Computational Vision
U. Minn. Psy 5036
Daniel Kersten
Lecture 16: Shape from X

Initialize

■ Spell check off

```
Off[General::spell1];
```

```
Off[General::spell1];
<< Graphics`
<< Graphics`Graphics3D`
<< Graphics`PlotField3D`
<< Graphics`PlotField`
<< Graphics`Polyhedra`
```

You have to make sure these 3rd party add-in packages are present, and then find the path to set them:

In[7]:=
```
SetDirectory[DirectoryName[Experimental`FileBrowse[False]]]
```

Out[7]= /Users/kersten/Sites/kersten−lab/courses/Psy5036W2006/Lectures/16. Shape−from−X

In[8]:=
```
<<ImplicitSolids.m
```

In[9]:=
```
<<RDS.m
```

Outline

Last time

■ **Geometry,shape and depth: Representation issues & generative models**

■ **Lambertian model & surface normal representations**

Surface normal representation: Local, dense, metric, viewer-dependent coordinate system (e.g. slant from observer).

Today

- **Inference of shape from shading--set in context of other cues to shape,**

- **Perception of shape from shading**

- **Formal ambiguities in the generative model, the bas-relief transform**

- **Task analysis: "Shape for X" problem**

"Shape from X"

- **What is shape?**

One mathematical definition: "Geometrical properties that are invariant over translation and scale".

A computational vision definition: "Whatever is left over after discounting material, illumination and viewpoint variations in the images of an object"

We will take a more general view here:

"geometrical relationships within a surface that are useful for visual functions". Later we will return to the "discounting" issues.

There are a variety of cues to shape. In natural images, these cues typically co-vary. Later on we will talk about cue integration. In a psychophysics lab, they can be manipulated independently.

Texture

Later on we will talk in more detail about how to model and infer surface material properties. Texture is one property of surfaces that is particularly informative for shape and object identity. Texture by definition has some degree of regularity. Texture can be highly regular, or only regular in the statistical sense. For the moment, suppose a surface has a deterministic regular homogeneous pattern of "texture elements", where each texture element is the same 3D size, and they are distributed homogeneously.

■ Slant & texture

Define a function to place circular disks on a regular grid

```
checker[x_,y_,space_,radius_]:=
If[(Mod[x,space]^2+Mod[y,space]^2<radius^2) ||
(Mod[-x,space]^2+Mod[-y,space]^2<radius^2) ||
(Mod[x,space]^2+Mod[-y,space]^2<radius^2) ||
(Mod[-x,space]^2+Mod[y,space]^2<radius^2),0,1];
```

```
Plot3D[{1, GrayLevel[checker[x,y,32,12]]},{x, 1, 255}, {y, 1, 255},
PlotPoints->128,Mesh->False, Axes->False, Boxed->False];
```

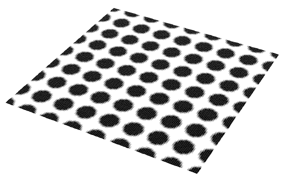

Try changing the orientation of the above texture

The generative assumptions above lead to regularities in the image that can be used to estimate surface slant--if the spatial scale of the texture is small in comparison to the surface curvature, so that the surface can be approximated locally by a plane.

Then, cues to surface slant:

1) spatial gradient of the density of texture elements: the image density increases with distance;

2) an individual element (such as a circular disk) gets smaller with distance, and its image height to width ratio gets smaller ("compressed") with increasing slant;

3) the ratio of the back width to the front width of an individual element gets smaller with slant (imagine a small square, linear pespective on small scale) . (See Knill).

Tilt is the direction that the surface slants away most rapidly. Here the density and sizes of the elements change the most rapidly.

■ Shape & texture

```
bump[x_,y_,R_]:=If[x^2+y^2>R^2,0,Sqrt[R^2-x^2-y^2]];
```

```
Plot3D[{bump[x-128,y-128,95], GrayLevel[checker[x,y,24,8]]},{x, 1, 255},
{y, 1, 255}, PlotPoints->128,Mesh->False, Axes->False, Boxed->False,
AspectRatio->Automatic, LightSources→{{{1.,0.,1.},RGBColor[1,1,1]}},
ViewPoint→{0,0,2},PlotRange->{0,95}];
```

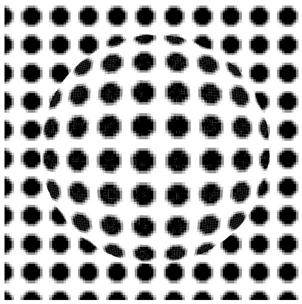

Change the above bump so that it is wider than it is tall

Stereo disparity

■ Random dot stereograms, "Magic eye" style

```
RDSPlot[-Cos[Sqrt[x^2+(2 y)^2]], {x,-10,10}, {y, -5, 5},
    AspectRatio -> 1, PlotPoints-> {400,400}, Density->0.1, Periods->6,
Guides->True];
    Show[%];
```

