Name____________________________

Computer Vision	**Neel Joshi, Ira Kemelmacher, Ian Simon** **TAs: Jiun-Hung Chen, Rahul Garg**
CSE 455	**Winter 2010**

Midterm

Handed out: Friday, Feb 5
Due: Friday, Feb 12, at the beginning of class
Late exams will not be accepted

Directions: Please print out the exam and provide answers to the questions in the space provided. This take-home midterm is open book and open notes, but do not discuss it or collaborate with other students. In addition, the instructors and TAs will not answer questions about the exam, other than for simple clarifications, which we will then post to the entire class.

Make sure you turn in all 14 pages.

Put your name at the top of every page.

Name____________________________

Problem 1: Reconstruction Filters (15 pts)

Consider a sampled 1D signal f, where f is nonzero only at integers. To resample this signal, we can first convolve it with a reconstruction filter h, yielding a continuous function g. We can then take arbitrary samples of g. In class we looked at linear interpolation using a hat function as the reconstruction filter. In this question we will look at other possible reconstruction filters.

Part 1. Instead of using the hat function for reconstruction, consider using a box function:

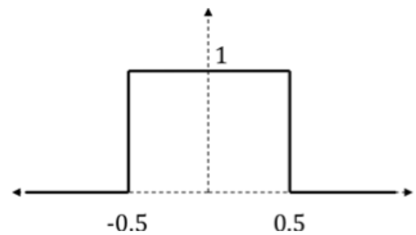

Draw the continuous function that results from applying this reconstruction filter to the following sampled data:

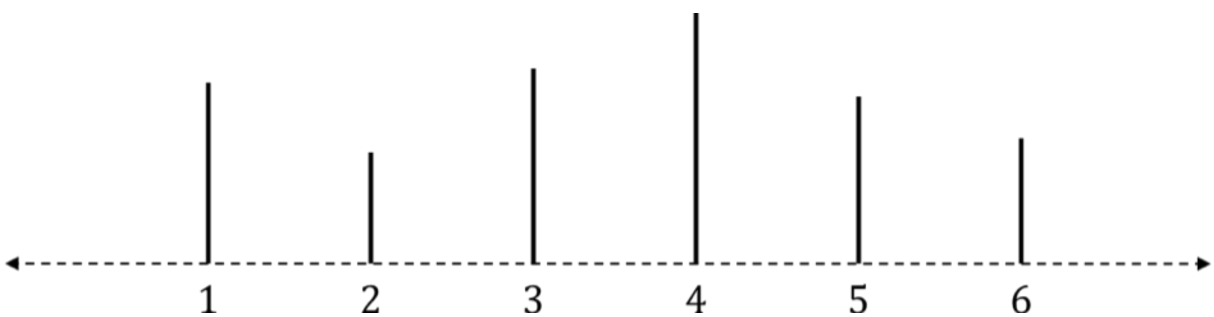

When computing reconstructed function $g = f * h$, suppose $i \leq x < i + 1$, where i is an integer. Describe how to compute $g(x)$ as a function of $f(i)$, $f(i + 1)$, and α.

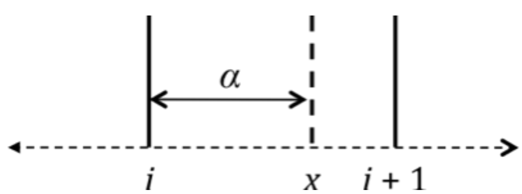

What kind of interpolation is this?

Name____________________________

Part 2. Now consider using a stretched-out hat function for reconstruction:

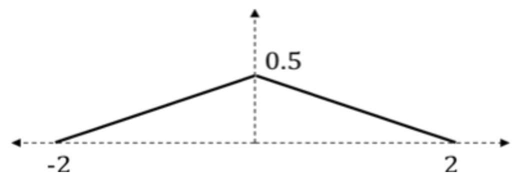

When computing reconstructed function $g = f * h$, suppose $i \leq x < i + 1$, where i is an integer, and $x = i + \alpha$. How many samples of f are needed to determine $g(x)$? Give a formula for $g(x)$ as a function of these samples and α.

