

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

DEPARTMENT OF COMPUTER SCIENCE AND ELECTRICAL ENGINEERING

6.801/6.866 MACHINE VISION QUIZ I

Handed out: 2004 Oct. 21st

Due on: 2003 Oct. 28th

Problem 1: Uniform reflecting properties are a prerequisite for the usual shape from shading methods. Consider now a surface covered by a material of spatially varying reflectance. Suppose that the brightness can be treated as the product of a spatially varying ‘reflectance’ or ‘albedo’ $\rho(x, y)$, and a ‘geometric factor’ $R(p, q)$ that depends only on surface orientation.

There are two unknowns— $z(x, y)$ and $\rho(x, y)$ —at every position on the surface, so a single image will not provide enough information to recover both (consider, for example, a photographic print of a rounded object where the brightness variations could either be from a rounded object of uniform albedo or from a flat object of varying albedo).

Now suppose we take two images under different lighting conditions.

- (a) Combine the two resulting image irradiance equations in such a way as to eliminate $\rho(x, y)$. Suppose the new—now ‘ ρ -free’—equation can be written in the form

$$E'(x, y) = R'(p, q).$$

- (b) Show that if the underlying surface actually is a Lambertian reflector, then the isophotes in gradient space are straight lines. When will they be *parallel* straight lines?
- (c) Show that the isophotes all go through a common point in gradient space in the case that they are *not* parallel. Where in gradient space would you expect the highest accuracy in recovering surface orientation? That is, where is ‘brightness’ (in the ρ -free equation) most affected by small changes in surface orientation? Relate this back to the original imaging situation. How should the lighting be arranged to obtain high accuracy?

Problem 2: An edge detection method starts by finding the brightness gradient (E_x, E_y) at each picture cell using some local estimator of the partial derivatives of image brightness $E(x, y)$. The magnitude of the brightness gradient is then computed. Next, “non-maximum suppression” keeps for further consideration only pixels where the magnitude of the gradient is a local maximum in the direction of the local brightness gradient.

The maximum of the gradient need, of course, not fall exactly on a discrete pixel location, when the gradient magnitude is considered as a continuous function of position, Edge locations can be estimated to finer than pixel resolution if this continuous function of image position is first estimated from the discrete samples of the magnitude of the brightness gradient at pixels.

We can approximate the magnitude of the gradient locally as a power series

$$M(x, y) = ax^2 + bxy + cy^2 + dx + ey + f$$

where (x, y) is the displacement from the center of a 3×3 region of pixels.

- (a) Show that the parameters $a, b, c, d, e,$ and f can be estimated from $M, M_x, M_y, M_{xx}, M_{xy}$ and M_{yy} evaluated at the origin i.e. at the center pixel.
- (b) The following six stencils (a.k.a. computational molecules) can be used to estimate the parameters a, b, c, d, e and f . Indicate which stencil corresponds to which parameter.

| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|-----------|--|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----------|---|-----|--|-----|-----|--|-----|-----|--|-----|-----------|--|-----|-----|-----|-----|-----|-----|-----|-----|-----|
| (1) k_1 | <table style="border-collapse: collapse; margin: auto;"> <tr><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;">+ 1</td></tr> <tr><td style="padding: 2px 10px;">- 2</td><td style="padding: 2px 10px;">- 2</td><td style="padding: 2px 10px;">- 2</td></tr> <tr><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;">+ 1</td></tr> </table> | + 1 | + 1 | + 1 | - 2 | - 2 | - 2 | + 1 | + 1 | + 1 | (2) k_2 | <table style="border-collapse: collapse; margin: auto;"> <tr><td style="padding: 2px 10px;">- 1</td><td style="padding: 2px 10px;"></td><td style="padding: 2px 10px;">+ 1</td></tr> <tr><td style="padding: 2px 10px;"></td><td style="padding: 2px 10px;"></td><td style="padding: 2px 10px;"></td></tr> <tr><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;"></td><td style="padding: 2px 10px;">- 1</td></tr> </table> | - 1 | | + 1 | | | | + 1 | | - 1 | (3) k_3 | <table style="border-collapse: collapse; margin: auto;"> <tr><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;">+ 1</td></tr> <tr><td style="padding: 2px 10px;"></td><td style="padding: 2px 10px;"></td><td style="padding: 2px 10px;"></td></tr> <tr><td style="padding: 2px 10px;">- 1</td><td style="padding: 2px 10px;">- 1</td><td style="padding: 2px 10px;">- 1</td></tr> </table> | + 1 | + 1 | + 1 | | | | - 1 | - 1 | - 1 |
| + 1 | + 1 | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| - 2 | - 2 | - 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + 1 | + 1 | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| - 1 | | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + 1 | | - 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + 1 | + 1 | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| - 1 | - 1 | - 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| (4) k_4 | <table style="border-collapse: collapse; margin: auto;"> <tr><td style="padding: 2px 10px;">- 1</td><td style="padding: 2px 10px;">+ 2</td><td style="padding: 2px 10px;">- 1</td></tr> <tr><td style="padding: 2px 10px;">+ 2</td><td style="padding: 2px 10px;">+ 5</td><td style="padding: 2px 10px;">+ 2</td></tr> <tr><td style="padding: 2px 10px;">- 1</td><td style="padding: 2px 10px;">+ 2</td><td style="padding: 2px 10px;">- 1</td></tr> </table> | - 1 | + 2 | - 1 | + 2 | + 5 | + 2 | - 1 | + 2 | - 1 | (5) k_5 | <table style="border-collapse: collapse; margin: auto;"> <tr><td style="padding: 2px 10px;">- 1</td><td style="padding: 2px 10px;"></td><td style="padding: 2px 10px;">+ 1</td></tr> <tr><td style="padding: 2px 10px;">- 1</td><td style="padding: 2px 10px;"></td><td style="padding: 2px 10px;">+ 1</td></tr> <tr><td style="padding: 2px 10px;">- 1</td><td style="padding: 2px 10px;"></td><td style="padding: 2px 10px;">+ 1</td></tr> </table> | - 1 | | + 1 | - 1 | | + 1 | - 1 | | + 1 | (6) k_6 | <table style="border-collapse: collapse; margin: auto;"> <tr><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;">- 2</td><td style="padding: 2px 10px;">+ 1</td></tr> <tr><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;">- 2</td><td style="padding: 2px 10px;">+ 1</td></tr> <tr><td style="padding: 2px 10px;">+ 1</td><td style="padding: 2px 10px;">- 2</td><td style="padding: 2px 10px;">+ 1</td></tr> </table> | + 1 | - 2 | + 1 | + 1 | - 2 | + 1 | + 1 | - 2 | + 1 |
| - 1 | + 2 | - 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + 2 | + 5 | + 2 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| - 1 | + 2 | - 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| - 1 | | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| - 1 | | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| - 1 | | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + 1 | - 2 | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + 1 | - 2 | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + 1 | - 2 | + 1 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

- (c) Determine suitable values for the six constants $\{k_i\}$, assuming that the pixel spacing is ϵ . (Hint: apply the stencil to the test functions $x^i y^j$ for $i = 0, 1, \dots$ and $j = 0, 1, \dots$).
- (d) Show that the distance from the origin to the maximum in brightness gradient in the direction of the local brightness gradient is

$$\rho = -\frac{M_x E_x + M_y E_y}{M_{xx} E_x^2 + 2M_{xy} E_x E_y + M_{yy} E_y^2} \sqrt{E_x^2 + E_y^2}$$

- (e) This result only applies when the denominator of the above expression is negative. Why?
- (f) How large can ρ reasonably become? Keep in mind that a picture cell where this computation is performed is one that has survived the non-maximum suppression operation.

Problem 3a: In a simple model of brightness in a Scanning Electron Microscope we have a reflectance map

$$R(p, q) = p^2 + q^2.$$

At a singular point, $E_x = 0$ and $E_y = 0$. Assume the singular point is at $(0, 0)$. Suppose that the surface near the singular point can be approximated by a power series of the form

$$z = z_0 + s_x x + s_y y + ax^2 + 2bxy + cy^2$$

(a) First, show that $s_x = 0$ and $s_y = 0$ at the singular point.

(b) Then show that

$$E_{xx} - 2E_{xy} + E_{yy} = k((a - b)^2 + (b - c)^2)$$

for some k . What is the value of k ?

(c) If $E_{xx} = E_{xy} = E_{yy} = 8$, what are a , b , and c in the series expansion given above for the surface height near the singular point? Is there a unique answer? Verify that

$$b(E_{xx} - E_{yy}) = (a - c)E_{xy}$$

Problem 3b: In solving the shape from shading problem, we derived the characteristic strip equations from the image irradiance equation

$$E(x, y) = R(p, q).$$

Here we go in the other direction, showing that solutions of the characteristic strip equation satisfy the image irradiance equation.

Show that — when using the method of characteristic strip expansion to solve the shape from shading problem — $E(x, y) - R(p, q)$ is constant along a characteristic strip. That is:

$$\frac{d}{d\xi}(E(x, y) - R(p, q)) = 0$$

along the strip. Conclude that, if $E(x, y) = R(p, q)$ at the beginning of the strip, then $E(x, y) = R(p, q)$ all along the strip. (If nothing else, this is a sanity check on the characteristic strip equations).

