PLEASE HAND IN

UNIVERSITY OF TORONTO
Faculty of Arts and Science

DECEMBER EXAMINATIONS 2007

CSC 487/2503H F
St. George Campus

Duration — 3 hours

Aids allowed: none

PLEASE HAND IN

Student Number: |_|_|_|_|_|_|_|_|_|

Family Name: ____________________

Given Name: ____________________

*Do **not** turn this page until you have received the signal to start.*
(In the meantime, please fill out the identification section above, and read the instructions below.)

This examination consists of 6 questions on 15 pages (including this one). *When you receive the signal to start, please make sure that your copy of the examination is complete.* If you need more space for one of your solutions, use the reverse side of the page and *indicate clearly the part of your work that should be marked.* Also, the last two pages are provided for scratch work.

Write your student number at the bottom of pages 2-15 of this test.

1: ______/15

2: ______/15

3: ______/15

4: ______/15

5: ______/15

6: ______/15

TOTAL: ______/90

Good Luck!

Question 1. [15 MARKS]

Give a short answer for each of the following questions.

Part (a) [3 MARKS] Describe the Canny edge detection algorithm, including how the edge strength and orientation is calculated.

Part (b) [3 MARKS] What is the spatial orientation tensor and how is it used?

Part (c) [3 MARKS] When computing image segmentations, describe one difficulty you might expect to have when using simple connected components. How does NCut deal with this issue?

Part (d) [3 MARKS] In an orthographic multiframe factorization algorithm, describe what is meant by an affine reconstruction. How is this upgraded to a metric reconstruction?

Part (e) [3 MARKS] What special difficulties would you expect to have in computing optical flow given spatially aliased images? What special difficulties would you expect to have in computing optical flow given a temporally aliased image sequence? What can you do to deal with these situations?

Question 2. [15 MARKS]

Suppose we place a green pepper in a well light room which also contains a colour TV (which is on and showing Seinfeld reruns). We take a colour image of the green pepper. In this image we notice the reflection of the TV is visible on a small patch of the surface of the green pepper. Your task is to remove this reflection from the image of the pepper (much as you did in assignment 1).

Part (a) [3 MARKS] Describe a mathematical model for the distribution of RGB responses for the image of this green pepper. Your model should include the smoothly shaded portion of the pepper as well as the reflection of the TV. Make sure you clearly define all the components of this model.

Part (b) [2 MARKS] In assignment 1 you derived equations which provided an estimate the colour of the green pepper with the highlight removed, and you could apply these equations at each pixel separately. Can you do that here? Explain why or why not.

Question continued on next page.

Part (c) [5 MARKS] Suppose you take several images, with the camera on a tripod and with the only change being the image displayed on the TV, how might you approach the problem of removing the highlight? Be as specific as possible.

Part (d) [5 MARKS] How could you approach this problem given just a single image? Be as specific as possible. [Hint: Make use of the fact that the shading on the pepper is slowly varying and the reflection of the TV covers only a small patch.]

Question 3. [15 MARKS]

Let's say we wish to fit a 1D curve $c(t) = a_0 + a_1 t + a_2 t^2$ to a set of N observed 1D image locations, $\{x_j\}_{j=1}^N$, in frames t_j, for $j = 1...N$. That is, you are given $\{(x_j, t_j)\}_{j=1}^N$. Further assume that we measure error between the curve and the measurement as $\rho(c(t_j) - x_j)$ for a robust estimator ρ:

Part (a) [2 MARKS] Define the *influence function* and the *weight function* in terms of a general robust estimator $\rho(e)$.

Part (b) [2 MARKS] Using this function, $\rho(e)$, write an objective function that you might try to minimize in order to find $c(t)$.

Part (c) [2 MARKS] Specify the mathematical form of a robust estimator function $\rho(e)$ that you might use if you are told that the data set has outliers.

Question continued on next page.

Part (d) [5 MARKS] For the specific robust estimator in Part (c) above, derive an iteratively reweighted least squares algorithm for minimizing the objective function.