If x is an integer, under what circumstances does $f(x) = g(x)$ as reconstructed by the stretched-out hat function? Under what circumstances does $f(x) = g(x)$ if we reconstruct using the ordinary hat function?

Problem 2: Directional Derivative (5 pts)

Given an image I, we can compute its derivative in the x or y direction using a Sobel filter. Suppose that we want to compute the derivative in an arbitrary direction (u,v). Give a 3-by-3 filter that computes this derivative, in terms of u and v. Assume $u^2 + v^2 = 1$. Hint: recall that the derivative in direction (u,v) is the dot product of vector (u,v) with the gradient.

Name____________________________

Problem 3: Convolution (15 pts)

Part 1. Describe the effect of convolving an image with each of the following filters, using concepts like blurring, derivatives, etc. where applicable:

$$\frac{1}{4}\begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\frac{\sqrt{2}}{4}\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix}$$

$$\begin{bmatrix} 0 & 0 & 0 \\ -1 & 3 & -1 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\frac{\sqrt{2}}{2}\begin{bmatrix} 1 & 0 & 1 \\ 0 & -4 & 0 \\ 1 & 0 & 1 \end{bmatrix}$$

$$\frac{1}{2}\begin{bmatrix} 0 & 1 & 0 \\ 0 & 2 & 0 \\ 0 & -1 & 0 \end{bmatrix}$$

Name___________________________

Part 2. Give a filter for which convolution of an image with that filter will yield each of the following effects, or state that it cannot be achieved with convolution:

a) Detect a 3-by-3 pattern in an image with intensities between 0 (black) and 1 (white). Your filter should give the maximum response for this pattern:

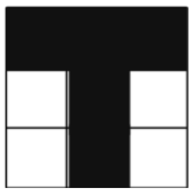

b) Simulate the effect of convolution with two arbitrary 2-by-2 filters in succession:

a	b
c	d

e	f
g	h

c) Replace each pixel with the brightest pixel in its 3-by-3 neighborhood.

d) Subtract from each pixel the mean intensity in its 3-by-3 neighborhood.

e) Shift an image to the right by 3 pixels.

f) Rotate an image clockwise by 90°.

Name____________________________

Problem 4: Old-School Special Effects (10 pts)

Director Dolly Spielberg loves Alfred Hitchcock movies and wants to film a movie that consists of numerous Dolly-Zoom shots. She sets up a scene with two colored boxes at two different depths, and starts shooting. The first frame of the shot looks like the following image:

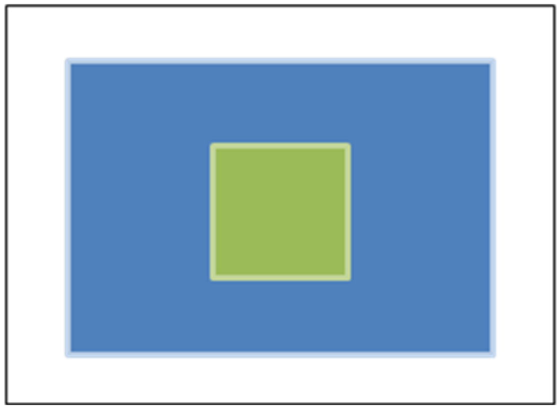

As the camera rolls she changes the camera depth and focal length. The following images are a few selected frames from her shot. For each one indicate whether the camera position is further, closer, or is unchanged, and focal length increased, decreased, or is unchanged when compared to the first image shown above.

a)

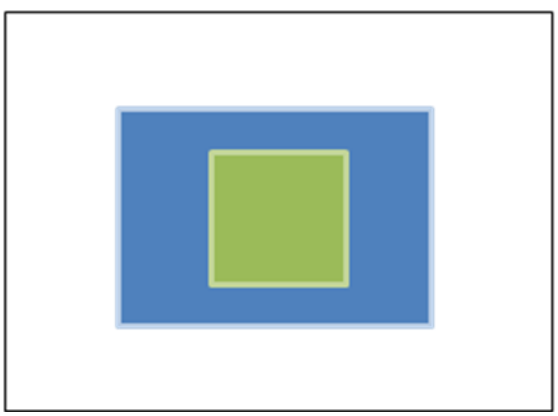

Camera Depth (Circle One)		
Further	Closer	Unchanged

Focal Length (Circle One)		
Increased	Decreased	Unchanged

Name_______________________________

b)

