90^{\circ}$). *Justify your answers* to each of the following questions (i.e., show your work). When giving a color as an answer, it is sufficient to just give the (R,G,B) coordinates.

a) (3 points) If we display solid white – (R,G,B) = (1.0, 1.0, 1.0) – and hold against the screen a polarizer in the vertical ($\beta = 90^{\circ}$) orientation, what color would we see? (We assume in this problem that each color channel ranges from 0.0 to 1.0, so (1.0, 1.0, 1.0) is bright white.)

b) (3 points) What color would we see if we then rotated the polarizer to a diagonal ($\beta = 45^{\circ}$) orientation?

c) (3 points) What color would we see if we then rotated the polarizer to a horizontal ($\beta = 0^{\circ}$) orientation?

d) (3 points) If we keep the filter oriented as in **c**), what color would we see if we now displayed solid green, i.e., (R,G,B) = (0, 1, 0)?

Problem 3: 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 4: 3D Affine Transformations (21 points)

The equation $\hat{\mathbf{n}} \bullet \vec{\mathbf{x}} = d$ describes the plane pictured below which has unit length normal $\hat{\mathbf{n}}$ pointing away from the origin and is a distance d from the origin (in the direction of the normal vector).

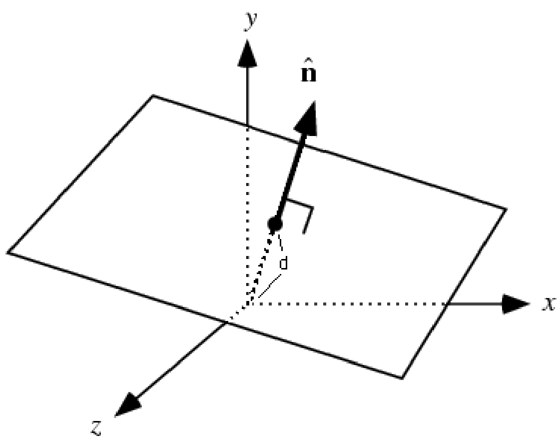

$$M_{xy} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Now consider a plane with normal lying in the y-z plane. The normal will have the form (0, sinθ, cosθ) for some θ. The equation for the plane is then $y\sin\theta + z\cos\theta = d$. Write out the product of 4x4 matrices that would perform a reflection across this plane. One of these matrices will be a reflection matrix; you must use the matrix M_{xy} above, which performs a reflection across the x-y plane. You must write out the elements of the matrices and the product order in which they would be applied, but you do not need to multiply them out. Justify your answer with words and/or drawings.

Problem 5: Hierarchical modeling (26 points)

Suppose you want to model the pincer with accordion joint illustrated below. The model is comprised of 8 parts, using primitives **A** and **B**. The model is shown in two poses below, with the controlling parameters of the model illustrated on the far right. The illustration on the right also shows a point P that the model is reaching toward, as described in sub-problem **c**).

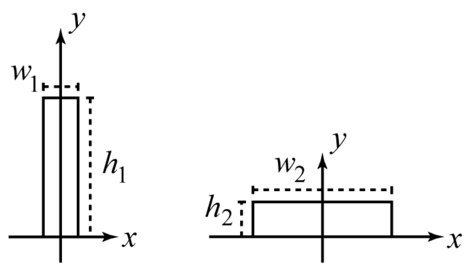

Primitive **A** Primitive **B**

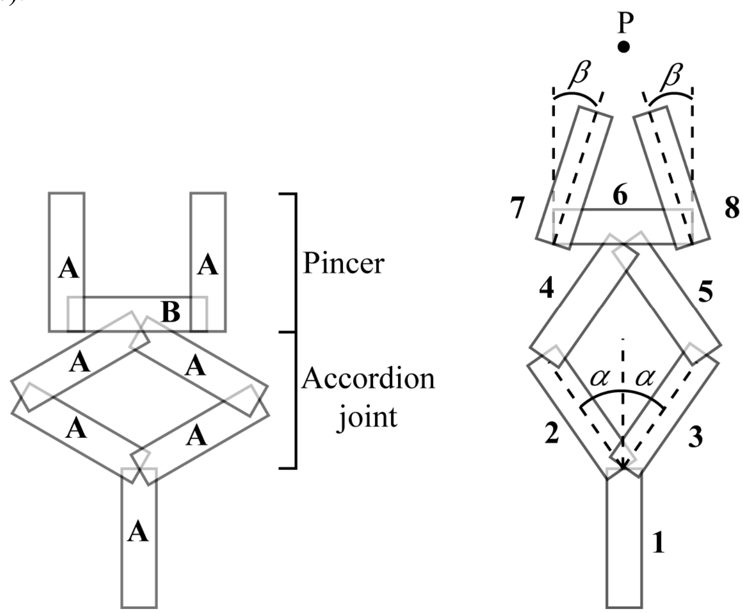

Assume that α and β can take values in the range [0, 90°]. Also assume that all parts use primitive **A**, except for part **6**, which uses primitive **B**. The model on the left shows the primitives used, the model on the right shows the enumeration (naming) of the parts.

The following transformations are available to you:

- R(θ) – rotate by θ degrees (counter clockwise)
- T(a, b) – translate by $\begin{bmatrix} a \\ b \end{bmatrix}$

a) (16 points) Construct a tree to describe this hierarchical model using part **1** 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 the matrices). Remember that the order of transformations is important! Show your work wherever the transformations are not "obvious." Your tree should conatin a bunch of boxes (or circles) each containing one part number (1…8); 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.

b) (3 points) Write out the full transformation expression for part **7**.

c) (7 points) Suppose the primitives are infinitesimally thin, $w_1 = h_2 = 0$, and have lengths $h_1 = 10$ and $w_2 = 12$. Assume that part **1** sits right on the origin in world coordinates. What would the α and β parameters have to be so that the model extends out and closes the pincer just enough to precisely grasp the point $P = \begin{bmatrix} 0 & 28 \end{bmatrix}^T$, in world coordinates. Show your work.