ECE 253a MIDTERM EXAM SOLUTIONS

1. **Order of Operations**

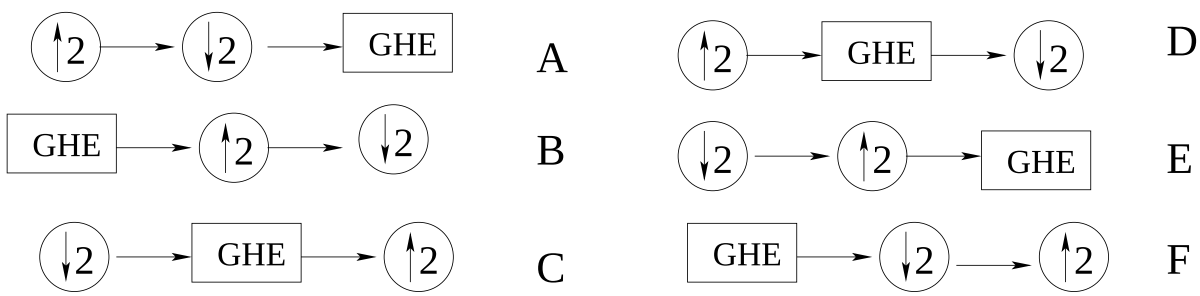

(a) When you start with a sampled image, and you wish to "interpolate" it to get a continuous image field, in class we distinguished between *approximators* and true *interpolators*. The true interpolators are guaranteed to give you back the original sampled image if you re-sample the continuous image field onto the original sampling grid. The approximators might not. Pixel replication and bilinear interpolation are both interpolators, so the upsampling followed immediately by downsampling is an identity operation. The upsampling preserves the exact pixel values of the original image (spacing them out by a factor of 2 in each direction) and interpolates values for the in-between positions. The downsampling will remove those interpolated in-between positions, and just give back the original pixel values. So, we can conclude that systems A and B just collapse down to being GHE alone (output equals Y).

When you upsample using pixel replication, the shape of the histogram does not change. There will be 4 times as many pixels in the image, but the probability mass function is unchanged. That means that the CDF (cumulative distrubution function) which is the GHE remapping function is also unchanged. So, system D, which has pixel replication before GHE and then has downsampling, will also yield output image Y.

If you downsample before the GHE, you have, in the general case, changed the histogram, and therefore the GHE remapping operation. So system C and system E will not yield output image Y, however, systems C and E will have the same output as each other, call it Z.

Lastly, system F will produce something which is not Y, but is also not Z.

(b) The answer here is different, because (unlike with pixel replication) the histogram after upsampling with bilinear interpolation is not identical to the histogram of the starting image, so the GHE will be different.

Grading for this problem: 1 point each for saying that systems A, B and D produce Y. 1 point for saying that systems C and E produce the same thing as each other. 1 point for saying that system F is different from all of the above. 1 point for part (b).

2. **Color: Hue and Saturation**

(a) When you blur the hue, and then transform back to RGB, the image looks like this:

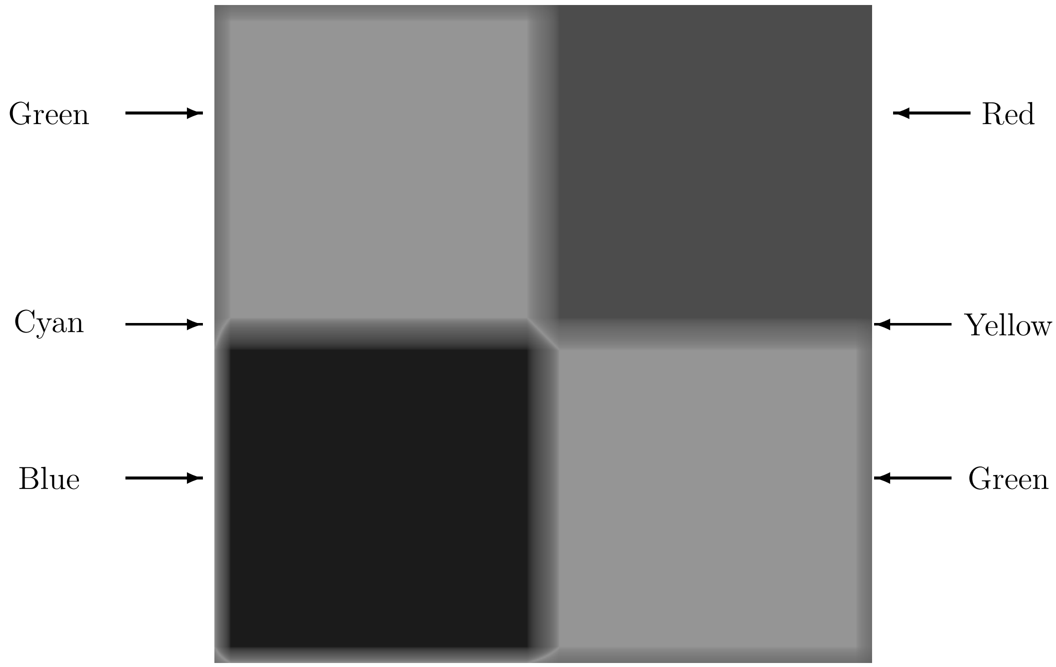

Between the green and blue regions, there is a strip of width 24 pixels (that is, 12 pixels on each side of the center line) which passes in color from green, through cyan, to blue. Between the red and green regions, there is a strip of width 24 pixels which passes in color from red, through yellow, to green.