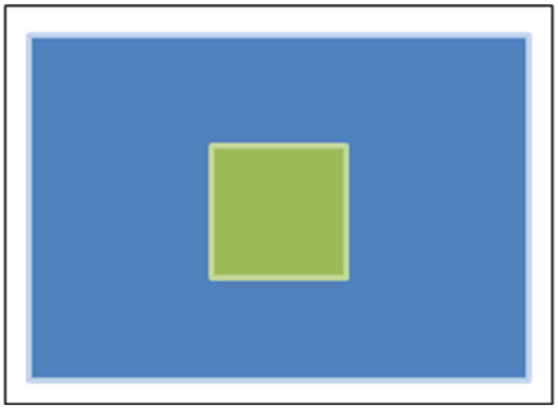

Camera Depth (Circle One)		
Further	Closer	Unchanged

Focal Length (Circle One)		
Increased	Decreased	Unchanged

c)

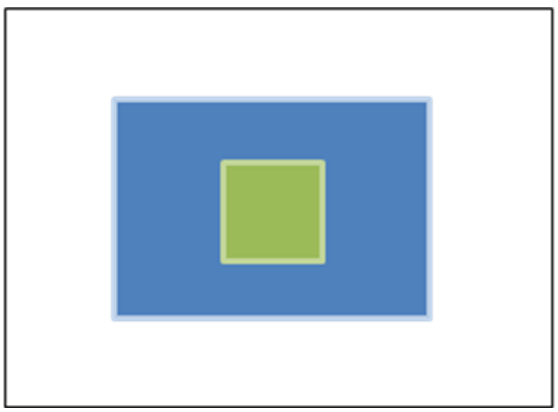

Camera Depth (Circle One)		
Further	Closer	Unchanged

Focal Length (Circle One)		
Increased	Decreased	Unchanged

d)

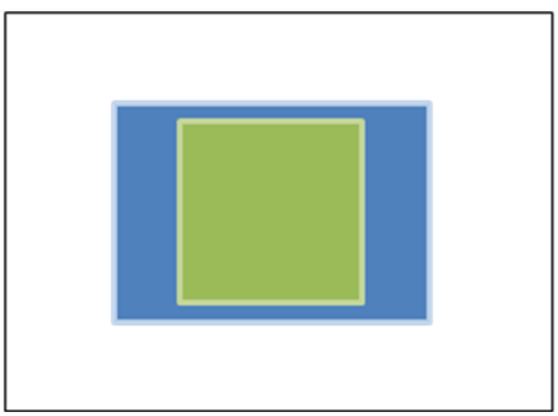

Camera Depth (Circle One)		
Further	Closer	Unchanged

Focal Length (Circle One)		
Increased	Decreased	Unchanged

Name___________________________

Problem 5: Optical Design (15 pts)

Dr. Mike R. Scope is building a very special instrument for his Biology lab. The device needs a lens system with very specific properties. Help Dr. Scope pick a lens and figure out the design of his new scope.

The device should have a single lens that images a slide with sample on it

The lens will project light onto a screen 5m away, and at the distance of 5m, the magnification should be 10X, as shown in the diagram below:

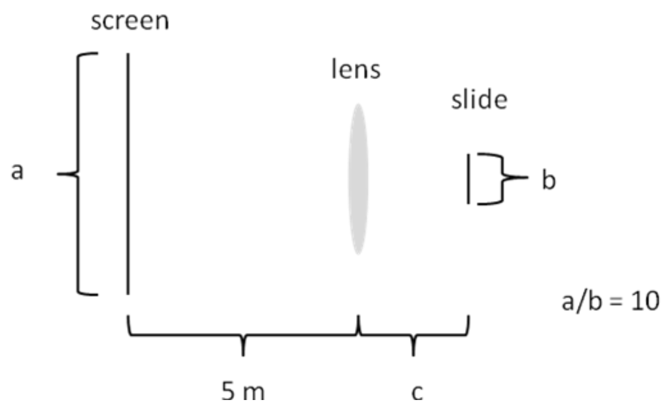

a) How far behind the lens should the slide be mounted (i.e., what is the value for *c*)? Show your work.

b) What focal length lens should he use? Show your work.

