Final Exam Solutions

1. **Vector Noise Filtering: (6 points)**

   (a) Vector median filtering: The distances between all pairs of vectors are as follows: $d(1,2) = |3-2.5| + |1.5-2| = 1$, $d(1,3) = 1.5$, $d(1,4) = 2+0.5 = 2.5$, $d(1,5) = 2+1.5 = 3.5$, $d(2,3) = 0.5+1 = 1.5$, $d(2,4) = 1.5$, $d(2,5) = 2.5$, $d(3,4) = 3$, $d(3,5) = 2$, $d(4,5) = 1$.

   With these distances, we can compute, for each vector, the sum of the distances to all the other vectors: $D_1 = \sum_{j=2}^{5} d(1,j) = 1 + 1.5 + 2.5 + 3.5 = 8.5$
   $D_2 = 1 + 1.5 + 1.5 + 2.5 = 6.5$
   $D_3 = 1.5 + 0.5 + 3 + 2 = 7$
   $D_4 = 2.5 + 1.5 + 3 + 1 = 8$
   $D_5 = 3.5 + 2.5 + 2 + 1 = 9$
   So the vector median is x_2.

   (b) The output of the vector directional filter is x_3 since that is the vector in the middle.

   Grading: 3 points for part (a) and 3 for (b). 1 point off if there's a math error.

2. **Interpolation: (5 points)**

   We label two pixels P and R as shown:

   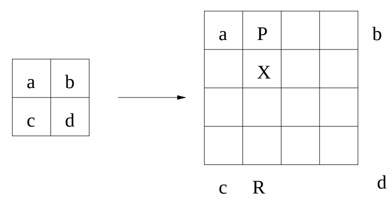

   By linear interpolation between a and b, we know that

   $$P = \frac{3}{4}a + \frac{1}{4}b$$

   Likewise, linear interpolation between c and d yields:

   $$R = \frac{3}{4}c + \frac{1}{4}d$$

   Lastly, we interpolate linearly between P and R to get:

   $$X = \frac{3}{4} \times P + \frac{1}{4} \times R = \frac{9}{16}a + \frac{3}{16}b + \frac{3}{16}c + \frac{1}{16}d$$

3. **Compactness: (6 points)**

   (a) For any shape or region, compactness is defined as:

   $$\gamma = \frac{perimeter^2}{4\pi(Area)}$$

   Your textbook defines it on page 822 without the 4π normalization factor, so either response is OK for this exam. For a square of side length a, the compactness is:

   $$\gamma_{sq} = \frac{(4a)^2}{4\pi a^2} = \frac{4}{\pi}$$

(b) For a rectangle with side lengths a and b, the compactness is:

$$\gamma_{rect} = \frac{(2a+2b)^2}{4\pi ab} = \frac{(a+b)^2}{\pi ab}$$

(c) Comparing these two, we see that

$$\gamma_{rect} > \gamma_{sq}$$

because, for any positive a and b that are not equal, we have

$$(a-b)^2 > 0 \quad \Rightarrow \quad a^2 + b^2 + 2ab > 4ab \quad \Rightarrow \quad (a+b)^2 > 4ab \quad \Rightarrow \quad \frac{(a+b)^2}{\pi ab} > \frac{4}{\pi}$$

Grading: 2 points for each part. Part (c) has 1 point for stating the relation, and another for proving it.

4. **Hurst Transform: (6 points)**

Negative slope: NO. In plotting log(range) vs. log(distance), the range is monotonically non-decreasing, since all pixels examined at closer distances are included at larger distances. The minimum possible slope is zero. This will happen if the range is constant at all distances.

Negative intercept: YES. Although not a very common occurence, it is possible for the range to increase more than exponentially with increasing distance. In this case, the log plot will be bent upwards, and the best fit straight line could have a negative intercept. In class we did not discuss *how* the straight line was to be fit to the points; we merely said that a straight line is fit. But still it should be fairly clear that it is possible that a straight line fit to an upward curve might have a negative intercept. For example, consider the 5×5 block of pixels shown here on the left. The points and the best fit line (fit using least squares regression) are shown on the right. The slope is 5.3991, and the intercept is -0.2532.

255	80	20	80	255
80	4	0	4	80
20	0	0	1	20
80	4	0	4	80
255	80	20	80	255

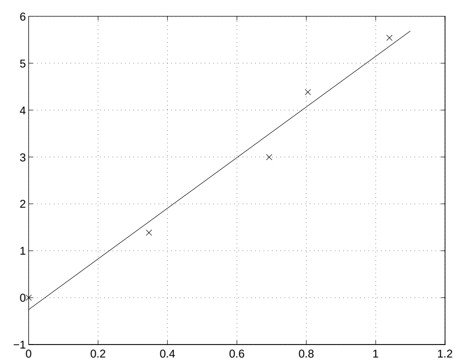

In fact, because of the powerful effect an outlier can have on a regression line, you can even squeak out an intercept of -1.1 for the center pixel using a block of pixels like this:

255	0	0	0	0
0	0	0	0	0
0	0	0	0	0
0	0	1	0	0
0	0	0	0	0

Grading: 3 points for each part. 1 point counts for stating YES or NO correctly but with no reason provided. 2 additional points for correct explanation.

