

Contents

2	Light and Shading	1
3	Color	6
4	Linear Filters	11
5	Local Image Features	14
6	Texture	18
9	Segmentation by Clustering	20
10	Grouping and Model Fitting	23
11	Tracking	28
12	Registration	31
13	Smooth Surfaces and Their Outlines	33
15	Learning to Classify	35
16	Classifying Images	37

CHAPTER 2

Light and Shading

PROBLEMS

- 2.1. We see a diffuse sphere centered at the origin, with radius one and albedo ρ , in an orthographic camera, looking down the z -axis. This sphere is illuminated by a distant point light source whose source direction is $(0, 0, 1)$. There is no other illumination. Show that the shading field in the camera is

$$\rho\sqrt{1-x^2-y^2}$$

Solution:

The surface is $(x, y, \sqrt{1-x^2-y^2})$. We get two tangent vectors by partial differentiation; they are $(1, 0, p)$ and $(0, 1, q)$ where $p = -x/\sqrt{1-x^2-y^2}$ and $q = -y/\sqrt{1-x^2-y^2}$. You can verify that the unit normal is

$$\vec{N} = \frac{(-p, -q, 1)}{\sqrt{1+p^2+q^2}}$$

and the shading must be $\rho((0, 0, 1) \cdot \vec{N})$, which yields

$$\frac{1}{\sqrt{1+p^2+q^2}} = \sqrt{1-x^2-y^2}$$

- 2.2. What shapes can the shadow of a sphere take if it is cast on a plane and the source is a point source?

Solution:

These are conic sections with one important exception - you only get one half of the hyperbola.

- 2.3. We have a square area source and a square occluder, both parallel to a plane. The source is the same size as the occluder, and they are vertically above one another with their centers aligned.

- (a) What is the shape of the umbra?

Solution:

Square. This is rather a special case. You can construct the umbra by constructing all points on the plane that can see no part of the source. This is a square directly below the occluder.

- (b) What is the shape of the outside boundary of the penumbra?

Solution:

This is quite a nasty question. Easiest way to construct the answer is to think about an arbitrary point on the occluder; construct a cone over the source, whose vertex is this point. Now intersect that cone with the plane — the resulting square region on the plane consists of all points such that this point on the occluder blocks some point on the source. Finally, take the union of all such regions (i.e. over all points on the occluder); that's the penumbra. The envelope of the boundaries of these regions is the boundary of the penumbra. It's a square.

- 2.4. We have a square area source and a square occluder, both parallel to a plane. The edge length of the source is now twice that of the occluder, and they are vertically above one another with their centers aligned.

(a) What is the shape of the umbra?

Solution:

Depending on the distances between area source, occluder and plane, either there isn't one, or it's square.

(b) What is the shape of the outside boundary of the penumbra?

Solution:

same as in previous exercise; This is quite a nasty question. Easiest way to construct the answer is to think about an arbitrary point on the occluder; construct a cone over the source, whose vertex is this point. Now intersect that cone with the plane — the resulting square region on the plane consists of all points such that this point on the occluder blocks some point on the source. Finally, take the union of all such regions (i.e. over all points on the occluder); that's the penumbra. The envelope of the boundaries of these regions is the boundary of the penumbra. It's a square.

- 2.5. We have a square area source and a square occluder, both parallel to a plane. The edge length of the source is now half that of the occluder, and they are vertically above one another with their centers aligned.

(a) What is the shape of the umbra?

Solution:

Square

(b) What is the shape of the outside boundary of the penumbra?

Solution:

This is quite a nasty question. Easiest way to construct the answer is to think about an arbitrary point on the occluder; construct a cone over the source, whose vertex is this point. Now intersect that cone with the plane — the resulting square region on the plane consists of all points such that this point on the occluder blocks some point on the source. Finally, take the union of all such regions (i.e. over all points on the occluder); that's the penumbra. The envelope of the boundaries of these regions is the boundary of the penumbra. It's a square.

- 2.6. A small sphere casts a shadow on a larger sphere. Describe the possible shadow boundaries that occur.

Solution:

The question is slightly poorly posed, *mea culpa*, because it doesn't say what the source is. Assume it's a point source. There are two interesting cases: the point source is on the line connecting the two sphere centers; or it's off. If it's on, then the shadow must be a circle. If it's off, it's the intersection of a sphere and a right circular cone.

- 2.7. Explain why it is difficult to use shadow boundaries to infer shape, particularly if the shadow is cast onto a curved surface.

Solution:

As the previous exercise suggested, quite simple geometries lead to very complex shadows; it becomes hard to tell whether the shadow is complex because it was a simple shadow cone intersected with a complex surface, or because it was a complex shadow cone intersected with a simple surface.

