Stanford Honor Code Statement:

The Honor Code is an undertaking of the students, individually and collectively:

1. that they will not give or receive aid in examinations; that they will not give or receive impermissible aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;

2. that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code;

The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid, as far as practicable, academic procedures that create temptations to violate the Honor Code.

While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to establish optimal conditions for honorable academic work.

I acknowledge and accept the Honor Code.

Signature: _____

Name (Print): _____

Date: _____

CS148 Midterm, Summer 2010

You have 2 hours to complete this exam
The exam is closed-everything, including calculators

The question paper has 10 pages & 6 questions (including extra credit). Avoid attaching extra sheets as far as possible. You may use blank scratch sheets, but don't submit them. You don't need to show steps in a solution unless they're explicitly asked for.

NAME: _____

1. Light and Color *(8 points)*

1a) Andrew is playing with cyan, magenta and yellow paint. Assuming the paints mix perfectly (no chemical reactions), what colors would he observe when he mixes the following? *(4 points)*

Cyan + Magenta = _____

Cyan + Yellow = _____

Yellow + Magenta = _____

Cyan + Magenta + Yellow = _____

1b) Bob decides to stand up to the CIE and choose his own set of primaries in an additive color space (mixing beams of light). Bob's primaries are (R + B) / 2, (B + G) / 2 and (G + R) / 2. In other words, he constructs each of his primaries by taking two CIE primary beams at half the reference intensities, and mixing them together. Is the gamut of his color space smaller than, larger than, or equal to CIE RGB? Briefly justify your answer. *(4 points)*

2. Cameras, Displays and Compression *(17 points)*

2a) Alice is exploring Central Asia, and is trying to take a picture of wild horses running across the plains, with high peaks in the distance. She places her camera on a tripod, focuses on the horses (which are fairly close to her), and clicks the shutter. Upon reviewing the shot, she finds the exposure is correct (so the picture is neither too bright nor too dark), but both the horses and the peaks are blurry. How should she adjust her camera variables to take a picture which is equally well-exposed, but in which everything is sharp? (Tick one option for each variable.) *(3 points)*

Shutter Speed:	Shorter exposure time _____	Longer exposure time _____
Aperture:	Increase _____	Decrease _____
Sensitivity (ISO):	Increase _____	Decrease _____

2b) Recall that a Bayer filter is a mask of red, green and blue filtering pixels placed over a digital sensor. State one advantage and one disadvantage of this filter. *(2 + 2 points)*

Advantage:

Disadvantage:

2c) Camera lenses typically contain multiple individual elements (separate pieces of glass), even when the lens has a single focal length and does not zoom. Why do you think this is so? *(2 points)*

2d) Daila tosses a fair coin a million times to produce a sequence of heads (0) and tails (1). Would you expect this sequence to have high or low Kolmogorov complexity? Briefly justify your answer (an informal argument is fine). *(2 points)*

2e) Eve decides to construct her own "wavelet" encoding for sequences of positive numbers. For every successive pair of numbers <A, B>, she computes the "average coefficient" as sqrt(A * B), and the "detail coefficient" as sqrt(A / B), where sqrt denotes the (positive) square root. In other words, she replaces the arithmetic mean (A + B) / 2 of Haar wavelets with the geometric mean.

- Compute the "wavelet" transform of [4 4 1 1] using this encoding (Hint: this is another sequence of the same size.) *(4 points)*

- Would this be a good way to compress data on a typical PC? If so, why, and if not, why not? *(2 points)*

3. Image Processing *(7 points)*

3a) Is the dynamic range of an image that has 16 bits per channel necessarily greater than that of an image with 8 bits per channel? Argue for or against. *(2 points)*

3b) Here is a signal *f* that needs to be smoothed (blurred). On the empty graph on the right, sketch a one-dimensional kernel *g* that you think will best smooth *f* when convolved with it. The overall scale of your drawing is not important. *(2 points)*