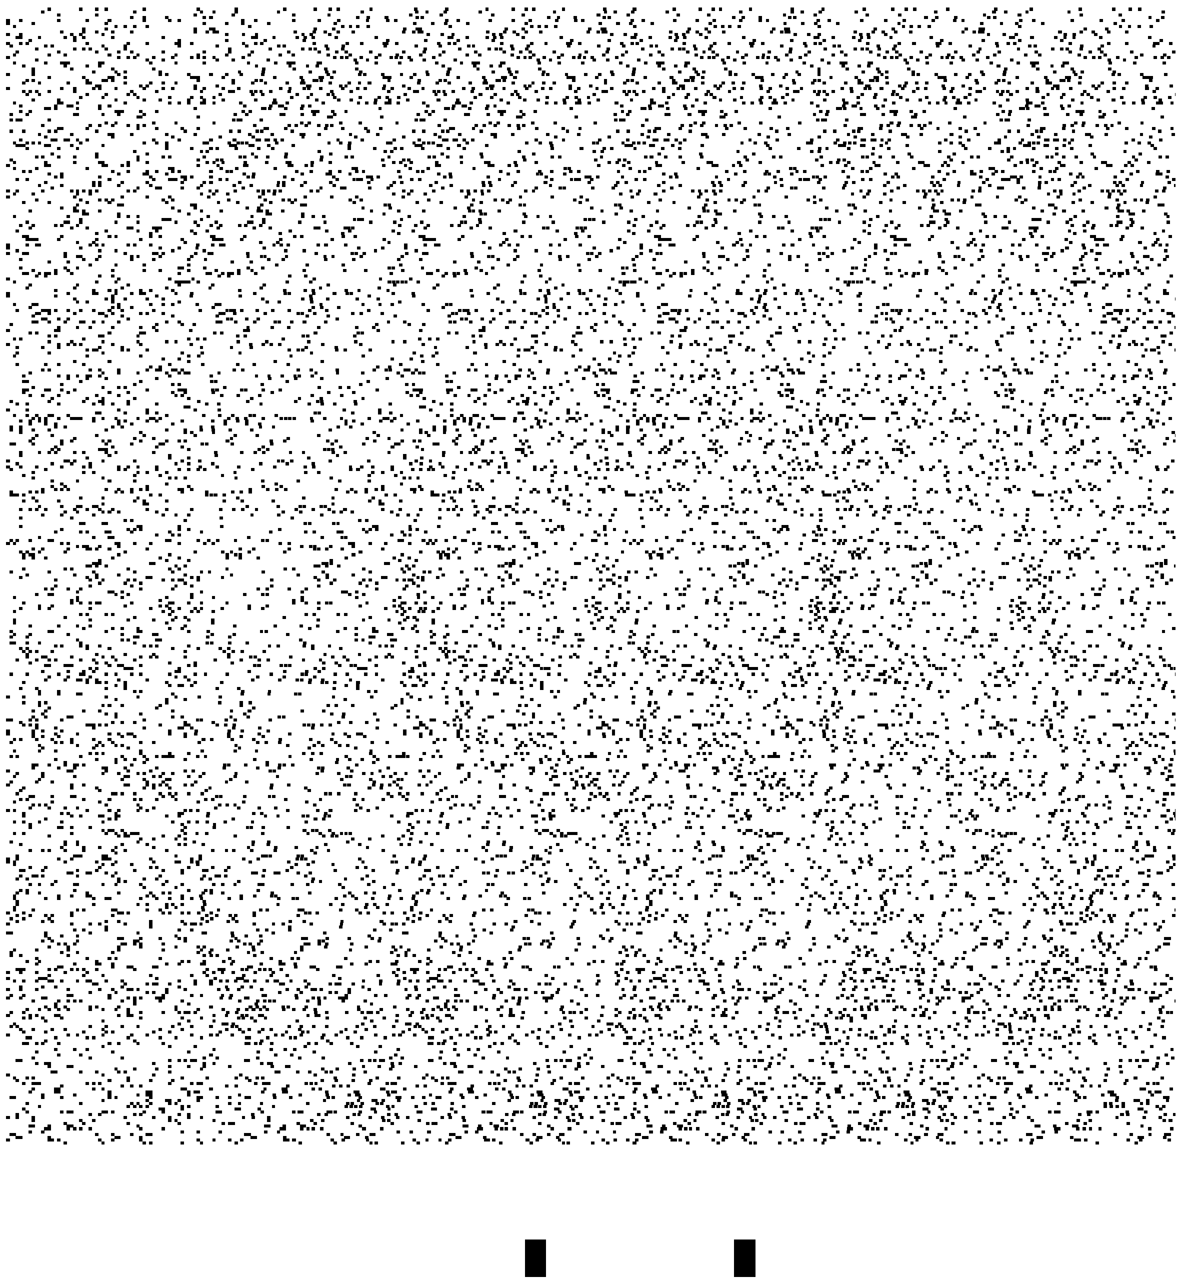

Cross your eyes so that you see three dots (rather than two) at the bottom. The above surface should appear like the one below, but viewed from above.

■ Stereo + shading + surface contours

```
pL=Plot3D[Cos[Sqrt[x^2+y^2]], {x,-10,10}, {y, -10, 10},
    AspectRatio -> Automatic, PlotPoints-> {25,25}, ViewPoint→
{1.3,-2.4,2.}];
```

```
pR=Plot3D[Cos[Sqrt[x^2+y^2]], {x,-10,10}, {y, -10, 10},
    AspectRatio -> Automatic, PlotPoints-> {25,25}, ViewPoint→
{1.1,-2.4,2.}];
```

```
Show[GraphicsArray[{pL, pR}]];
```

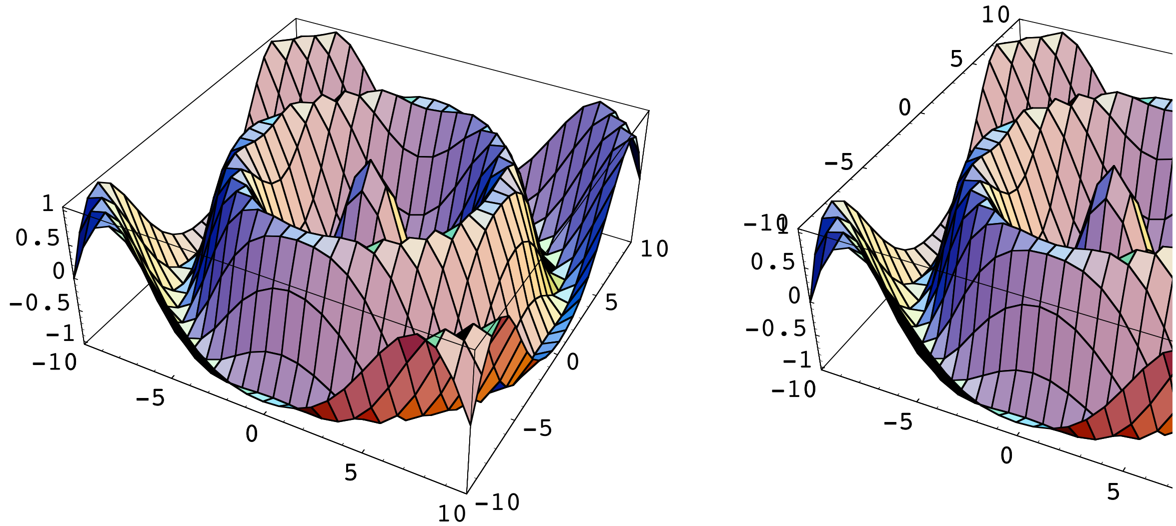

Motion

We'll look at structure & shape from motion in detail later. With some care, one can put up an animation in which any single frame shows no apparent structure, but when moving, the shape becomes clear. E.g. rotating "glass" cylinder with dots on it. Or, "biological motion".