Part (e) [4 MARKS] Describe how you would automatically generate initial guesses for your iterative method in part (d) above, given that about half the data is expected to be outliers.

Question 4. [15 MARKS]

Consider estimating the translational component of optical flow given two images $I(\vec{x}, t)$ and $I(\vec{x}, t+1)$. You can assume that these images have been cropped so that we only need to estimate a constant displacement between these patches.

Part (a) [3 MARKS]

What is the brightness constancy constraint?

Part (b) [3 MARKS]

Given an initial guess $\vec{u}_0$ for the displacement, so $I(\vec{x}+\vec{u}_0, t+1) \approx I(\vec{x}, t)$, what are the linearized brightness constancy constraints that can be used to update the displacement? Derive a linear system of equations for the update, say $\vec{v}$.

Question continued on next page.

Part (c) [3 MARKS]

How can a robust estimator $\rho(e)$ be applied to the linearized brightness constancy constraints? In particular, write out the objective function to be minimized by the displacement update $\vec{v}$.

Part (d) [3 MARKS]

Derive an iteratively reweighted least squares algorithm for minimizing your objective function in part (c) above.

Part (e) [3 MARKS]

Given the converged estimate of the update v from part (d) above corresponds to a translation of several pixels, what should you do next? Explain.

Question 5. [15 MARKS]

During computer assisted surgery a surgeon wishes to have a computer track the low dose x-ray image of a 3D linear probe as the surgeon moves the probe within the patient's body. Assume that the probe appears to be a line segment of roughly uniform brightness in the x-ray image. Moreover assume the brightness of the probe is just slightly brighter than the average brightness of the background tissue in the x-ray image, and that the brightness of background tissue varies considerably across the image. The probe is not the brightest object in the image (nor is it the darkest object). The length of the line segment corresponding to the 3D probe is the length of the perspective image of the probe, so it varies in time.

We wish to try a particle filter to perform this tracking.

Part (a) [2 MARKS]

What state space would you use, and why?

Part (b) [2 MARKS]

What would you use to model the dynamics. Explain.

Part (c) [3 MARKS]

What would you use for the likelihood function? How would you train the likelihood function?

Question continued on next page.

Part (d) [4 MARKS]

How would you maintain an estimate for the state parameters, along with their uncertainties? Explain in detail.

Part (e) [2 MARKS]

What sort of diagnostics would you use to identify when the tracker was working well (and not working well)?

Part (f) [2 MARKS]

Describe situations in which you would expect this to work well, and also other situations in which you would expect this tracker to have trouble.

Question 6. [15 MARKS]

Suppose we have a stereo pair of cameras viewing the interior of a cluttered room and a laser pointer which emits light spanning a 3D plane (instead of a single ray, like most pointers). The intersection of this plane of light with the various surfaces in the room is visible as a "light stripe" in each camera.

Part (a) [2 MARKS] Describe all situations in which this light stripe appears to be a single straight line in one of the images of the stereo pair.

Part (b) [3 MARKS] Describe all situations in which this light stripe appears to be a single straight line in both images of the stereo pair (but not necessarily the same line in each image).

Part (c) [4 MARKS] Given that the plane of light intersects multiple smooth surfaces in the room, and appears as a single linear light stripe in both stereo images, formulate linear constraints on the F-matrix for this stereo pair in terms of the position and orientation of these two light stripe lines.

question continued ... on next page.

Part (d) [2 MARKS] In the situation described in part (c) above, what do you know about the position of the left and right epipoles? Explain.

Part (e) [4 MARKS] Explain how you could use this equipment to determine the F-matrix for the stereo pair of cameras from two pairs of light stripe images. [Hint: Consider rotating the laser pointer after you take the first pair of stereo images.] Be as specific about the algorithm you would use to solve for F (i.e., clearly formulate the constraints on F and describe how you would solve them).

For scratch work.

For scratch work.

Total Marks = 90