5. **Quadtree Segmentation: (5 points)**

   We need to establish what order the four quadrants of a given block are represented in the quaternary tree. We will use the following order for the branches going from left to right:

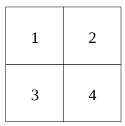

   The image as a whole does not have all the pixels of the same intensity, so we have to subdivide it into four quadrants. That is the first level of the tree. One of those four quadrants (the 3rd one) has all of its pixels the same color, so that branch is done. The other 3 branches each get subdivided into 4 quadrants. There are two places where further subdivision is needed. The final quadtree looks like this:

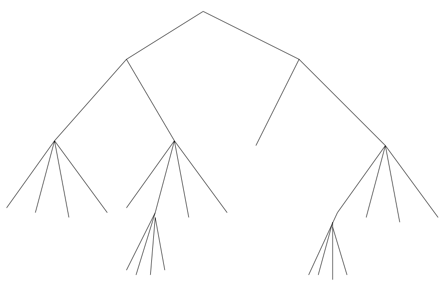

6. **Hough transform: (6 points)**

   (a) This is textbook problem 10.23. Point 1 has coordinates x=0 and y=0. Substituting this into the equation for the normal representation of a line:

$$x \cos\theta + y \sin\theta = \rho$$

   yields $\rho = 0$, which, in a plot of ρ versus θ is a straight line.

   (b) Only the origin (0,0) would yield this result when substituted into the equation.

7. **Histograms after Filtering: (6 points)**

   i) For both the median filter and spatial averaging filter, the input is always three 0's, three 0.5's and three 1's, and the output of both will always be 0.5. So the answer is histogram (d) for both of the filters.

   ii) Histogram (a) for median filter. Because the stripes are wide, this is going to be a root signal for the median filter, so the output histogram is the same as the input histogram.

   Histogram (f) for spatial averaging. When the filter is in the homogenous region (with relatively high probability), the output is either 0, 0.5 or 1, so these are the tall bins. When the filter straddles a stripe boundary, we get an intermediate value. The intermediate bins are not equal height, because there are two filter positions which give you 1/3 and only one which gives you 1/6.

iii) (c) for median filter. The probability that the output equals 0 is the same as the probability that it equals 1, and this is lower than the probability of the output being equal to 0.5.

Histogram (h) for the spatial averaging filter. There are 19 different possible outputs (0 to 1 with a step size of 1/18), and the distribution is like a binomial distribution.

Grading: 1 point for each part. No explanations needed.

8. **Binary Dilation: (6 points)**

   (a) Set B_1 must have at least a side length of a in order for the dilation to have no interior hole. The resulting dilated set would be a square with side length 3a, as shown below left. The original set A is shown with dashed lines for reference.
   (b) The disk must have a diameter of at least a. The dilation would be a square with rounded corners with side lengths 3a, as shown below middle. The rounded corners would be such that the circle exactly fits, as shown below right. All four corners are the same.

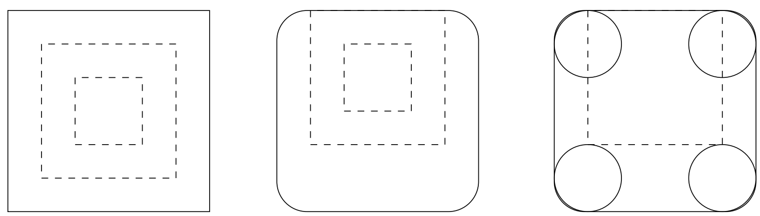

   Grading for this problem: (a) 1 point for saying smallest size is a = side length. 1 point for drawing the correct shape = square. 1 point for giving the dimension of the square = 3a. (b) same breakdown as for part a. It doesn't count as the correct shape if the bottom two corners are different from the top 2 corners (all corners should be the same).

9. **Distance Transformation: (6 points)**

   To get the desired result, we have to initialize this such that all the pixels of the background have value 0, and all the pixels inside A have some suitable large value (say 100). Then the results of the forward pass and subsequent backward pass of the Chamfer3-4 algorithm are as follows:

			3							
3	3	3	4	3	3	3	3	3	3	3
3	6	6	7	6	6	6	6	6	6	4
3	6	9	10	9	9	9	9			
3	6			3	6	9	4			
3	4			3	6	8	4	3	3	3

			3							
3	3	3	4	3	3	3	3	3	3	3
3	6	6	6	6	6	6	4	3	3	3
3	4	3	3	4	7	6	3			
3	3			3	6	6	3			
3	3			3	3	3	3	3	3	3

10. **Detecting gaps in lines: (6 points)**

    This is problem 10.2 from the textbook.

    I wasn't aiming for an answer that uses the hit-or-miss transformation, but I'll accept an answer that uses it if you did it correctly. With the hit-or-miss transformation, the masks are:

$$
\begin{array}{ccc} 0 & 0 & 0 \\ 1 & 0 & 1 \\ 0 & 0 & 0 \end{array} \qquad
\begin{array}{ccc} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{array} \qquad
\begin{array}{ccc} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{array} \qquad
\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{array}
$$

For example, the first one would find a place where there is a gap in a horizontal line. If there is a "hit" we consider that a 1-pixel gap has been detected in the center location.