In the following, we show shading, varioius kinds of contour information, and motion. Play the following pair as an animation:

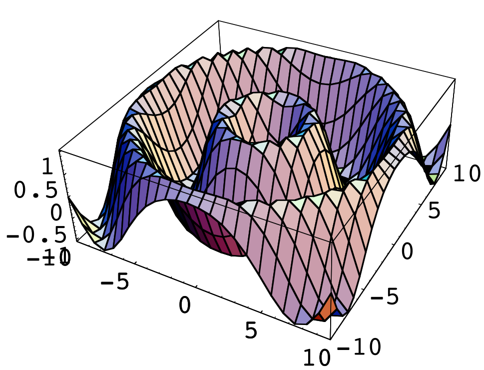

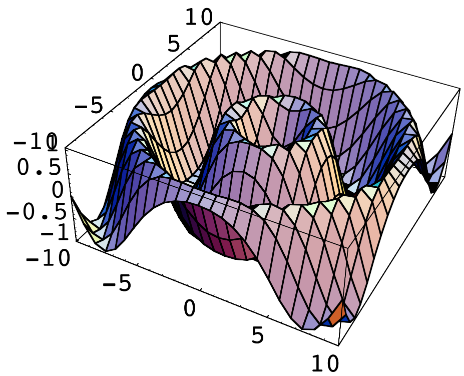

Contours

■ Surface contour (markings)

```
Plot[{Sin[x],Sin[x]+1,Sin[x]+2, Sin[x]+3, Sin[x]+4, Sin[x]+5},{x,0,10},
Axes->False];
```

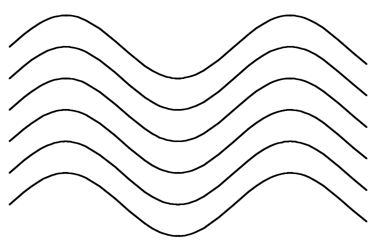

■ Surface contours (creases or orientation discontinuities)

The lines mark surface orientation discontinuities.

```
Show[ Graphics3D[ GreatDodecahedron[ ] , Shading->False, Boxed->False]];
```

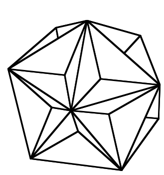

■ Bounding contours, (depth discontinuities), e.g. silhouette

Orientation discontinuities can coincide with depth discontinuities.

```
Show[ Graphics3D[ GreatDodecahedron[ ] , Lighting->False, Boxed->False]];
```

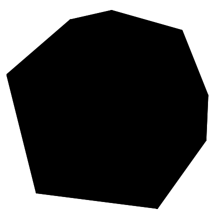

■ ...and smooth occluding contours

In this smooth shape, orientation of a unit normal vector changes smoothly at the depth discontinuities of the boundary.

```
dumbell := (x^2 + y^2 + 4 y^2/(1+32.0 x^2) + z^2 + 4 z^2/(1+32.0 x^2));
```

```
SphericalRender[dumbell,1,32,AxesLabel->{"X","Y","Z"},Lighting->False,
Boxed->False]
```

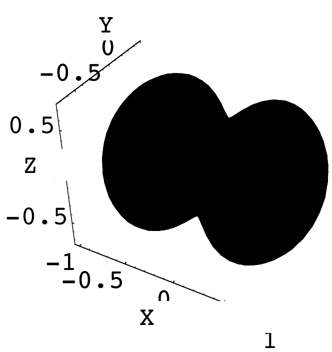

Shading...

Shape from shading

Introduction

The ambiguity of shading has had a newsworthy impact for well over a century, beginning with the "canals of Mars". From 1877, Giovanni Virginio Schiaparelli (an Italian astronomer) reported about a hundred of them attracting world-wide attention. The american astronomer Percival Lowell thought the markings were vegetation, several kilometres wide,bordering irrigation canals dug by intelligent life to carry water from the poles. However, most astronomers couldn't see the canals. Photography through the Earth's atmosphere offered no solution because the lines were near the limit of resolution of the human eye and the physical optics. The controversy was settled in 1969 when photographs were taken from several hundred kilometres above the surface of Mars by the Mariner 6 and 7 spacecraft. These showed many craters and other formations but no canals.

However, the phenomenon of seeing intriguing shapes from shading hadn't gone away even in the late 20th century. With the Viking orbiter, NASA photographs have more interesting pictures from Mars:

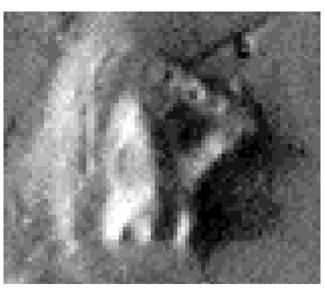
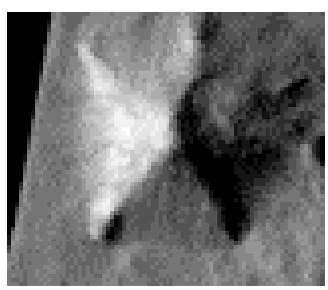
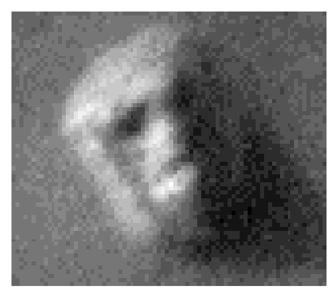

Are these shapes really there? If not, why do we see the particular shapes that we do?

The ambiguities

■ Overview

The following figure from William Freeman at MIT illustrates the basic light-direction/shape ambiguity problem of shape from shading:

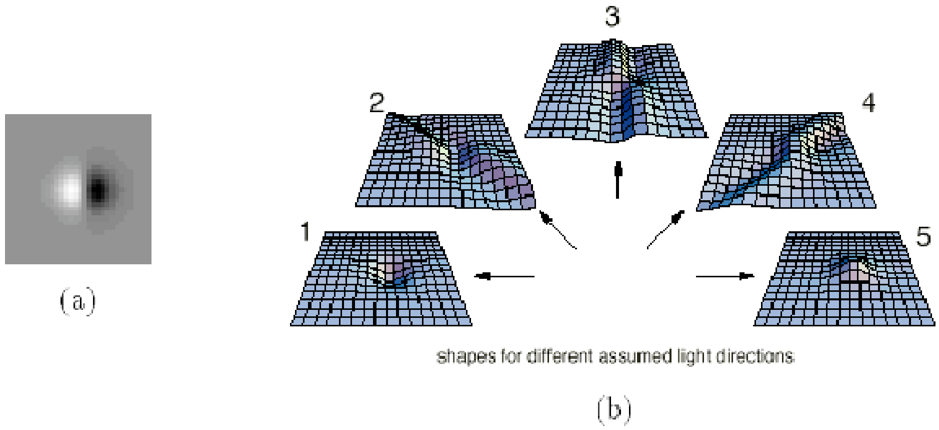

Figure 6: (a) Perceptually, this image has two possible interpretations. It could be a bump, lit from the left, or a dimple, lit from the right. (b) Mathematically, there are many possible interpretations. For a sufficiently shallow incident light angle, if we assume different light directions, we find different shapes, each of which could account for the observed image.

