

Computational Vision
U. Minn. Psy 5036
Daniel Kersten
Lecture 26: Spatial Layout, Structure from Motion

Initialize

Outline

Last time

Object recognition continued

Today

Computational theory for estimating relative depth, camera motion

Challenges to computational theories of depth and spatial layout

Spatial layout: Where are objects? Where is the viewer?

Recall distinctions: Between vs. within object geometry.

Where are objects?

■ Absolute

Distance of objects or scene feature points from the observer.

"Physiological cues": Binocular convergence--information about the distance between the eyes and the angle converged by the eyes. Crude, but constraining. Errors might be expected to be proportional to reciprocal distance. Closely related to accommodative requirements.

"Pictorial cue"--familiar size

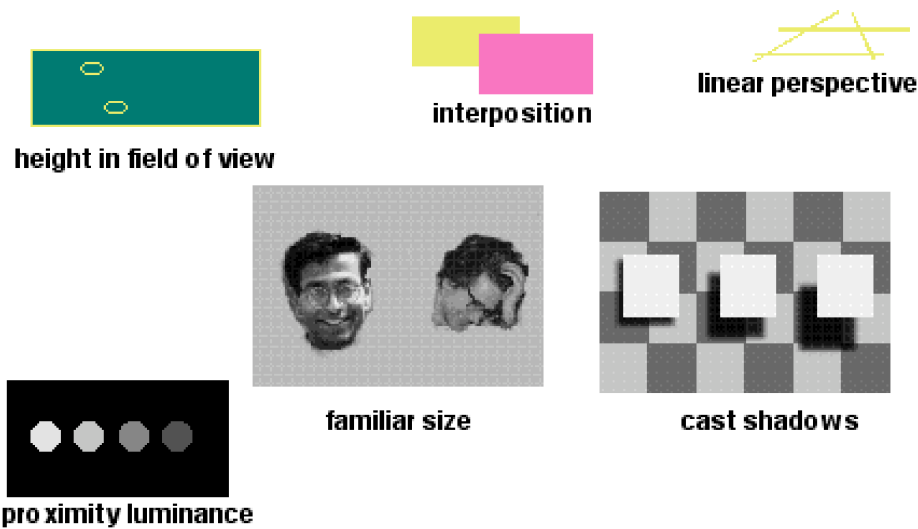
■ Relative

Distance between objects or object feature points. Important for scene layout.

Processes include: Stereopsis (binocular parallax) and motion parallax.

Also information having to do with the "pictorial" cues: occlusion, transparency, perspective, proximity luminance, focus blur, also familiar size & "assumed common physical size", "height in picture plane", cast shadows, texture & texture gradients for large-scale depth & depth gradients

■ Examples of pictorial information for depth



■ Cooperative computation & cue integration

...over a dozen cues to depth. Theories of integration (e.g. stereo + cast shadows). Theories of cooperativity (e.g. motion parallax \Leftrightarrow transparency).

Vision for spatial layout of objects, navigation, heading and for reach

Where is the viewer? And where is the viewer headed?

Computing scene structure from motion information provides information for vision. Can't say where the viewer is in absolute terms, but can say something about the relative depth relationships between objects, and can say something about heading direction, and time to contact.

Calculating structure from motion and heading from the motion field

Estimation of relative depth and eye (or camera) motion

Introduction

Earlier we saw:

- 1) how local motion measurements constrain estimates of optic flow, and thus the motion field.
- 2) how a priori slowness and smoothness constraints constrain dense and sparse estimates of the flow field (e.g. Weiss et al.).

How can we use an estimate of the motion field to estimate useful information for navigation--such as relative depth, observer motion, and time to collision??

Goals

Estimate relative depth, and eye's motion from motion field, estimates of time-to-contact

Ultimately we would like to gain some understanding of the environment from the moving images on our retinas. There are approaches to structure from motion that are not based directly on the motion field, but rather based on a sequence of images in which a discrete set of corresponding points have been identified (Ullman, S., 1979; Dickmanns).


Alternatively, suppose we have estimated the optic flow, and assume it is a good estimate of the motion field--what can we do with it? Imagine the observer is flying through the environment. The flow field should be rich with information regarding direction of heading, time-to-contact, and relative depth (Gibson, 1957).

In this section we study the computational theory for the estimation of relative depth, and camera or eye-point heading from the optic flow pattern induced by general eye motion in a rigid environment. We follow a development described by Longuet-Higgins, H. C., & Prazdny, K. (1980). (See also Koenderink and van Doorn, 1976, Horn, Chapter 17, Perrone, 1992 for a biologically motivated model, and Heeger and Jepson, 1990).