There are some configurations where these masks won't work perfectly. For example, if we had a horizontal line crossing a vertical line just above a place where the vertical line has a break, the masks listed would not find that break. Additional masks or "don't care" states could be added to help deal with various problems of crossing lines.

The approach I was looking for does a convolution with the following 4 masks:

$$
\begin{array}{ccc} 0 & 0 & 0 \\ 1 & -2 & 1 \\ 0 & 0 & 0 \end{array} \qquad
\begin{array}{ccc} 0 & 1 & 0 \\ 0 & -2 & 0 \\ 0 & 1 & 0 \end{array} \qquad
\begin{array}{ccc} 0 & 0 & 1 \\ 0 & -2 & 0 \\ 1 & 0 & 0 \end{array} \qquad
\begin{array}{ccc} 1 & 0 & 0 \\ 0 & -2 & 0 \\ 0 & 0 & 1 \end{array}
$$

Then each mask would

- yield a value of zero if centered on an empty space, or if centered on a pixel of an unbroken 3-pixel segment oriented in the direction favored by the mask
- yield a value of -2 when centered on a pixel of an unbroken 3-pixel segment oriented in a direction NOT favored by the mask
- yield a value of $+1$ when placed next to a line of a different direction
- yield a value of $+2$ when centered on a one-pixel gap in a 3-pixel segment oriented in the direction favored by the mask

So the results can be thresholded at $+2$. The method will be confused in various cases (e.g., two parallel lines with a 1-pixel wide gap between them). Again, more complicated masks, or additional masks could be used that would help distinguish between some of these cases.

11. **Registration: (6 points)**

(a) Under the SAVD error metric, the new image will line up with the first one such that: 25 $\leftrightarrow$ 0 and 60 $\leftrightarrow$ 50 and 75 $\leftrightarrow$ 100. Here I am using the symbol $\leftrightarrow$ to mean that the regions are in spatial correspondence, that they line up with each other. It is a unique rotation that works.

(b) Now there is not a unique minimum. If we have 0 $\leftrightarrow$ 0 and 50 $\leftrightarrow$ 50, then we will be stuck with 0 $\leftrightarrow$ 100 which produces an error of 100. Let's call this case 1. If instead we had 0 $\leftrightarrow$ 0 and 0 $\leftrightarrow$ 50 and 50 $\leftrightarrow$ 100, we also get an error of 100 (2×50). Let's call this case 2. Any intermediate position between case 1 and case 2 also produces an error of 100.

(c) With the MSE, case 1 is really bad, because we now get 100^2. Case 2 is not as bad, since we just get 2×50^2. So, with the MSE, we will back to a unique rotation that works, which is case 1.

12. **Clustering: (5 points)**

    For the top row, left to right, the answers are YES, NO, YES. For the bottom row: YES, NO.

    The two NO cases are really the same case. The are both NO because, although the cluster centers are in fact the centroids of the data points mapping into them, the decision boundary is NOT the perpendicular bisector of the line segment which connects the two cluster centers. The other 3 cases which are YES have both conditions met: the cluster centers are the centroids of the data points mapping into them, and also the decision boundary is the perpendicular bisector of the line segment which connects the two cluster centers.

    Grading: 1 point each, no explanation needed.

13. **Color: (6 points)**

    For the RGB group: The middle one is B because it has large values in the blue sky/water area. The one on the right is R, because it has high values in the flowers, but low values for the rest of the foliage (which is green). So the one on the left is Green.

    For the HSI group: The left one is S, because the white flowers in the original image (which are highly unsaturated) are the darkest things in this picture. The middle one is I, because that looks like a normal grayscale version of the image. The right is H, which one can say by process of elimination, or else by using the fact that blue hue = 2/3, and green is 1/3, and we can see that the water/sky region is a medium gray, whereas most of the hillside is a darker gray. Also remember that the zero for hue is red, so the pink flowers are both very dark (min=zero) and very bright (max=one) splotches on the hue scale, depending on which side of pure red they're on.

    Grading for this problem: 1 point for each sub-part, no explanation needed.

14. **Texture analysis with autocorrelation: (5 points)**

    The two periodic ones are the easiest. Patch A is a tight weave, so corresponds to the bottom left ACF with lots of bumps. Patch B is a coarse weave so is the ACF in the middle which only has a few bumps.

    Patches C, D, and E are all aperiodic. Of these, D has the finest texture, and it corresponds to the top left ACF. Patches C and E coarser, more slowly varying textures. C is top right ACF and E is the bottom right ACF. These are hard to tell apart... perhaps the only thing one can say is the ACF in the bottom right is more bumpy or granular, which corresponds to the seeds of E.

    For grading: writing C when it's actually E is half credit and likewise writing E when it's actually C is half credit, since those are fairly close. Otherwise, one point for each, no explanation needed.