(Figure from: Mitsubishi Tech Report, TR93-15. http://www.merl.com/people/freeman/publications.html. See too: W. T. Freeman, Exploiting the generic viewpoint assumption, International Journal Computer Vision, 20 (3), 243-261, 1996.)

One solution to resolve the ambiguity is to use a surface smoothness constraint (e.g. as a Bayesian prior). Another is to approximate the shading as linear. A third is to learn a shape estimator. These methods include adopting prior constraints, "active vision" in which looks for new information to relieve ambiguity, and the combination of information from other modules (e.g. illumination direction estimator).

Marginalization over illumination direction also provides a means to reduce shading ambiguity (Freeman, 1994). The principle is analogous to the general viewpoint constraint, but now applied to another secondary variable, illumination:

Vision should pick the shape interpretation that is most robust to variations in illumination direction.

Recall that scene-from-image problems, from a formal perspective, can treated as an inverse problem. For shape from shading, we will follow a similar strategy. As we've noted before, the theoretical approach to inverse problems can be formalized in terms of Bayesian estimation.

Classic regularization approach to shape from shading (Ikeuchi and Horn, 1981)

Write down the generative model for image luminance L in terms of geometrical gradient space parameters p, q:

$$L_R = re\,\hat{N} \cdot \hat{E}$$

where r is the reflectance, and e the strength of illumination. N and E are as in the last lecture. E is a point light source (we'll assume e=1), and E's direction can be expressed in gradient space:

$$\hat{E} = \frac{(p_e, q_e, -1)}{\sqrt{p_e^{\,2} + q_e^{\,2} + 1}}$$

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

Taking the dot product, we have:

$$L_R(p,q) = \frac{r(pp_e + qq_e + 1)}{\sqrt{(p^2 + q^2 + 1)(p_e^{\,2} + q_e^{\,2} + 1)}}$$

Our data is L(*x,y*). We require that L(*x,y*)=LR(*p,q*). One problem is that we do not know r or (*ps,qs*). (Note that in general, r is not going to be a constant, but rather some pattern of pigmentation, r(x,y)) Even if we do, we still have many p's and q's which will satisfy L-LR=0. Isophotes (lines of constant luminance in the image) do not necessarily imply constant (p,q):

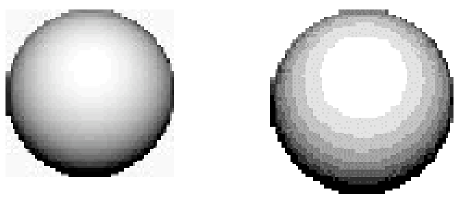

Shape from shading is under-constrained or "ill-posed" as a mathematical problem.

A solution to ambiguity (that we'll explore in detail later in optic flow motion field estimation) is to impose a smoothness constraint (a local constraint) and boundary conditions (global constraint).

Implementing a smoothness constraint.

Let us follow the approach of Ikeuchi and Horn (1981), and enforce a surface smoothness assumption (or constraint) by requiring that the squared values of the Laplacians of p and q be small. The Laplacians are:

$$\nabla^2 p = \frac{\partial^2 p}{\partial x^2} + \frac{\partial^2 p}{\partial y^2}$$

$$\nabla^2 q = \frac{\partial^2 q}{\partial x^2} + \frac{\partial^2 q}{\partial y^2}$$

We then construct a local cost function function:

$$e(x, y) = (L(x, y) - L_R(p, q))^2 + \lambda((\nabla^2 p)^2 + (\nabla^2 q)^2)$$

We want e(x,y) to be small. This will be the case if (p,q) is chosen so that the first term tends toward zero, and if the spatial derivatives are small (i.e. because we've chosen p,q to change slowly). But we want this local cost function to be small over the whole image.

So we construct a global cost functional by integrating over the whole image:

$$e[p(x,y),\; q(x,y] = \iint (L(x,y) - L_R(p,q))^2 + \lambda((\nabla^2 p)^2 + (\nabla^2 q)^2)\; dx\, dy$$

The goal is to find (p,q) at each point (x,y) which makes the global error or cost, e[p,q] as small as small as possible. λ is a Lagrange multiplier--a weighting function on the smoothness. For example, it should be big if we have reason to believe our data are noisy--i.e. we want to trust our prior smoothness constraint over the unreliable data.

The boundary values can also be used to constrain the solution, and may be necessary at times to produce a unique solution. They are also important to prevent smoothing over discontinuities. What are the boundary conditions?

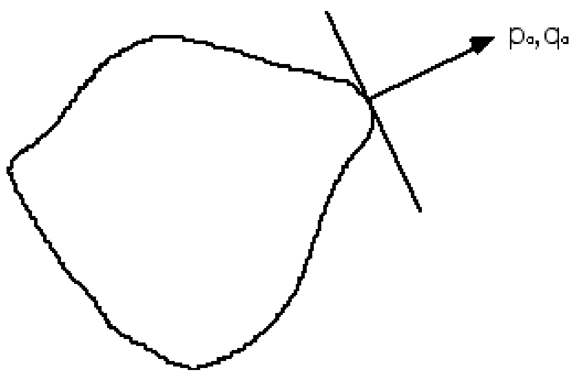

The surface normal can be read directly off the image if the boundary contours are known. But we have a problem, at smooth occluding bounding contours, po and qo (the spatial rates of change of the surface depth away from the viewpoint) are infinite. And if the bounding contours are not smooth, but sharp, the surface normal is not defined there.

