

Homework #1

Displays, Image Processing, Affine Transformations, Hierarchical Modeling

Assigned: Wednesday, October 12th

Due: Wednesday, October 19th
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 follow the Gilligan's Island rule**, *answer the questions on your own, and show your work.*

Please write your name on your assignment!

* **The Gilligan's Island Rule:** This rule says that you are free to meet with fellow student(s) and discuss assignments with them. Writing on a board or shared piece of paper is acceptable during the meeting; however, you should not take any written (electronic or otherwise) record away from the meeting. After the meeting, engage in a half hour of mind-numbing activity (like watching an episode of Gilligan's Island), before starting to work on the assignment. This will assure that you are able to reconstruct what you learned from the meeting, by yourself, using your own brain.

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 panel in front of the unpolarized backlight – you removed the entire assembly with polarizers, crystals, color filters, everything – and looked right 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 LCD 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 a linear polarizer. 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 \mathbf{u} and \mathbf{v} which are of non-zero length and not parallel to each other. Which of the following is true and which is false:

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

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

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

4. $\mathbf{v} \cdot \frac{\mathbf{u}}{\|\mathbf{u}\|} = \mathbf{u} \cdot \frac{\mathbf{v}}{\|\mathbf{v}\|}$

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

Problem 2: Image processing (25 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.

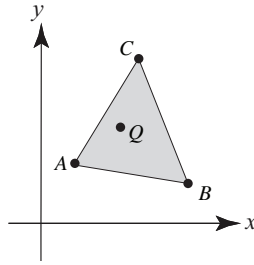
- a) (3 points) The image you're editing is too dark and noisy, and you decide you need to blur the image a little with a 3×3 filter to reduce noise, and amplify the overall brightness of the image by a factor of 4. Suggest a single convolution filter that perform this task when applied to the image. (Technically, pixel values could go out of range, i.e., brighter than 255 in one or more color channels at a pixel; assume that any needed clamping will be taken care of later, after filtering).
- b) (3 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) (3 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. (The origin is in the lower left corner.) Suggest a convolution filter that would accomplish some diagonal blurring along that direction by averaging across m pixels.
- d) (3 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) (7 points) We can think of sharpening as "removing blur" from an image. Suppose you take an image f and convolve it with a 3×3 mean filter. Now you multiply the resulting blurred image by α and subtract the result from the original image. You can think of the parameter α as controlling how much blur you remove. The resulting image will be a bit sharper, but also a bit darker, which can be corrected with suitable brightening. Solve for a single, normalized filter that performs all of these steps when convolved with the image, with α left as a variable in your filter. Note that the normalization will actually take care of the needed brightening. Show your work.
- f) (2 points) What filter do you get from part (e) if you set $\alpha = 1/2$? Try applying it using your Impressionist project. Does it sharpen?
- g) (4 points) Suppose you follow the same sharpening process as in part (e), but this time you blur the original image with the following filter:

$$\begin{bmatrix} 0 & 1/5 & 0 \\ 1/5 & 1/5 & 1/5 \\ 0 & 1/5 & 0 \end{bmatrix}$$

then multiply the result by $5/6$, subtract it from the original image, and fix the brightness. Solve for the (normalized) filter that will accomplish these steps all at once when convolved with the original image. Show your work. The result should resemble a filter we have discussed before — which one?

Problem 3: Triangle coordinates (24 points)

Consider triangle ΔABC 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 ΔABC , 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.

- (3 points) Using cross and/or dot products, devise a test (with the help of equations) to determine if point Q lies inside of ΔABC . 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.]
- (2 points) Using cross and/or dot products, compute the unit-length normal to ΔABC . 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.]
- (3 points) Using cross and/or dot products, compute the area of the triangle, $\text{Area}(\Delta ABC)$. 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 $\text{Area}(\Delta_{ABC})$ 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 (a weighted sum of affine points that yields an affine point), 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:

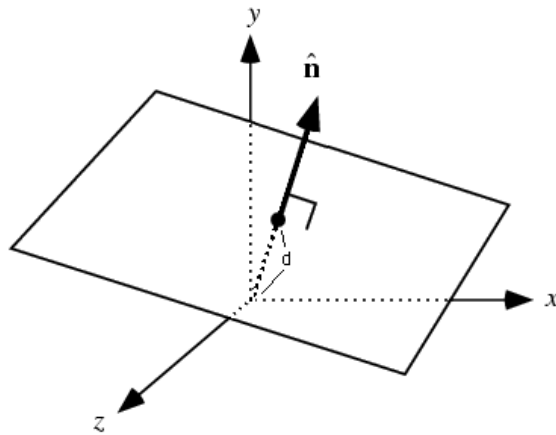
$$p_0 = \frac{\det \begin{bmatrix} r_0 & m_{01} & m_{02} \\ r_1 & m_{11} & m_{12} \\ r_2 & m_{21} & m_{22} \end{bmatrix}}{\det \begin{bmatrix} m_{00} & m_{01} & m_{02} \\ m_{10} & m_{11} & m_{12} \\ m_{20} & m_{21} & m_{22} \end{bmatrix}} \quad p_1 = \frac{\det \begin{bmatrix} m_{00} & r_0 & m_{02} \\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

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 (17 points)

The equation $\hat{\mathbf{n}} \cdot \bar{\mathbf{x}} = d$ 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). Any point $\bar{\mathbf{x}} = [x \ y \ z]$ on the plane must satisfy the plane equation $\hat{\mathbf{n}} \cdot \bar{\mathbf{x}} = d$.

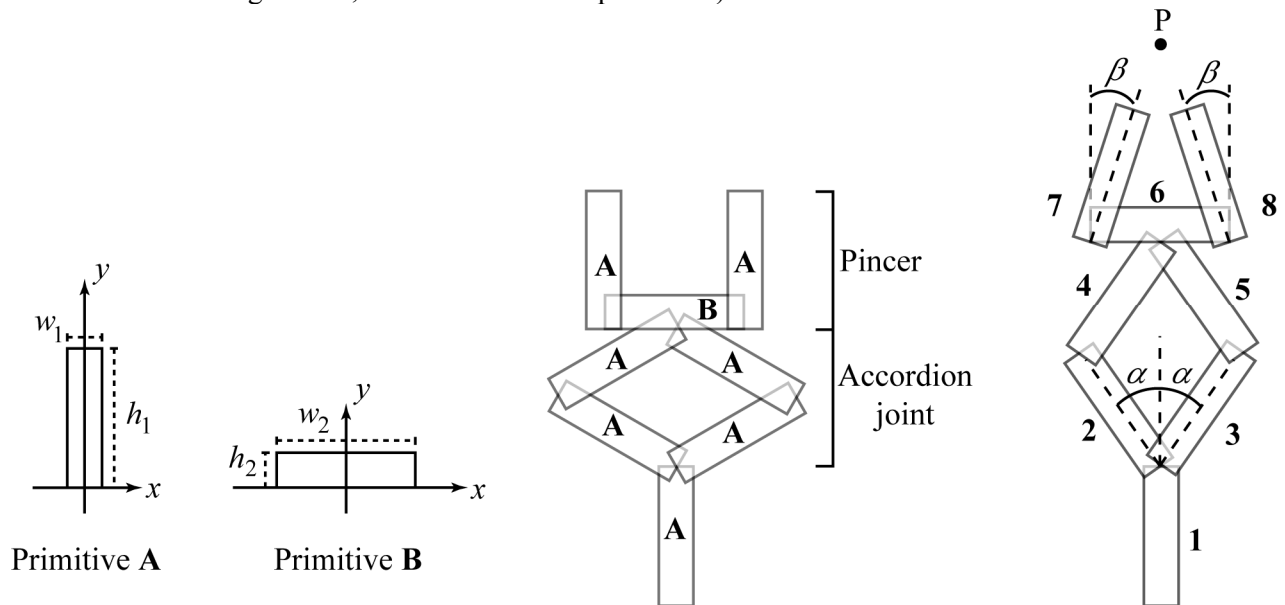


$$M_{xy} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Now consider a plane with normal lying in the y - z plane. The normal will have the form $(0, \sin\theta, \cos\theta)$ 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 (24 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).



Assume that α and β can take values in the range $[0, 90^\circ]$. 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:

- $R(\theta)$ – rotate by θ degrees (counter clockwise)

- $T(a, b)$ – translate by $\begin{bmatrix} a \\ b \end{bmatrix}$

- (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 contain 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. 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.
- (2 points) Write out the full transformation expression for part **7**.
- (6 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 $P = [0 \quad 28]^T$, in world coordinates. Show your work.