Problem 4: The image brightness gradient is large where an edge is thought to occur. Derivatives of brightness can, however, be used for purposes other than finding edges.

We saw this when we computed the perimeter of a “binary” object using

$$\iint_D \sqrt{E_x^2 + E_y^2} dx dy$$

where $E(x, y) = 1$ in the object and $E(x, y) = 0$ in the background, with a narrow, smooth transition region in between.

For this exercise, consider an image with a circularly symmetric bright spot. Consider polar coordinates with $x = r \cos \theta$ and $y = r \sin \theta$. Let

$$E(x, y) = f(r) \quad \text{where} \quad r = \sqrt{x^2 + y^2}.$$

Suppose that $f(r) = 1$ for $r < R - \epsilon$, and $f(r) = 0$ for $r > R + \epsilon$. Then the above integral of the magnitude of the brightness gradient yields $2\pi R$ —i.e. the perimeter of the circular disk as $\epsilon \rightarrow 0$.

(a) Show that

$$E_x = f'(r) \frac{x}{r} \quad \text{and} \quad E_y = f'(r) \frac{y}{r}$$

and

$$\begin{aligned} E_{xx} &= f''(r) \frac{x^2}{r^2} + \frac{f'(r)}{r} \frac{y^2}{r^2}, \\ E_{xy} &= f''(r) \frac{xy}{r^2} - \frac{f'(r)}{r} \frac{xy}{r^2}, \\ E_{yy} &= f''(r) \frac{y^2}{r^2} + \frac{f'(r)}{r} \frac{x^2}{r^2}. \end{aligned}$$

(b) Show that

$$(E_x \quad E_y) \begin{pmatrix} E_x \\ E_y \end{pmatrix} = (f'(r))^2,$$

and that

$$(E_y \quad -E_x) \begin{pmatrix} E_{xx} & E_{xy} \\ E_{xy} & E_{yy} \end{pmatrix} \begin{pmatrix} E_y \\ -E_x \end{pmatrix} = (f'(r))^3 / r,$$

Conclude that

$$\iint_D \frac{E_{xx} E_y^2 - 2E_{xy} E_x E_y + E_{yy} E_x^2}{E_x^2 + E_y^2} dx dy = -2\pi$$

for the image $E(x, y)$ as given above, *independent of R* .

- (c) How is the result changed if we replace the bright disk with a dark disk in a bright background (i.e. replace $E(x, y)$ with $(1 - E(x, y))$)? What happens if the bright circular disk has an internal dark circular disk-shaped depression where E is zero? What happens if there is more than one bright circular disk in the image?
- (d) Based on the above, what measure or property do you think is computed by this integral when applied to a slightly blurred version of a binary image? Hint: It may help to consider an approximation of an arbitrary, simple, smooth, closed curve by a sequence of a large number of short circular arcs with tangent continuity.
- (e) Show that the integrand in the integral above is the directional second derivative of brightness in a particular direction. What direction?

Problem 5a: Consider binary image data streaming in along a row. We feed the bits into an adder/accumulator that is reset to zero at the start of the row. Clearly at the end of the row the accumulator contains the number of bits that were ‘1’ in that row. Now consider a second adder/accumulator that takes its input not from the image data, but from the output of the first accumulator. Like the first accumulator, it adds its current input to the current total when a new pixel is scanned in.

- Show that the contribution of a single bit in the image to the output of the second accumulator depends on its position in the image row—a bit at the end of the row being added in only once, while a bit at the start of the row is added in many times.
- How would you use such a device (modified if needed) to compute the first moment without using multiplication?
- Suggest how you could obtain the second (and higher) moments using additional accumulators, again without needing multiplication.

Problem 5b: Consider the binary image of a circular disk that is approximately d pixels across. We are going to estimate the expected error in finding the horizontal (x) component of the centroid. We break the image up into rows. Within each row that intersects the disk there is a contiguous set of pixels where $b_{ij} = 1$.



Now the value of first and last pixels in this sequence are somewhat uncertain in that additive image measurement noise might push the brightness measurements below threshold, or push adjacent pixels, now classified as “background,” above threshold. The centroid of the row of pixels then has an error with standard deviation some multiple of the size of a pixel. Due to this error, the centroids of different rows intersecting the disc will not be exactly the same. In computing the centroid of the disc, we are in essence averaging these noisy estimates.

How do you expect the error in the x -component of the centroid to depend on the size of the image of the disc?