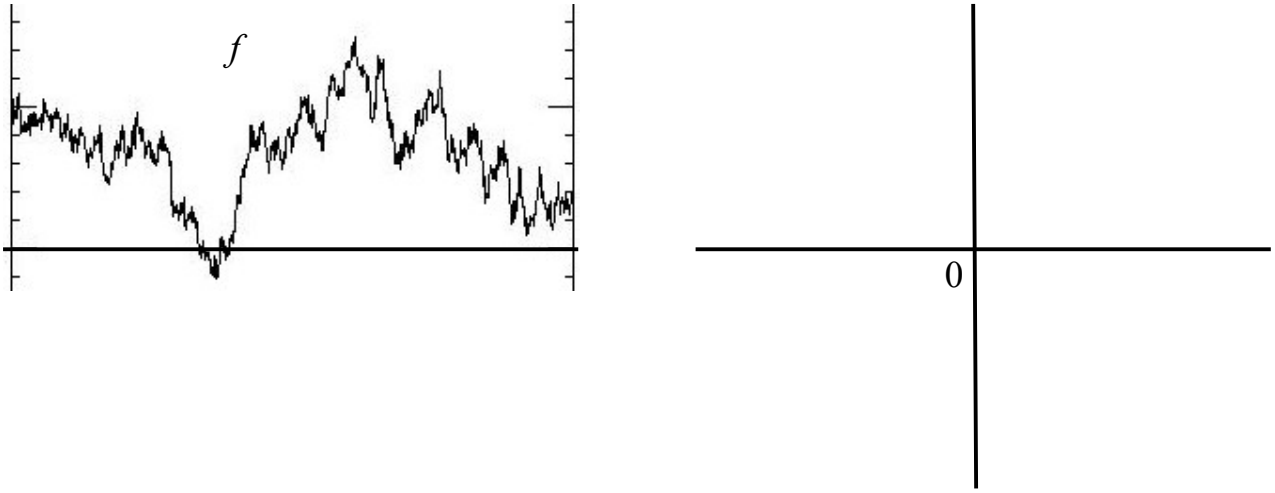

3c) Carol shrinks an image to half its size (one-quarter its area) by deleting every second row and every second column of pixels. Is this the best way to resize the image? If not, how would you do it? *(3 points)*

4. Geometry *(6 points)*

We'll use the following notation for 2D transformation matrices:

T(x, y): Translates a point by (x, y)
R(a): Rotates counter-clockwise around the origin by a degrees
S(sx, sy): Scales the coordinates of a point, relative to the origin, by (sx, sy)

You are given a unit square with minimum corner (0, 0):

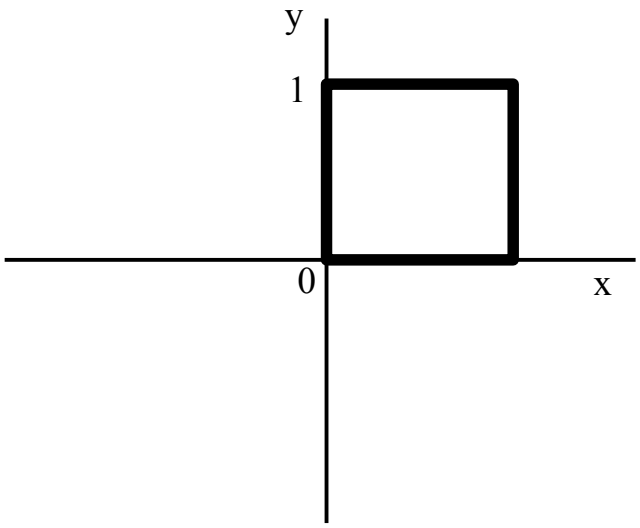

In the graph below, sketch what happens to the square when we apply
T(0, -0.5) * R(45) * S(0.5, 1) to it. Remember: we apply transforms right-to-left. Overall scale of your drawing is not important.