One solution is to use coordinates that don't blow up at boundaries. Stereographic coordinates have also been suggested as a means to avoiding infinities at the boundaries (Ikeuchi and Horn, 1982).

$$f = \frac{2p\;\sqrt{1+p^2+q^2} - 1}{p^2+q^2}$$

$$g = \frac{2q\;\sqrt{1+p^2+q^2} - 1}{p^2+q^2}$$

Another representation that will avoid the boundary problem is slant and tilt introduced earlier (See too, Mamassian and Kersten (1995)).

It can be shown that minimizing the above error function is equivalent to maximizing the posterior probability of the map of surface p(x)'s and q(x)'s conditional on L(x,y). Using Bayes, the generative model determines the likelihood, and the laplacians of the gradient parameters determines the smoothness prior.

Linear approximation to the Lambertian generative model for shading

The image formation constraint that we have used is non-linear. It is worth asking to what degree a linear approximation would be adequate-- a question which has been addressed by Knill and Kersten (1990) and by Pentland (1990). An alternative approach to shape from shading was taken in a series of computational studies that sought to learn scene parameters from images (See too: Kersten, D., O'Toole, A. J., Sereno, M. E., Knill, D. C., & Anderson, J. A. ,1987; Lehky, S. R., & Sejnowski, T. J., 1988.)

Let's derive a linear approximation to:

$$L_R(p,q) = \frac{r(pp_e + qq_e + 1)}{\sqrt{(p^2 + q^2 + 1)(p_e{}^2 + q_e{}^2 + 1)}}$$

```
Lmodel[pn_,qn_,pe_,qe_]:= ({pn,qn,-1}/Sqrt[pn^2+qn^2 + 1]).({pe,qe,-1}/Sqrt[pe^2+qe^2 + 1]);
```

where we use the notation p->pn, q->qn.

■ Taylor's series

We can expand the image luminance in a Taylor series about {pn,qn}={0,0}. Recall that this is done by calculating successive derivatives at {pn,qn}={0,0}, and then substituting {pn,qn}->{0,0}. The first derivative with respect to pn is:

```
D[Lmodel[pn,qn,pe,qe],pn]/.{pn->0,qn->0}
```

$$\frac{\text{pe}}{\sqrt{\text{pe}^2 + \text{qe}^2 + 1}}$$

Note that this is the cosine of the slant of the light source (see previous lecture).

Mathematica's **Series[]** function puts the Taylor series together for us. Here is the expansion up to linear terms:

```
Series[Lmodel[pn,qn,pe,qe], {pn, 0, 1}, {qn, 0, 1}]
```

$$\left(\frac{1}{\sqrt{\text{pe}^2 + \text{qe}^2 + 1}} + \frac{\text{qe qn}}{\sqrt{\text{pe}^2 + \text{qe}^2 + 1}} + O(\text{qn}^2)\right) + \left(\frac{\text{pe}}{\sqrt{\text{pe}^2 + \text{qe}^2 + 1}} + O(\text{qn}^2)\right)\text{pn} + O(\text{pn}^2)$$

You can improve the approximation by including quadratic terms: **Series[Lmodel[pn,qn,pe,qe], {pn, 0, 2}, {qn, 0, 2}].**

■ **Try it out on bump[x_,y_]:=(1/4) (1-1/(1+Exp[-10 (Sqrt[x^2+y^2]-1)]));**

```
bump[x_,y_]:=(1/4) (1-1/(1+Exp[-10 (Sqrt[x^2+y^2]-1)]));
pn[x_,y_]:= Evaluate[D[bump[x,y],y]];
qn[x_,y_]:= Evaluate[D[bump[x,y],x]];
```

Light source:

```
{pe,qe}={1,1};
```

$$\text{Lapprox}[x\_, y\_] := \frac{1}{\sqrt{1 + pe^2 + qe^2}} + \frac{qe * qn[x, y]}{\sqrt{1 + pe^2 + qe^2}} + \frac{pe * pn[x, y]}{\sqrt{1 + pe^2 + qe^2}};$$

```
Lapprxom[x_,y_] := (1+qe*qn[x,y] + pe*pn[x,y])/Sqrt[1+pe^2 + qe^2];
```

```
DensityPlot[Lapprxom[x, y], {x, -3, 3}, {y, -3, 3}, Mesh → False,
  Frame → False, PlotPoints → 128, PlotRange → {0, 1}];
```

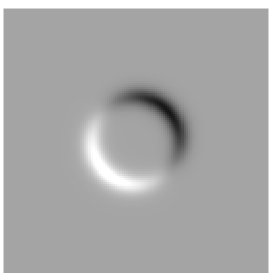

Use Plot3D to make a grayscale shaded plot of the bump function. Compare.

We can compare this approximation with the full lambertian rendering of this bump using the code in the "Bas relief transform" section later.

How could you use the above linear generative model for "shading from shape" to devise a solution to the inverse problem of "shape from shading"?

Is shape from shading really ill-posed? Formal (exact) solution to the classic problem: lambertian, point light source

Given a Lambertian reflectance, and a point light from the camera, the shape from shading problem is well-posed and a solution can be found (Dupuis and Oliensis, 1994). Oliensis showed that 1) this problem is generally **well-posed,** contrary to previous belief---and that when the light comes from the camera direction a shaded image determines the depicted surface **uniquely;** 2) that the image of a surface's **occluding boundary** (the edge where it rounds away from the viewer) is **not useful** for determining the surface---again contrary to previous belief; 3)that for an arbitrary shaded image there is generally no physical (differentiable) surface corresponding to it (**nonexistence**).

These are theoretical results--it is quite another matter to discover the built-in knowledge that the human visual system has regarding shape from shading, for example the boundary information is quite important (see below).

Shape from shading on a "cloudy day"--diffuse lighting

Most theoretical work on shape from shading has started from point light sources, but diffuse lighting is the more realistic case. The generative physical problem becomes complicated. Ray-tracing and radiosity methods provide forward generative models (eg. Greenberg, 1989; Foley et al., 1990). But the straightfoward inverse problem is impractical--can't check out all the ray-tracing bounces!

For an elegant solution to the problem of shape-from-shading on a cloudy day, see: Langer, M. S. & Zucker, S. W. (1994)

Perception

Perception of shape from shading & lighting direction

Some studies seem to indicate that human observers are not very accurate at estimating the local orientation of surface normals. We have a bias towards the fronto-parallel plane.

We also seem to have difficulty in estimating the slant of a light source (Todd, J. T., & Mingolla, E., 1983; Mingolla, E., & Todd, J. T. 1986). We are better at the tilt.