Name____________________________

Given recent budget cuts at Dr. Scope's University, he can't afford a custom focal length lens. Fortunately, he has a gift certificate from his mother for *Cheap Joe's Optical Supply* that only sells lens in multiples of 10mm focal lengths. He still wants 10X magnification at 5m from the lens, but is flexible on other aspects of the system.

c) What will he need to change from the initial design?

d) When Dr. Scope goes to Cheap Joe's he finds that the store only sells lenses that are either twice the diameter or half the diameter of the lens he originally wanted. Advise Dr. Scope on the advantages and disadvantages of his two options. Keep in mind that Dr. Scope sometimes puts samples on his slides that are not flat and because of the budget cuts, he has to keep his power bill down and can only illuminate the slide with a dim light.

Lens Twice the Diameter:

Lens Half the Diameter:

Name____________________________

Problem 6: Homographies (15 pts)

Part 1. Specify the homography (a 3-by-3 matrix) corresponding to each kind of warp:

a) A translation: $x' = x + u$; $y' = y + v$

b) A mirror reflection about the line $x = a$, as shown below. It may help if you think of this as a sequence (multiplication) of simpler operations. Show your work.

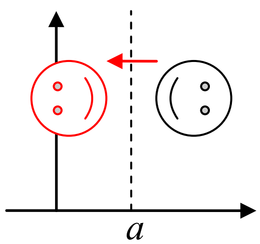

Part 2. Homographies will often fix one or more points in the image. That is, the point will stay in the same place after the homography has been applied. For each of the following transformations, specify which points, if any, are fixed, including any fixed points at infinity. Two examples are given for you. Ignore the point (0, 0, 0).

A 2D rotation
finite fixed points: **the center of rotation**
infinite fixed points: **none**

Scaling the image
finite fixed points: **the image center (0, 0, 1)**
infinite fixed points: **all points at infinity (u, v, 0)**

a) A 2D translation by *(u, v)*

finite fixed points:

infinite fixed points:

b) A mirror reflection about the line $x = a$

finite fixed points:

infinite fixed points:

Name___________________________

Problem 7: Features (15 pts)

Part 1. In class, we saw that the Harris corner detector finds points that are local maxima of the Harris operator. The SIFT feature detector finds features that are maxima (or minima) of a different function. SIFT first convolves the image with a *difference of Gaussians* (DoG) filter, which is similar to the Laplacian of Gaussian filter described in class. SIFT then finds maxima and minima of the convolved image. In 1D, the DoG filter looks like this (note that some values are negative):

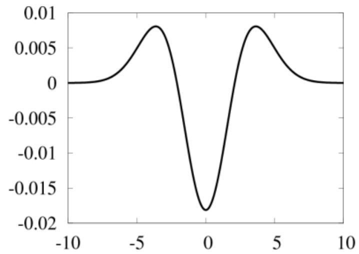

a) Consider the three images below labeled Image A, Image B, and Image C. This question asks you to apply the 1D DoG filter to each image at the position $x = 0$, and determine the sign of the resulting number (called the *response* of the filter). For which images is the filter response at x=0 positive? For which images is it negative? Check one of the boxes for each image; the correct box for Image A has already been checked for you. (The DoG filter has been replicated above each image to help you visualize the convolution). Note that the x-axis of each 1D image represents position, and the y-axis represents intensity (an intensity of 0 indicates black, and an intensity of 1 indicates white).