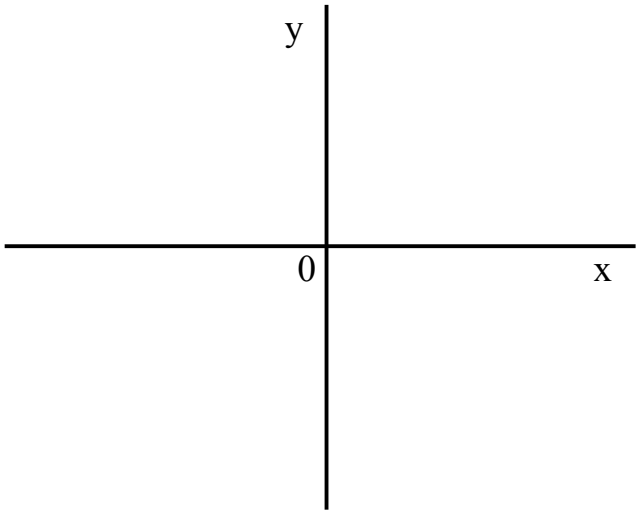

5. Rendering *(12 points)*

Anand wants to render an outdoor scene with a fog effect. He observes that fog has the following characteristics:

- it has a color of its own, usually white but sometimes other colors
- shapes are faded to the fog color with distance

5a) Help Anand write a **fog**(*d*, *src_color*, *fog_color*, *fog_density*) function that takes 4 parameters:

- *d*: the distance to an object
- *src_color*: the original (shaded) color of the object
- *fog_color*: the color of the fog
- *fog_density*: how quickly shapes fade to the fog color – greater density implies a quicker fade

fog(...) returns the perceived color of the object seen through the fog. Write either a math formula or pseudocode for a plausible **fog**(...) function. (There is no single correct answer. Your solution should satisfy Anand's two criteria at least.) *(4 points)*

5b) Extend the basic raytracing loop, in pseudocode, to handle a uniform fog throughout the scene. You may assume you are given the **fog**(...) function as a black box. Ignore refractions. *(8 points)*

Note: Yes, we do want you to write the whole loop, ignoring refractions. Syntax is not important, and you can assume you have all the math functions available to you.

6. Extra Credit *(+10 points)*

Recall that a 2D shear matrix (for non-homogenous coordinates) is one that has the form $\begin{bmatrix} 1 & s \\ 0 & 1 \end{bmatrix}$ (horizontal shear) or $\begin{bmatrix} 1 & 0 \\ s & 1 \end{bmatrix}$ (vertical shear). There is a connection between shearing and rotation, as can be seen in this sequence.

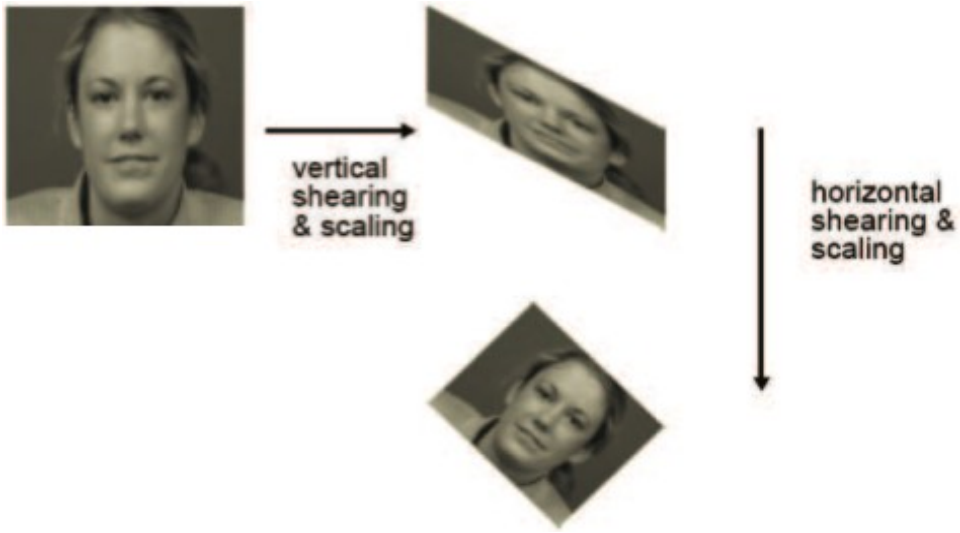

This has been exploited in image processing, as shearing maps are easily performed by shifting scan-lines (columns or rows). Explore this connection by writing the 2D rotation matrix $\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$ as a product of a horizontal shear, then a vertical shear, and finally another horizontal shear. Show the steps in your reasoning.

(This page intentionally left blank for the answer to the extra credit question.)