One problem with our above shape from shading analysis is that we assumed light source direction is known.

In fact, *human observers do seem to assume that the light source is from above*. This can be seen in the age-old "crater illusion" in which we have vertical luminance gradients of intensity giving an impression of a convex object on the left, and a concavity on the right (Ramachandran, V. S., 1988). (See previous lecture and code below that makes a bump lit from above and one that is upside-down). Interestingly, there also seems to be a slight bias towards assuming the light source is coming from the left (Sun & Perona, 1998).

Illumination direction & shape ambiguity

- **Don't open cell unless curious--just uses code from last time to make a hemisphere for the shape-convexity ambiguity demo**

- Lambertian sphere

```
bump[x_, y_] := If[x^2 + y^2 > 1, 0, Sqrt[1 - x^2 - y^2]];
g1 = Plot3D[bump[x, y], {x, -3, 3}, {y, -3, 3}, PlotPoints → 64,
   PlotRange → {0, 3}, Mesh → False,
   LightSources → {{{1., 0., 1.}, RGBColor[1, 0, 0]},
     {{1., 0., 1.}, RGBColor[0, 1, 0]}, {{1., 0., 1.}, RGBColor[0, 0, 1]}},
   ViewPoint → {1, 1, 1}, AxesLabel → {"x", "y", "z"}, Ticks → False,
   AspectRatio → 1];
```

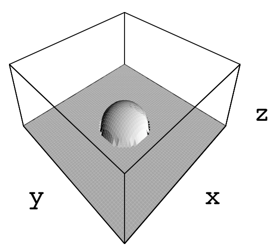

- Calculate surface normals of surface

```
Clear[nx, ny]; little = 0.001; big  =  1000;

nx[x_, y_] := Evaluate[D[bump[x, y], x]];
ny[x_, y_] := Evaluate[D[bump[x, y], y]];
nx[x_, y_] := big /; x^2 + y^2 == 1
ny[x_, y_] := big /; x^2 + y^2 == 1
nx[x_, y_] := 0. /; {x^2 + y^2 > 1}
ny[x_, y_] := 0. /; x^2 + y^2 > 1
```

nx[x,y] is $\frac{\partial z}{\partial x}$, and similarly for ny. The rate of change of depth range is greatest as the face slopes away from the viewpoint:

■ Lambertian rendering: specification for normals, light, reflectance

■ Unit surface normals

Normalize the surface normal vectors to unit length:

```
normface[x_, y_] := -{nx[x, y], ny[x, y], 1} /
    (Sqrt[nx[x, y]^2 + ny[x, y]^2 + 1]);
PlotVectorField[{normface[x, y][[1]], normface[x, y][[2]]},
  {x, -1.2, 1.2}, {y, -1.2, 1.2}];
```

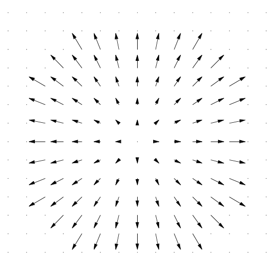

■ Render bump surface

```
a[x_, y_] := 1;
s = {-10, 100, 10}; s = N[s / Sqrt[s.s]];
b1 = DensityPlot[a[x, y] * normface[x, y].s, {x, -1.2, 1.2},
    {y, -1.2, 1.2}, Mesh → False, PlotPoints → 64, Frame → False,
    DisplayFunction → Identity, PlotRange → {-1, 1}];
s = {-10, -100, 10}; s = N[s / Sqrt[s.s]];
b2 = DensityPlot[a[x, y] * normface[x, y].s, {x, -1.2, 1.2},
    {y, -1.2, 1.2}, Mesh → False, PlotPoints → 64, Frame → False,
    DisplayFunction → Identity, PlotRange → {-1, 1}];
```

```
Show[GraphicsArray[{b1,b2}],DisplayFunction→$DisplayFunction];
```

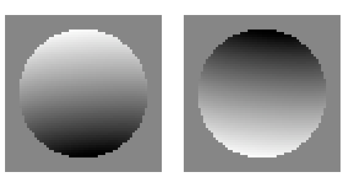

Shading and contour interaction

A challenging problem for future research is understanding how shape-from-contour interacts with shape from shading. Contour can override shading information. Let's cut out sections of part of an apparent sphere:

```
s={-10,100,10}; s=N[s/Sqrt[s.s]];
temp=Transpose[Table[normface[x,y].s,{x,-1.2,1.2,0.05},{y,-1.2,1.2,0.05}]
];
b1=ListDensityPlot[temp,Mesh→False,Frame→False, PlotRange→{-1,1}];
```

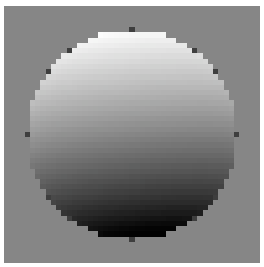

..this patch is an apparent cylinder:

```
s={-10,100,10}; s=N[s/Sqrt[s.s]];
temp=Transpose[Table[If[(x<-1/4 || y<-1/4) || (x>3/4 ||
y>3/4),0,normface[x,y].s],{x,-1.2,1.2,0.05},{y,-1.2,1.2,0.05}]];
ListDensityPlot[temp,Mesh→False,Frame→False, PlotRange→{-1,1}];
```

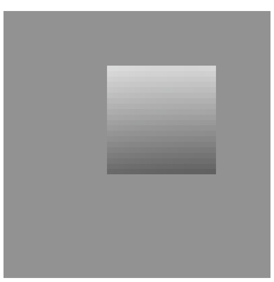

This patch is an apparent diamond:

```
s={-10,100,10}; s=N[s/Sqrt[s.s]];
temp=Transpose[Table[If[Abs[x]+Abs[y]>.5,0,normface[x,y].s],{x,-1.2,1.2,0
.02},{y,-1.2,1.2,0.02}]];
ListDensityPlot[temp,Mesh→False,Frame→False, PlotRange→{-1,1}];
```

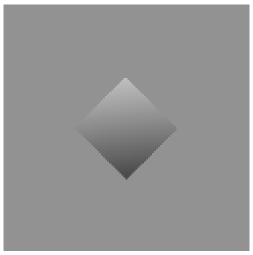

Somehow the visual system incorporates boundary conditions. This could be done either through low-level "propagation" analogous to the Ikeuchi & Horn algorithm, or it could be done by using the contour to "index" shape classes (polyhedra, cylinders, spheres, pills, chiclets, etc..) with which to infer the interior shape.

