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 Final, Summer 2010

You have 2 hours and 15 minutes to complete this exam
The exam is closed-everything, including calculators

The question paper has 12 pages, 5 questions and 50 points (including the page for the Honor Code statement and the extra credit question). 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, Color, Cameras and Images *(12.5 points)*

1a) Recall that CMYK separation is not unique, i.e. the same image can be broken down into cyan, magenta, yellow and black channels in more than one way.

i) How many ways are there to do this decomposition? _____ *(1 point)*

ii) For reproduction on a typical inkjet printer, would you prefer the separation that has the maximum black component, or minimum black component? Explain your answer. *(1 point)*

1b) The ProPhoto RGB color space has a larger gamut than sRGB. Graham owns a camera capable of capturing images covering the entire ProPhoto RGB gamut, but his neighbor John's camera only captures the sRGB range. The cameras are otherwise identical. In the local photo exhibition, images from both cameras will be displayed on the *same* sRGB display. Does Graham's camera have any advantage over John's for this application? Justify your answer. (You may assume the pixel values are represented with arbitrarily precise real numbers in each case, so precision is not a differentiating factor.) *(2 points)*

1c) Terry spends a week processing a JPEG image. Each day, he loads the image into Photoshop, makes some minor edits, and then resaves the image as JPEG. At the end of each day, he quits Photoshop and shuts down his computer. Why is this a bad workflow? How could Terry do better? Explain clearly why your approach would improve matters. *(2 points)*

1d) Draw the convolution of the following triangular function with itself. Your sketch should be on the same graph, and roughly to scale. *(2 points for the correct basic shape + 2 points if the plot crosses the axes at approximately the correct coordinates)*

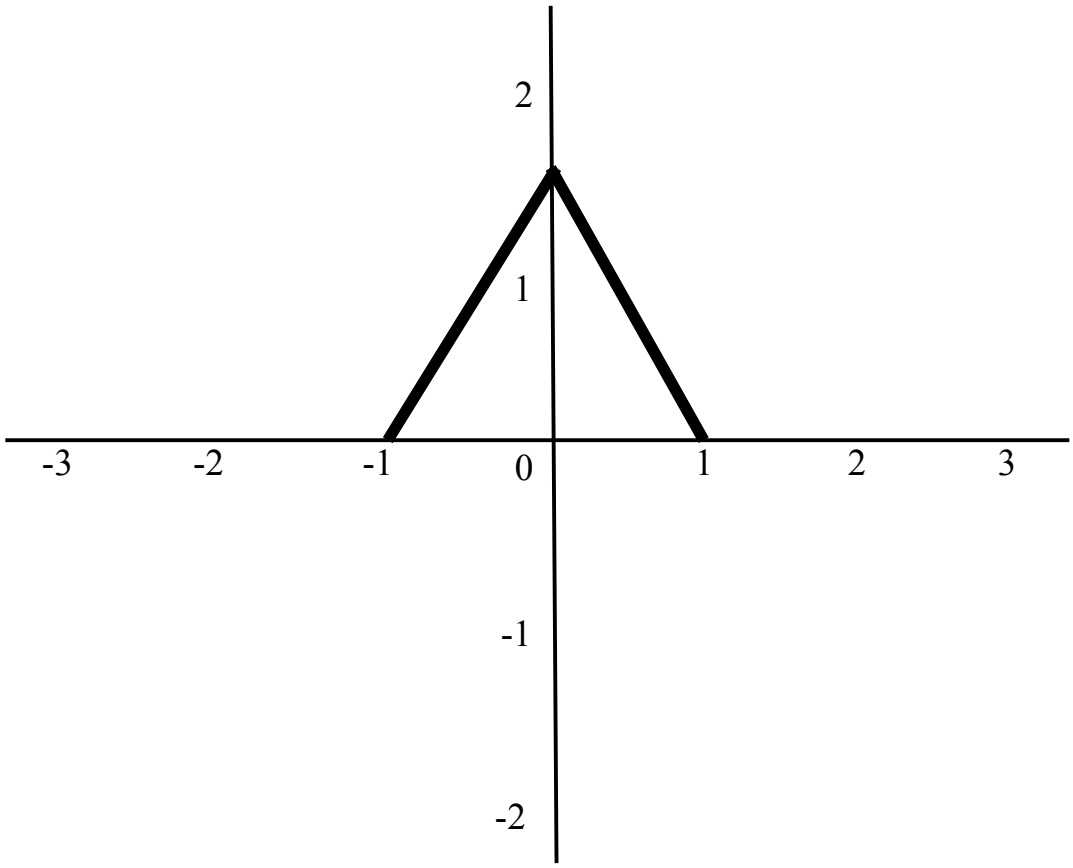

1e) What is the function of the alpha channel in an image? What do alpha = 0 and alpha = 1 correspond to? *(1 point)*

1f) Recall that in blue screen matting, the actor stands in front of a blue screen and the part of the image corresponding to this screen is automatically removed, to be replaced with a different background. Assume all objects in the scene are opaque, no part of the actor or his/her clothing is blue, and the screen appears an absolutely uniform blue color in the image (no variations in shading). Why is it not sufficient to simply delete all pixels that are approximately the same color as the screen? *(1.5 points)*