- 2.8. As in Figure 2.18, a small patch views an infinite plane at unit distance. The patch is sufficiently small that it reflects a trivial quantity of light onto the plane. The plane has radiosity $B(x, y) = 1 + \sin ax$. The patch and the plane are parallel to one another. We move the patch around parallel to the plane, and consider its radiosity at various points.

- (a) Show that if one translates the patch, its radiosity varies periodically with its position in x .

Solution:

If you translate the patch by $(2\pi/a)$ units, it's in the same position with respect to the radiosity pattern as it was; and the plane is infinite.

- (b) Fix the patch's center at $(0, 0)$; determine a *closed form* expression for the radiosity of the patch at this point as a function of a . You'll need a table of integrals for this (if you don't, you're entitled to feel very pleased with yourself).

Solution:

This is worked out in detail in (Haddon and Forsyth 1997)

- 2.9. If one looks across a large bay in the daytime, it is often hard to distinguish the mountains on the opposite side; near sunset, they are clearly visible. This phenomenon has to do with scattering of light by air—a large volume of air is actually a source. Explain what is happening. We have modeled air as a vacuum and asserted that no energy is lost along a straight line in a vacuum. Use your explanation to give an estimate of the kind of scales over which that model is acceptable.

Solution:

During the day, the sun illuminates the air, so the air between the viewer and the far side of the bay glows (from the light scattered into the viewing direction). This reduces the contrast of the stuff on the other side of the bay — the viewer sees a bright light superimposed on so can't distinguish the edges of objects, etc. In the evening, the air does not glow, and so you can see the far side of the bay. Generally, you can think of air as a vacuum on scales up to kilometers; if you're very lucky and the air is very clear, perhaps 100 km.

- 2.10. Read the book *Colour and Light in Nature*, by Lynch and Livingstone, pub-

lished by Cambridge University Press, 1995.

PROGRAMMING EXERCISES

- 2.11.** An area source can be approximated as a grid of point sources. The weakness of this approximation is that the penumbra contains quantization errors, which can be quite offensive to the eye.

(a) Explain.

Solution:

As you move along the penumbra, you see an integer number of point sources; at various points, one comes in or out and so there is a visible jump in the intensity.

- (b) Render this effect for a square source and a single occluder casting a shadow onto an infinite plane. For a fixed geometry, you should find that as the number of point sources goes up, the quantization error goes down.
- (c) This approximation has the unpleasant property that it is possible to produce arbitrarily large quantization errors with any finite grid by changing the geometry. This is because there are configurations of source and occluder that produce large penumbrae. Use a square source and a single occluder, casting a shadow onto an infinite plane, to explain this effect.

Solution:

The trick to this is to have the plane at an angle, so that the penumbra is stretched out.

- 2.12.** Make a world of black objects and another of white objects (paper, glue, and spraypaint are useful here) and observe the effects of interreflections. Can you come up with a criterion that reliably tells, *from an image*, which is which? (If you can, publish it; the problem looks easy, but isn't.)

Solution:

Can't help you here; I've been looking for a good criterion for absolute lightness for years, as the exercise suggests.

- 2.13.** (This exercise requires some knowledge of numerical analysis.) Do the numerical integrals required to reproduce Figure 2.18. These integrals aren't particularly easy: if one uses coordinates on the infinite plane, the size of the domain is a nuisance; if one converts to coordinates on the view hemisphere of the patch, the frequency of the radiance becomes infinite at the boundary of the hemisphere. The best way to estimate these integrals is using a Monte Carlo method on the hemisphere. You should use importance sampling because the boundary contributes rather less to the integral than the top does.

Solution:

It could help to look at (Haddon and Forsyth 1997)

- 2.14.** Set up and solve the linear equations for an interreflection solution for the interior of a cube with a small square source in the center of the ceiling.
- 2.15.** Implement a photometric stereo system.
- (a) How accurate are its measurements (i.e., how well do they compare with known shape information)? Do interreflections affect the accuracy?
- (b) How repeatable are its measurements (i.e., if you obtain another set of images, perhaps under different illuminants, and recover shape from those, how does the new shape compare with the old)?

- (c) Compare the minimization approach to reconstruction with the integration approach; which is more accurate or more repeatable and why? Does this difference appear in experiment?
- (d) One possible way to improve the integration approach is to obtain depths by integrating over many different paths and then average these depths (you need to be a little careful about constants here). Does this improve the accuracy or repeatability of the method?

CHAPTER 3

Color

PROBLEMS

- 3.1.** Sit down with a friend and a packet of colored papers, and compare the color names that you use. You need a large packet of papers—one can very often get collections of colored swatches for paint, or for the Pantone color system very cheaply. The best names to try are basic color names—the terms *red*, *pink*, *orange*, *yellow*, *green*, *blue*, *purple*, *brown*, *white*, *gray* and *black*, which (with a small number of other terms) have remarkable canonical properties that apply widely across different languages (the papers in Hardin and Maffi (1997) give a good summary of current thought on this issue). You will find it surprisingly easy to disagree on which colors should be called blue and which green, for example.

