CompSci 527 Final Exam Sample

The questions below are more involved than the ones in the exam. First, they are "cascaded," in that many answers depend on previous answers, and this is not good in an exam. Second, there are more questions below than there will be on the exam. On the other hand, this exam sample is designed to prepare you for the final. If you do well on this sample, you are likely to do well on the final.

The exam will be closed-book, closed-notes, and you will not be allowed to have anything other than the exam and a pen/pencil and an eraser on your desk. The amount of space provided under each question is *not* an indication of the length of the answer. Materials covered are sections 4-6 of the class syllabus web page, the last two homework assignments, and this sample. Parenthesized materials on the syllabus page and Appendices in the class notes are *not* required reading, except to the extent that they help you understand materials elsewhere.

1. This problem takes you through the computation of the set of all least-squares solutions to the following linear system:

$$\begin{aligned} 3x + 4y &= 2 \\ 3x + 4y &= 3 \end{aligned}$$

and the solutions to a related optimization problem. All the answers to the questions in this problem are numerical, and for the data given in the problem, and no more general answers are required. You may leave your answers in the form of fractions, with expressions like the following:

$$\frac{\sqrt{3}}{2}\begin{bmatrix} 2 \\ -5 \end{bmatrix},$$

but please simplify. If you cannot answer a question, write out a symbolic placeholder, and go on to the next question. For instance, if you don't know what $\mathbf{b}$ is in the first question, write

$$\mathbf{b} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}.$$

(a) What are A and $\mathbf{b}$ if we write the system in this problem in the following form?

$$A\mathbf{x} = \mathbf{b}$$

(b) What is the rank of A?

(c) Give a unit column vector $\mathbf{r}$ that spans the row space of A.

(d) Give a unit column vector $\mathbf{n}$ that spans the null space of A.

(e) Write the matrix V in the SVD $A = U\Sigma V^T$ of A.

(f) Compute the matrices U and Σ in the SVD of A. [Hint: compute $U\Sigma$ first.]

(g) Compute the pseudo-inverse $A^{\dagger}$ of A.

(h) Find the minimum-norm solution $\mathbf{x}^*$ of the system $A\mathbf{x} = \mathbf{b}$.

(i) Give an expression for the set S of all least-squares solutions of the system $A\mathbf{x} = \mathbf{b}$.

(j) Find all the solutions to

$$\hat{\mathbf{x}} = \arg \min_{\|\mathbf{x}\|=1} \|A\mathbf{x}\| \ .$$

2. A one-dimensional, real image

$$I(x) \ : \ D \to \mathbb{R} \quad \text{where} \quad D \subset \mathbb{Z}$$

can be thought of as a single row out of a regular, two-dimensional, gray image. The integer domain D is the interval of pixel positions on which the image is defined. We want to track points between two such images

$$I(x) \ : \ D \to \mathbb{R} \quad \text{and} \quad J(x) \ : \ D \to \mathbb{R}$$

by finding the displacement d that minimizes the Sum-of-Squared-Differences (SSD) residual

$$\epsilon(x_I, d) = \sum_x [J(x+d) - I(x)]^2 \ w(x - x_I) \quad \text{where} \quad w(x) = \begin{cases} 1 & \text{if } |x| \leq 3 \\ 0 & \text{otherwise} \end{cases}$$

is a window over 7 pixels, and x_I is the point in I we want to track. This problem is the one-dimensional formulation of point-feature tracking, and first and second derivative are analogous to gradient and Hessian for two-dimensional images.

(a) Write an expression for the Taylor series of $J(x + d)$ in d and around $J(x)$, truncated to just after the linear term. Use symbol $J'(x)$ to denote the derivative of $J(x)$ at x.

(b) Use your answer to the previous question to write an expression that approximates $\epsilon(x_I, d)$ with a quadratic function of d.

(c) Write an expression for the approximate derivative $\epsilon'(x_I, d)$ with respect to d, using the quadratic function from the previous question as an approximation for ϵ.

(d) Write an expression for the approximate second derivative $\epsilon''(x_I, d)$ with respect to d, using the same approximation for ϵ as in the last question.

(e) The Taylor series of $\epsilon(x_I, d)$ around 0 and truncated to the second term is

$$\epsilon(x_I, d) \approx \epsilon(x_I, 0) + \epsilon'(x_I, 0)d + \frac{1}{2}\epsilon''(x_I, 0)d^2 \ .$$

Find a formula for the minimum d^* of this truncated series with respect to d, assuming that $\epsilon''(x_I, 0) > 0$, and using the approximations for the derivatives of ϵ that you found when answering the last two questions.

(f) Consider the trivial case in which I and J are shifted versions of the same linear function

$$I(x) = ax + b \quad \text{and} \quad J(x) = ax + c\ .$$

What do you expect the exact solution of the following optimization problem to be in this case, and for any x_I?

$$d_x = \arg\min_d \epsilon(x_I, d)$$

(g) Verify that in the special case described in question 2f the value you gave for d_x is the same as the value of d^* you found when answering question 2e.