Another demo (due to David Knill) you can try to make yourself is generate a vertical sine-wave grating. Make the contours at the top and botoom have either 1) the same frequency as the grating or 2) twice the frequency.

One research problem is how do specularities affect the convex vs. concave perception in the above figures? It has been found that human observers can use the stereoscopic position of a specularity (which is in front of a concave surface, and behind a convex surface) to disambiguate the shape (Blake and Bülthoff, 1990; 1991).

■ Project idea

Here is a simple project idea. One would expect that the influence of contour on shape-from-shading would depend on whether the contour is "attached" to the shaded region or not. A simple way of manipulating this is to use stereo. For example, one could make the boundary of a sphere to appear either at the correct boundary (i.e. towards the back of the bulge), or in front of the bulge, as if it is a hole. Perceived shape should be affected. One might even be able to design a shading pattern that looks very different depending on what the boundary gets attached to. For a discussion of "intrinsic" vs. "extrinsic" contours, see: Nakayama, K., & Shimojo, S. (1992).

Suppose that the diamond boundary below appeared as a diamond-shaped whole. Would one still interpret the shading as a pill, or would the shape of the internal region now appear more spherical?

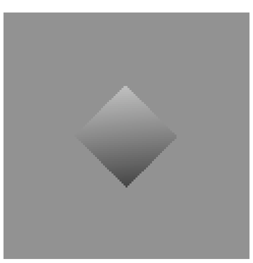

■ Psychophysics

Mamassian and Kersten (1995) have compared the kinds of psychophysical errors made in slant and tilt judgements with a shape-from-shading algorithm that assumes smoothness in slant and tilt and knows the boundary values. This algorithm is optimal in a Bayesian sense under the specified constraints. They investigated the perception of local surface orientation for a simple object with regions of elliptic (egg-shaped) and hyperbolic (saddle-shaped) points. The local surface orientation was measured by the slant and tilt of the tangent plane at different points of the surface under several different illumination conditions. They found an underestimation of the perceived slant, and a larger variance for the perceived tilt. They also found that subjects were better at estimating the surface orientation when the shape was locally egg-shaped rather than saddle-shaped or cylindrical. From converging evidence based on (i) the light direction most consistent with the observer's settings, (ii) a supplementary experiment where the object is displayed as a silhouette, and (iii) the computer simulations of the shape from shading algorithm, they concluded that the occluding contour was the dominant source of information used by the observers.

At this point, very little is known about the neurophysiology of shape from shading/contour, although there has been speculation about the role of V1 spatial filters (e.g. simple cells) in the local estimate of shape; see some ideas by Lehky and Sejnowski (1988), and by Pentland (1982). Knill and Kersten showed that orthogonal oriented linear filters could be used to estimate surface normals for Lambertian surfaces; however, we have very little idea of how neural systems may represent parameters of shape. Even basic issues such as viewpoint dependent (via extrinsic geometry) vs. viewpoint dependent (via intrinsic geometry) representations are unclear.

■ Neuroimaging

There has been an increase in interest in studying the neural basis of object and shape perception using fMRI. For example, see: Kourtzi and Kanwisher (2000) and Moore and Engel, 2001).

■ How to integrate? Shape from X1 and X2 and X3

Later in the course, we will look at the problem if cue integration for depth and shape.

■ How to pick? Shape from X1 or X2 or X3

How should a visual system decide what kind of shape cue it has (shading from lambertian material, shading from shiny material, shading due to texture), and thus what kind of algorithm to use? One approach is to run a set of experts in parallel and pick the solution from the one that seems to "know what it is doing". Another approach is to use robust algorithms that are not sensitive to the details of the cues.

Next time

Bayes formulation of shape from shading & bas-relief

Introduction to motion analysis

Appendices

```
Import[Experimental`FileBrowse[False]];
```

- (under construction) we can also plot the unit surface normals on the surface itself in 3D

```
bumparray =
  Flatten[Table[N[{{x, y, bump[x, y]}, normface[x, y]}],
    {x, -1 + ee, 1 - ee .0 .1}, {y, -1 + ee, 1 - ee, 0.1}], 1];
```

```
ListPlotVectorField3D[bumparray, ViewPoint → {1, -1, 3}, AspectRatio → 1,
  PlotRange → {{-1 + ee, 1 - ee}, {-1 + ee, 1 - ee}, {0, 1}}];
```

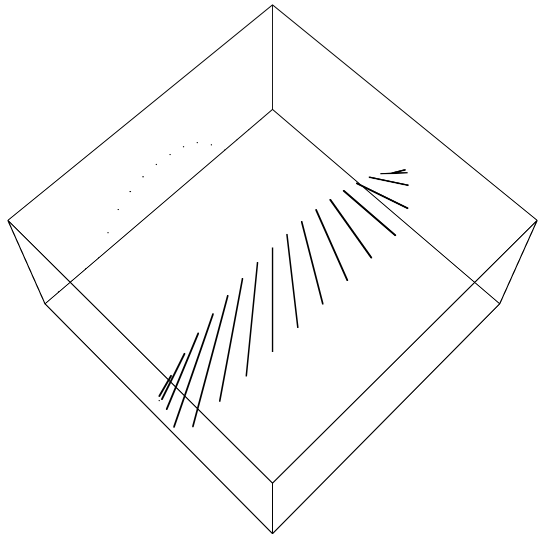

Try the stereo package: ShowStereo.m or ShowSereo2.m

```
<<ShowStereo2.m
```

-- ShowStereo Function Loaded --

```
p1=Plot3D[Cos[Sqrt[x^2+y^2]], {x,-10,10}, {y, -10, 10},
    AspectRatio -> Automatic, PlotPoints-> {25,25}];
```

```
ShowStereo[p1,DisplayFunction->$DisplayFunction];
```

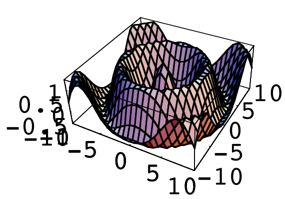

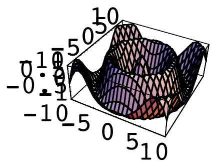

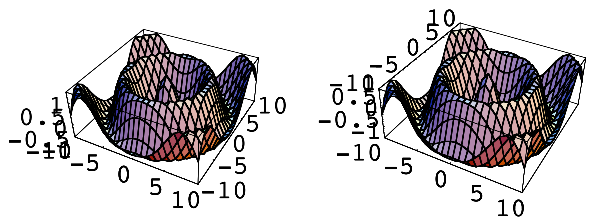

Options[ShowStereo]={AngularSeperation→0.06`,StereoWindowSize→{1265,470},StereoWindowPlacement→{{7,512},{7,-18}},SaveStereoPair→False}

References

Belhumeur, P. N., Kriegman, D. J., & Yuille, A. (1997). *The Bas-Relief Ambiguity*. Paper presented at the IEEE Conf. on Computer Vision and Pattern Recognition.