Solution:

This one is really worth doing. In my experience, quite substantial disagreements appear, and you won't get a sense of this without trying it.

- 3.2.** Derive the equations for transforming from RGB to CIE XYZ and back. This is a linear transformation. It is sufficient to write out the expressions for the elements of the linear transformation—you don't have to look up the actual numerical values of the color matching functions.

Solution:

Write $R(\lambda)$ for the R primary, $G(\lambda)$ for the green primary, $B(\lambda)$ for the B primary. A color $E(\lambda)$ has RGB coordinates r, g, b if it matches $rR(\lambda) + gG(\lambda) + bB(\lambda)$. Write $\phi_x(\lambda)$ for the colormatching function for the X primary, $\phi_y(\lambda)$ for the colormatching function for the Y primary, and $\phi_z(\lambda)$ for the colormatching function for the Z primary. Then $rR(\lambda) + gG(\lambda) + bB(\lambda)$ will have X coordinate $\int (rR(\lambda) + gG(\lambda) + bB(\lambda))\phi_x(\lambda)d\lambda$. By this reasoning,

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} \int R(\lambda)\phi_x(\lambda)d\lambda & \int G(\lambda)\phi_x(\lambda)d\lambda & \int B(\lambda)\phi_x(\lambda)d\lambda \\ \int R(\lambda)\phi_y(\lambda)d\lambda & \int G(\lambda)\phi_y(\lambda)d\lambda & \int B(\lambda)\phi_y(\lambda)d\lambda \\ \int R(\lambda)\phi_z(\lambda)d\lambda & \int G(\lambda)\phi_z(\lambda)d\lambda & \int B(\lambda)\phi_z(\lambda)d\lambda \end{pmatrix} \begin{pmatrix} r \\ g \\ b \end{pmatrix}$$

and this works the other way round if one swaps $R(\lambda)$ for $X(\lambda)$, etc. and $\phi_x(\lambda)$ for $\phi_r(\lambda)$.

- 3.3.** Linear color spaces are obtained by choosing primaries and then constructing color matching functions for those primaries. Show that there is a linear transformation that takes the coordinates of a color in one linear color space to those in another; the easiest way to do this is to write out the transformation in terms of the color matching functions.

Solution:

Take the result of the previous exercise; now use the primaries of the first color space and the color matching functions of the second.

- 3.4. Exercise 3 means that, in setting up a linear color space, it is possible to choose primaries arbitrarily, but there are constraints on the choice of color matching functions. Why? What are these constraints?

Solution:

When you choose colormatching functions, the new ones must be a linear combination of the colormatching functions for any other linear color space. For example, think of RGB and XYZ. The r coordinate of a color matching some $E(\lambda)$ is obtained by forming $\int \phi_r(\lambda)E(\lambda)d\lambda$; it can also be obtained as a linear combination of the X, Y, and Z coordinates (as above). This is true for any $E(\lambda)$, and so means that $\phi_r(\lambda)$ is a linear combination of $\phi_x(\lambda)$, $\phi_y(\lambda)$, and $\phi_z(\lambda)$.

- 3.5. Two surfaces that have the same color under one light and different colors under another are often referred to as *metamers*. An *optimal color* is a spectral reflectance or radiance that has value 0 at some wavelengths and 1 at others. Although optimal colors don't occur in practice, they are a useful device (due to Ostwald) for explaining various effects.

(a) Use optimal colors to explain how metamerism occurs.

Solution:

For example, mixing red and green *light* gets yellow. So if we have one surface albedo that is zero everywhere except in the yellow range, and another that is zero except in green and red, the two surfaces will look the same (yellow) in white light. But now if we obtain a light that is zero in the red range, one surface will look yellow and the other will look green.

(b) Given a particular spectral albedo, show that there are an infinite number of metameric spectral albedoes.

Solution:

Assume we have illumination that is uniform over albedo (white). Now this spectral albedo reflects some color of light. Choose three distinct wavelengths (say, the RGB primaries) as primaries, and determine the weights for these primaries to match the reflected light. You can then build an albedo that matches by setting the albedo to these weights at the chosen wavelengths, and zero elsewhere. There is clearly an infinite number, even though for some triples the construction will fail because subtractive matching applies.

(c) Use optimal colors to construct an example of surfaces that look different under one light (say, red and green) and the same under another.

Solution:

See the solution to the question before the last.

(d) Use optimal colors to construct an example of surfaces that swap apparent color when the light is changed (i.e., surface one looks red and surface two looks green under light one, and surface one looks green and surface two looks red under light two).

Solution:

Apply the strategy of the solution to the question before the last.

- 3.6. You have to map the gamut for a printer to that of a monitor. There are colors in each gamut that do not appear in the other. Given a monitor color that can't be reproduced exactly, you could choose the printer color that is closest.