In the center of the image, there is a region where the filter will be encountering all 4 quadrants, and so the result there is going to be a bit more complicated.

Near the boundary of the image, if the image is zero-padded prior to filtering, then there are going to be altered colors at the boundary as well. That is the situation shown in the figure above. But if the boundary were handlded by replicating the first/last row/column outward as many times as needed for the filtering, then there would be no altered color effects at the boundary. In any case, you were not expected to mention the boundary.

(b) Lastly, when you blur the saturation, there is no effect, because all the 4 quadrants are equally saturated to start with.

Grading for this problem: 2 points for getting the color right in part a (that is, saying that it would be a progressive shading through yellow, and through cyan), 1 point for getting the width right (24 pixels, that is, 12 on each side of center), 1 point for mentioning that something more complicated would happen right in the middle of the picture because all 4 quadrants would enter into the filtering, and 2 points for part (b).

3. **Skeletons**

One has to picture putting a ball (or disk) of largest possible size inside the object, and picture moving the disk around (shrinking it if necessary), to keep it just barely fitting inside. The skeleton is the set of center points of these maximal balls.

The barbell, line segment, and rectangle with rounded ends all have the same skeleton, namely the line segment. The disk's skeleton is a point, and the ring's skeleton is a circle. The square goes to the X-shape, and the rectangle to the sideways stick figure. The object consisting of three touching disks does not have a skeleton shown. The maximal disks that one could put in there would be exactly three disks, and so the skeleton would consist of 3 points, at the centers of the 3 disks.

Grading for this problem: 3/4 of a point for each of the 8 objects.

4. **Histogram Equalization**

By inspection, the average height of the bins in the image histogram is 750. But, if you look at most local regions (say, the bins from 10 to 20, or the bins from 90 to 100) the average is far from 750. After equalization, the average height will still be 750, but if you look in local regions, the average of 750 will be achieved.

At the low end, this is obtained by spacing out 3 tall bins (height 1000) with one empty bin (height 0) so the average of any consecutive group of 4 bins is 750. This will occupy all the bins from 1 to about 68 (which is about (50/3) x 4).

At the high end, a short bin of height 500 will alternate with a tall bin created by stacking two short bins on top of each other (height 1000) therefore the average of any consecutive group of 2 bins at the high end will be 750.

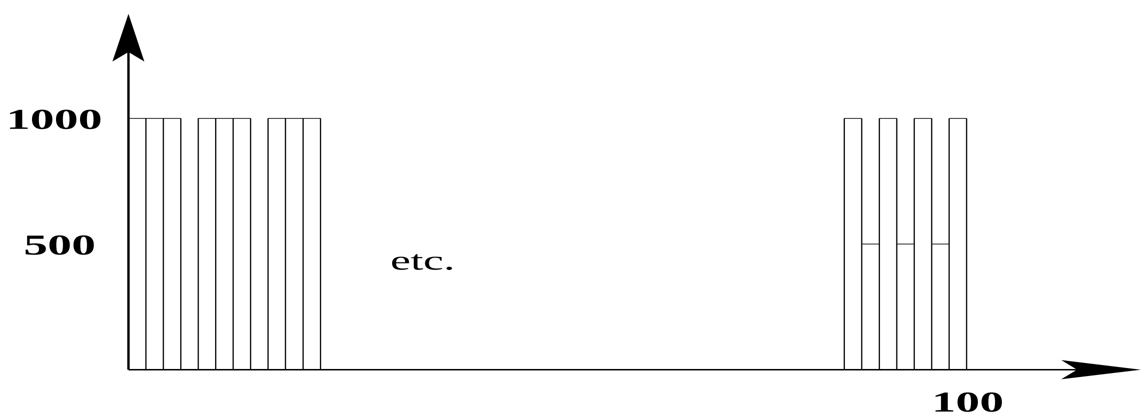

Grading for this problem: 1 point for getting average height (750) right. So if you mistakenly treat this as a continuous pdf and say the output histogram will be flat at 750, then you get only 1 point. 2 points more for getting the left-hand spacing right (bins spaced out with 3 tall ones and 1 empty) and 2 points for getting the spacing correct for the right-hand side (stacked and unstacked bins alternating). 1 more if you get the x position right where the transition is (about 68).

5. **Noise Reduction**

Image A is from the small spatial averaging filter (item 3). We can see this because each little white square is blurred out into a little white blob.

Image E is from the large spatial averaging filter (item 4). We can see this because the whole thing is very blurry.

All the others are some kind of median filtering (MF), because they still have some sharp edges.

Image B is from the 3x3 MF. We can tell this because there is still quite a bit of noise left, and the 3x3 MF is big enough to clean up the noise from a 2x2 square only if there is no adjacent 2x2 square which is also noisy. So any time there are adjacent noise squares, the 3x3 will leave some salt noise.

Image D is the sparse MF (item 6). It is intended for this kind of clustered noise.

Image C is from repeated use of the small MF (item 5). We can tell this because there are some moderately small structures (e.g., black bars on the plane's tail, or features in the mountains) which are not completely obliterated by the filtering.

Image F is from the large MF (item 2). Some people were unsure about C and F (switched them). One way to think about this is that if you have a structure like a very long bar of width 3 (so it is long in one direction but narrow in the other) then locally (in the middle of the bar) that object is unchanged by a 3x3 MF but will be obliterated by a 9x9 MF. So the 3x3 MF is able to preserve smaller structures than a 9x9 MF. If something is a root signal for a 3x3 MF, then it will be preserved, no matter how many times the 3x3 filter is applied.

Grading for this problem: 1 point for each picture.