Brady, M., & Yuille, A. (1983). An Extremum Principle for Shape from Contour. *Proceedings IJCAI*, 969-972.

Bülthoff, H. H., & Yuille, A. (1990). Shape-from-X: Psychophysics and computation. *Proceedings of the SPIE, 1383*, 165-172.

Dupuis, P., & Oliensis, J. (1994). An optimal control formulation and related numerical methods for a problem in shape reconstruction. *The Annals of Applied Probability, 4*(No, 2), 287-346.

Freeman, W. T. (1994). The generic viewpoint assumption in a framework for visual perception. *Nature, 368*(7 April 1994), 542-545.

Freeman, W. T., Pasztor, E. C., & Carmichael, O. T. (2000). Learning low-level vision. *Intl. J. Computer Vision, 40*(1), 25-47.

Gray, Alfred. Modern Differential Geometry of Curves and Surfaces with Mathematica, Second Edition. (CRC Press, 1997) (http://store.wolfram.com/view/ISBN0849371643/?39FED89E-38F0)

Horn, B., & Brooks, M. J. (1989). *Shape from shading*. Cambridge, Mass.: MIT Press.

Humphrey, G. K., Goodale, M. A., Bowen, C. V., Gati, J. S., Vilis, T., Rutt, B. K., et al. (1997). Differences in perceived shape from shading correlate with activity in early visual areas. *Curr Biol, 7*(2), 144-147.

Humphrey, G. K., Symons, L. A., Herbert, A. M., & Goodale, M. A. (1996). A neurological dissociation between shape from shading and shape from edges. *Behavioural Brain Research, 76*(1-2), 117-125.

Ikeuchi, K., & Horn, B. K. P. (1981). Numerical shape from shading and occluding boundaries. *Art. Intel., 17,*, 141-184.

Kersten, D., O'Toole, A. J., Sereno, M. E., Knill, D. C., & Anderson, J. A. (1987). Associative learning of scene parameters from images. *Appl. Opt., 26*, 4999-5006.

Knill, D. C., & Kersten, D. (1990). Learning a near-optimal estimator for surface shape from shading. *CVGIP, 50*(1), 75-100.

Knill, D. C. (1992). Perception of surface contours and surface shape: from computation to psychophysics. *J Opt Soc Am [A], 9*(9), 1449-1464.

Koenderink, J. J. (1998). Pictorial relief. *Philosophical Transactions of the Royal Society of London Series a-Mathematical Physical and Engineering Sciences, 356*(1740), 1071-1085.

Koenderink, J. J. (1990). *Solid Shape*. Cambridge, MA: MIT Press.

Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. (1992). Surface perception in pictures. *Perception & Psychophysics, 5*, 487-496.

Koenderink, J.-J., Kappers, A., & van Doorn, A. (1992). Local Operations: The Embodiment of Geometry. *Artificial and Biological Vision Systems*.

Kourtzi, Z., & Kanwisher, N. (2000). Cortical regions involved in perceiving object shape. *J Neurosci, 20*(9), 3310-3318.

Langer, M. S., & Zucker, S. W. (1994). Shape from Shading on a Cloudy Day. *Journal of the Optical Society of America*

A, 11(2), 467-478.

Langer M.S. and Buelthoff H. H. (2000), Depth discrimination from shading under diffuse lighting. *Perception. 29 (6) 649-660.*

Langer, M. S., & Bulthoff, H. H. (2001). A prior for global convexity in local shape-from-shading. *Perception, 30*(4), 403-410.

Mamassian, P., Kersten, D., & Knill, D. C. (1996). Categorical local-shape perception. *Perception, 25*(1), 95-107.

Mamassian P & Kersten D. (1996). Illumination, shading and the perception of local orientation. *Vision Research., 36*(15), 2351-2367.

Mingolla, E., & Todd, J. T. (1986). Perception of Solid Shape from Shading. *Biological Cybernetics, 53*, 137-151.

Moore, C., & Engel, S. A. (2001). Neural response to perception of volume in the lateral occipital complex. *Neuron, 29*(1), 277-286.

Norman, J. F., Todd, J. T., & Orban, G. A. (2004). Perception of three-dimensional shape from specular highlights, deformations of shading, and other types of visual information. *Psychol Sci, 15*(8), 565-570.

Oliensis, J. (1991). Uniqueness in Shape From Shading. *The International Journal of Computer Vision, 6*(2), 75-104.

Pentland, A. P. (1990). Linear shape from shading. *International Journal of Computer Vision, 4*, 153-162.

Pentland, A. P. (1988). *Shape Information from Shading: A Theory About Human Perception.*

TR-103. Paper presented at the M.I.T. Media Lab Vision Sciences.

Reichel, F. D., Todd, J. T., & Yilmaz, E. (1995). Visual-Discrimination of Local Surface Depth and Orientation. *Perception & Psychophysics, 57*(8), 1233-1240.

Sun, J., & Perona, P. (1998). Where is the sun? *Nature Neuroscience, 1*(3), 183-184.

Todd, J. T., & Mingolla, E. (1983). Perception of Surface Curvature and Direction of Illumination from Patterns of Shading. *Journal of Experimental Psychology: Human Perception & Performance, 9*(4), 583-595.

Todd, J. T., Norman, J. F., Koenderink, J. J., & Kappers, A. M. L. (1997). Effects of texture, illumination, and surface reflectance on stereoscopic shape perception. *Perception, 26*(7), 807-822.

Yuille, A. L. (1987). *Shape from Shading, Occlusion and Texture*. (No. A.I. Memo No. 885): M.I.T.

Yuille, A. L., & Bülthoff, H. H. (1996). Bayesian decision theory and psychophysics. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian Inference* (pp. 123-161). Cambridge, U.K.: Cambridge University Press.