2. Geometry, Curves and Surfaces *(12.5 points)*

2a) Vectors interpreted as positions and directions are different. Which of these operations are legal, and which are illegal? (Tick/check one in each case.) Consider both arguments and return type. *(8 * 0.5 = 4 points)*

Position = Scalar * Position	Legal _____	Illegal _____
Direction = Position – Position	Legal _____	Illegal _____
Position = Position + Direction	Legal _____	Illegal _____
Direction = Direction + Direction	Legal _____	Illegal _____
Position = Position – Position	Legal _____	Illegal _____
Position = Position + Position	Legal _____	Illegal _____
Direction = Position + Direction	Legal _____	Illegal _____
Direction = Scalar * Direction	Legal _____	Illegal _____

2b) Find a formula for the normal to the implicit 2D surface

$$ax^3 + bxyz + cz^2 + 1 = 0$$

embedded in 3D, at a point (x, y, z). Briefly outline how you arrived at your answer. (Your answer should be in terms of the variables x, y and z. We'll accept either of the two valid, opposing directions for the normal. The normal need not have unit length.) *(2 points)*

2c) A curved surface may be approximated in piecewise linear fashion by a polygon mesh. This process is called tessellation. A mesh consists of vertices, edges and faces.

i) Name a rendering trick to make a mesh appear to be a smoothly varying surface. (A description of such a technique is not a valid answer.) *(1 point)*

ii) Eric writes a procedure to tessellate an arbitrary curved surface. The vertices of the resulting mesh are guaranteed to sample the curved surface uniformly by area. When is this a good thing, and when is this a bad thing? Explain your answer. *(1 + 1 points)*

2d) For two curve segments to be joined together, define: *(1 + 1 points)*

C^0-continuity: _____

C^1-continuity: _____

2e) What is the principal difference between cubic Hermite splines and cubic Bézier splines (other than the mathematical descriptions, of course)? *(1.5 points)*

3. Rendering and GPGPU *(12.5 points)*

3a) GPUs are special-purpose processors designed to display lots of geometry very fast. Briefly answer the following questions about GPUs.

i) Why are scalar (serial processing) chips not getting significantly faster? *(1 point)*

ii) Give a characteristic of code that can run very fast on the GPU. *(1 point)*

iii) Which of the following would you expect to find in a GPU shader core circa 2010? *(10 * 0.5 = 5 points)*

Fetch-Decode Unit	Yes _____	No _____
L3 Cache	Yes _____	No _____
L2 Cache	Yes _____	No _____
Several Arithmetic Logic Units (ALUs)	Yes _____	No _____
Only a single ALU	Yes _____	No _____
Registers	Yes _____	No _____
Memory Reading/Writing Units	Yes _____	No _____

Memory Pre-Fetcher	Yes _____	No _____
Branch Predictor	Yes _____	No _____
Double Precision Floating Point Unit	Yes _____	No _____

iv) Name an API or language designed specifically for writing GPGPU programs. *(1 point)*

3b) Place these objects in the order in which they are processed in the graphics pipeline (the one that is processed earliest should be placed first): **pixels**, **primitives**, **vertices**, **fragments**. *(2 points)*

3c) Terry writes a z-buffer-based graphics application using OpenGL and executes it on a machine with a modern GPU. Terry's program simulates transparency by setting an alpha value for each fragment. If the fragment passes the depth test, the framebuffer pixel is updated by blending the fragment color with the existing color according to this alpha value. What could go wrong with this scheme? How would you fix it?
(1 point for error in Terry's scheme + 1 point for partial fix + 0.5 for complete fix)

4. Animation, Interaction, Visualization and Typography *(12.5 points)*

4a) Recall that Inverse Kinematics (IK) solves the following problem: Given a few constraints, make a body with many degrees of freedom move according to those constraints. This is usually a very underconstrained problem: the body has far more degrees of freedom than there are constraints. Describe any common way in which an IK system might further constraint the motion to derive a unique solution. *(1 point)*

4b) Distinguish between polling and interrupts. *(2 points)*

4c) Michael is studying the distribution of authors of books in his collection. For each author, he has two fields: birthplace (geographical location) and genre (a string). Sketch an effective visualization that simultaneously presents the distribution of both these attributes to Michael. Label your sketch to indicate particular features of your visualization. You may indicate, with text or crosshatching, the use of color if necessary. (This is an open-ended question – there is no one correct answer though some are obviously better than others. You should try to ensure Michael can glean as much information as possible from a quick glance.) *(3 points)*

4d) Here is a word typeset in a serif font.

i) Mark and label the cap height, x-height and baseline on the image. *(1.5 points)*
ii) Identify any two serifs (circle them and label them: “serif”). *(1 point)*
iii) Identify a ligature (circle it and label it: “ligature”). *(1 point)*

Please be unambiguous in your labeling, else you won't get credit.

4e) Define the following two terms, in the context of typography. *(1 + 1 points)*

i) Leading:

ii) Kerning:

4f) State one way in which Unicode is an improvement over ASCII. *(1 point)*

5. Extra Credit *(0 points)*

Answer one or more of the following:

i) Suggest a cool new assignment for CS148.
ii) What's a good place to eat out within 50 miles of where you are now, which won't bust our meagre grad student incomes? We're pretty familiar with University Ave, so out-of-the-way places are preferred.
iii) Tell us a joke we haven't heard before.

(To clarify: the answer, or lack of one, absolutely does not affect your grade.)