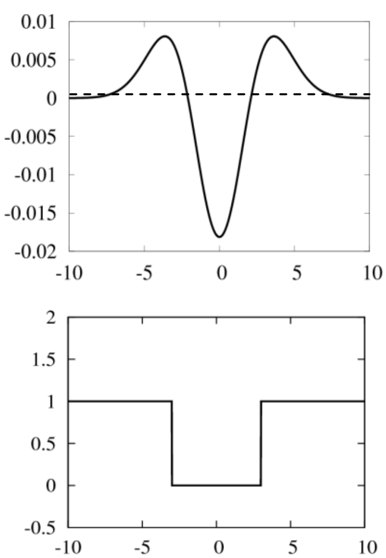

Image A

☒ Positive
☐ Negative

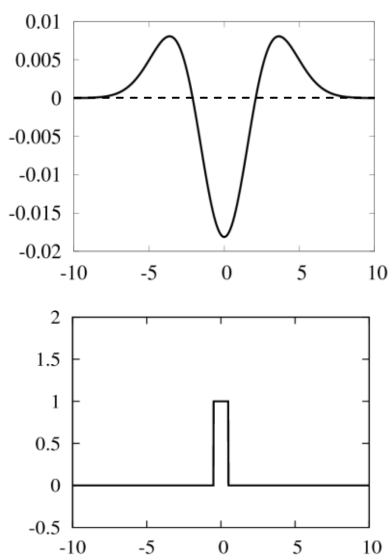

Image B

☐ Positive
☐ Negative

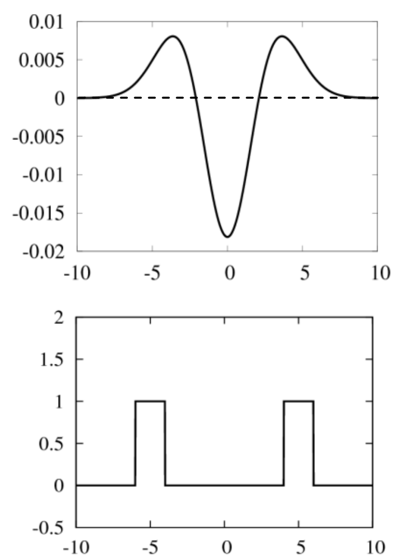

Image C

☐ Positive
☐ Negative

Name____________________________

b) In the space provided below, draw the 1D image (with intensity values in the range 0 and 1) which, when convolved with the DoG filter, will have the *maximum possible* response at x=0 (that is, greater than that of any other image). This will be the type of image feature that the 1D DoG detects most strongly.

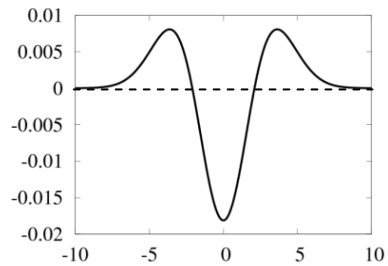

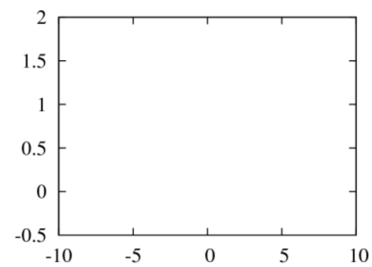

c) The 2D DoG filter is radially symmetric filter centered on (0,0), and looks like this:

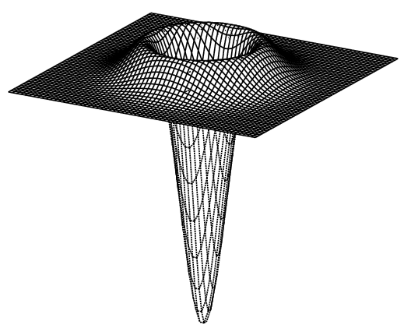

The intersection of any vertical plane passing through the center of the filter looks like the 1D DoG filter. Generalizing from the 1D case in part (c), what kind of 2D image feature would result in the most positive value when this feature is applied? Describe your answer in words or with a picture.