(h) How many iterations would the one-dimensional version of the Lucas-Kanade tracker take to find the exact solution for the special case in question 2f, starting with initial solution $d_0 = 0$? Explain briefly and clearly why this is the case.

(i) The case $a = 0$ is problematic, because it would lead to division by zero. What is this problem called in computer vision?

(j) A good feature to track for the two-dimensional Lucas-Kanade tracker is one that satisfies the following constraints:

$$\kappa_2(A_I(\mathbf{x}_I)) \leq \kappa_{\max} \quad \text{and} \quad \sigma_{\min}(A_I(\mathbf{x}_I)) \geq \sigma_0$$

where $\kappa_{\max}$ and σ_0 are suitable thresholds and

$$A_I(\mathbf{x}_I) = \sum_{\mathbf{x}} \nabla I(\mathbf{x})[\nabla I(\mathbf{x})]^T \, w(\mathbf{x} - \mathbf{x}_I) \; .$$

Which of the two constraints above is relevant for the one-dimensional case examined in this problem, and what is the value of the left-hand side of the relevant constraint in the special case examined in question 2f?

(k) What does the relevant constraint you identified in the previous question mean in terms of the images, for the special case in question 2f? Why is this constraint needed for good tracking?

3. This problem examines models for two distortion-free cameras, their essential matrix, and some aspects of the eight-point algorithm. Give all your answers in numerical form and for the specific cameras that are described in the first two questions of the problem, rather than in general terms.

   (a) A certain camera C has no lens distortion, a focal length of 10 mm, a sensor with 200 pixels per millimeter both vertically and horizontally, and principal point at the pixel at column 600, row 400 in the image. Write the projection matrix P of the camera in homogeneous coordinates.

   (b) A camera C' with the same projection matrix as camera C has its center of projection 100 millimeters to the right of camera C, along the X_1 axis of C. The optical axes of the two cameras are parallel to each other, their sensor rows are parallel to each other, and so are their sensor columns. Write the matrix G that transforms homogeneous coordinates $\mathbf{X}$ in the reference system of camera C to homogeneous coordinates $\mathbf{X}'$ in the reference system of camera C'.

(c) Write the projection matrix P' of camera C' so that if $\mathbf{X}$ is a vector of homogeneous coordinates of a point in the reference system of camera C then

$$\boldsymbol{\eta}' \sim P'\mathbf{X}$$

is a vector of homogeneous coordinates of the image point in camera C' (and in the reference system of C').

(d) A point $\mathcal{X}$ in the world has *Euclidean* image coordinates

$$\tilde{\boldsymbol{\xi}} = \begin{bmatrix} 1200 \\ 800 \end{bmatrix} \quad \text{and} \quad \tilde{\boldsymbol{\eta}}' = \begin{bmatrix} 1000 \\ 800 \end{bmatrix}$$

(in pixels) in the images taken by the cameras C and C' described earlier. Find the *Euclidean* image coordinates $\tilde{\mathbf{x}} = e(\mathbf{x})$ and $\tilde{\mathbf{y}}' = e(\mathbf{y}')$ of $\mathcal{X}$ in the canonical reference system.

[Hints: Find homogeneous coordinates first. Also,

$$\begin{bmatrix} a & 0 & b \\ 0 & a & c \\ 0 & 0 & 1 \end{bmatrix}^{-1} = \begin{bmatrix} \frac{1}{a} & 0 & -\frac{b}{a} \\ 0 & \frac{1}{a} & -\frac{c}{a} \\ 0 & 0 & 1 \end{bmatrix} \text{.]}$$

(e) The projection equations for camera C in its own canonical reference system are

$$\tilde{x}_1 = \frac{\tilde{X}_1}{\tilde{X}_3} \quad \text{and} \quad \tilde{x}_2 = \frac{\tilde{X}_2}{\tilde{X}_3}$$

in Euclidean coordinates. What are the projection equations for camera C' in the canonical reference system of camera C, for the particular geometry of the two cameras described in this problem? That is, what is the relationship between $\tilde{\mathbf{y}}'$ and $\tilde{\mathbf{X}}$? [Hint: There is no rotation between the cameras, so things are simple.]

(f) Solve the projection equations for the two cameras to find the vector $\tilde{\mathbf{X}}$ of the Euclidean coordinates of point $\mathcal{X}$ in the reference system of camera C. Specify the units of measurement for your solution.

(g) Write an essential matrix E for this pair of cameras.

(h) Write the equation in canonical Euclidean coordinates of camera C of the epipolar line of the image point that has the following canonical Euclidean coordinates in the second image:

$$\tilde{\mathbf{y}}' = \begin{bmatrix} 0.6 \\ 0.8 \end{bmatrix} .$$

(i) Let

$$\boldsymbol{\eta} = E(:)$$

be a vector of the entries of the essential matrix E you found above, listed in column-major order. For the two corresponding points with Euclidean coordinates $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}'$ that you found when answering question 3d, write out the vector $\mathbf{a}$ such that the epipolar constraint for these two points has the form

$$\mathbf{a}^T \boldsymbol{\eta} = 0$$

suitable for use in the eight-point algorithm.