Name____________________________

Part 2. When creating a descriptor for a detected feature point, a common approach is to sample the image in a window of some size around the feature's position. One decision that needs to be made is how big this window should be, as this can affect how many correct matches will be found.

Consider the two images of a whiteboard below, Image A and Image B, containing text, scribbles, and a drawing of the Mona Lisa. Image B is a shifted version of Image A; in addition, some parts of the whiteboard have been erased, others have been added to.

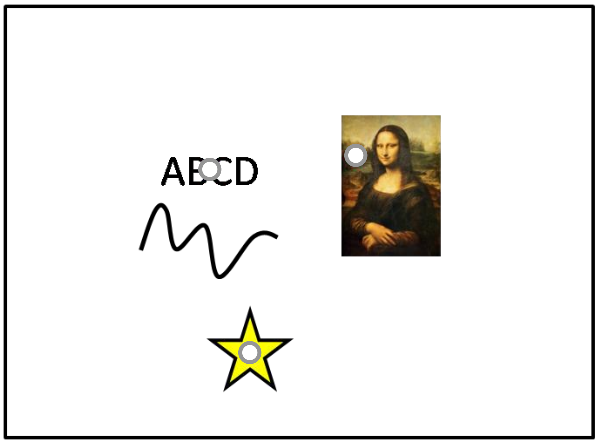

Image A

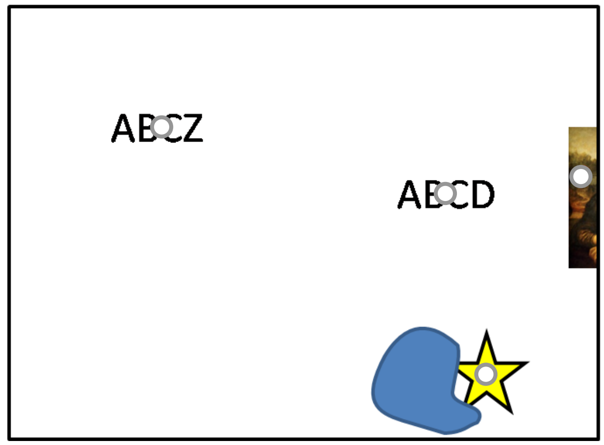

Image B

Suppose we detected a feature at each of the three white points in Image A, and each of the four white points in Image B. For each of the points in Image A, draw a square window of a "good" size, which will allow each feature in Image A to be unambiguously matched to the correct match in Image B, and will result in a low SSD value (assume the square windows are not rotated). Keep in mind possible problems with image boundaries, ambiguous matches, and changes between the two images. You will probably want different window sizes for different features.

Name____________________________

Problem 8: Segmentation (10 pts)

Part A: In this question you will manually run the steps of the K-Means algorithm for a very simple one-dimensional dataset. The data points are

2 5 12 20 23 26

Run the algorithm using two clusters, where the initial guesses for the cluster centers are **3.5** and **5.0**. For each step of the K-Means algorithm, **circle** the entries in each cluster and specify the new mean below. Assume that a step consists of assigning points to each cluster and computing the new mean. If it converges in fewer than five steps, you may leave the remaining steps blank.

After Step 1:
cluster1: 2 5 12 20 23 26 mean ___ cluster2: 2 5 12 20 23 26 mean ___

After Step 2:
cluster1: 2 5 12 20 23 26 mean ___ cluster2: 2 5 12 20 23 26 mean ___

After Step 3:
cluster1: 2 5 12 20 23 26 mean ___ cluster2: 2 5 12 20 23 26 mean ___

After Step 4:
cluster1: 2 5 12 20 23 26 mean ___ cluster2: 2 5 12 20 23 26 mean ___

After Step 5:
cluster1: 2 5 12 20 23 26 mean ___ cluster2: 2 5 12 20 23 26 mean ___

Part B: For the above dataset, give an example of an initial guess of two cluster centers for which the K-Means algorithm does not converge to the same segmentation as in Part A. Specify (1) the initial two means, and (2) the final two means. The cluster centers must be initialized to two different values, and the initial values cannot be (3.5, 5.0) or (5.0, 3.5).