Rather than following the derivation of Longuet-Higgins et al., we derive the relationship between the motion field and relative depth, and camera motion parameters using homogeneous coordinates.

Setting up the frame of reference and basic variables

Imagine a rigid coordinate system attached to the eye, with the origin at the nodal point. General motion of the eye can be described by the instantaneous translational (U,V,W) and rotational (A,B,C) components of the frame. Let P be a fixed point in the world at (X,Y,Z) that projects to point (x,y) in the conjugate image plane which is unit distance in the z direction from the origin:



Goal 1: Derive generative model of the motion field, where we express the motion field (u,v) in terms of Z, U,V,W,A,B,C .

- Express velocity V of world point $P, (X,Y,Z)$ in terms of motion parameters of the camera frame of reference

Let $\mathbf{r}(t)$ represent the position of P in homogeneous coordinates:

$$\mathbf{r}(t) = (X, Y, Z, 1)$$

An instant later, the new coordinates are given by:

$$\mathbf{r}(t + \Delta t) = \mathbf{r} + \Delta \mathbf{r} = (X + \Delta X, Y + \Delta Y, Z + \Delta Z, 1) = \mathbf{r} \Delta R_{\phi_x} \Delta R_{\phi_y} \Delta R_{\phi_z} \Delta T$$

where infinitesimal rotations and translations are represented by their respective 4x4 matrices. (Note that matrix operations do not in general commute). Then,

$$\Delta T = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -\Delta x_0 & -\Delta y_0 & -\Delta z_0 & 1 \end{bmatrix}$$

and

$$\Delta R_{\theta_x} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_x) & -\sin(\theta_x) & 0 \\ 0 & \sin(\theta_x) & \cos(\theta_x) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & -\Delta\theta_x & 0 \\ 0 & \Delta\theta_x & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} +$$

Using similar approximations for the other rotation matrices, and the relation

$$\Delta \mathbf{r} = \mathbf{r} \Delta R_{\theta_x} \Delta R_{\theta_y} \Delta R_{\theta_z} \Delta T - \mathbf{r} I$$

we will use *Mathematica* to show that

$$\{\Delta X, \Delta Y, \Delta Z, 0\} = \{X, Y, Z, 1\} \cdot \begin{pmatrix} 0 & -\Delta\theta_z & \Delta\theta_y & 0 \\ \Delta\theta_z & 0 & -\Delta\theta_x & 0 \\ -\Delta\theta_y & \Delta\theta_x & 0 & 0 \\ -\Delta x & -\Delta y & -\Delta z & 0 \end{pmatrix} + \text{higher order terms}$$

and derive expressions that describe the 3D velocity of P in camera coordinates.

- Let's use *Mathematica* to derive this formula and expressions for the velocity of P, i.e. $V = (dX/dt, dY/dt, \text{ and } dZ/dt)$

```
XRotationMatrix[theta_] :=
  {{1, 0, 0, 0}, {0, Cos[theta], -Sin[theta], 0},
   {0, Sin[theta], Cos[theta], 0}, {0, 0, 0, 1}};
YRotationMatrix[theta_] :=
  {{Cos[theta], 0, Sin[theta], 0}, {0, 1, 0, 0},
   {-Sin[theta], 0, Cos[theta], 0}, {0, 0, 0, 1}};
ZRotationMatrix[theta_] :=
  {{Cos[theta], -Sin[theta], 0, 0}, {Sin[theta], Cos[theta], 0, 0},
   {0, 0, 1, 0}, {0, 0, 0, 1}};
TranslateMatrix[x_, y_, z_] :=
  {{1, 0, 0, 0}, {0, 1, 0, 0}, {0, 0, 1, 0}, {x, y, z, 1}};
```

```
XRotationMatrix[theta]//MatrixForm
```

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & -\sin(\theta) & 0 \\ 0 & \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

```
YRotationMatrix[ $\theta$ ] // MatrixForm
```

$$\begin{pmatrix} \cos(\theta) & 0 & -\sin(\theta) & 0 \\ 0 & 1 & 0 & 0 \\ \sin(\theta) & 0 & \cos(\theta) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

```
ZRotationMatrix[ $\theta$ ] // MatrixForm
```

$$\begin{pmatrix} \cos(\theta) & \sin(\theta) & 0 & 0 \\ -\sin(\theta) & \cos(\theta) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Recall the **Series**[] function:

```
??Series
```

Series[f , { x , x_0 , n }] generates a power series expansion for f about the point $x = x_0$ to order $(x - x_0)^n$.
Series[f , { x , x_0 , n_x }, { y , y_0 , n_y }] successively finds series expansions with respect to x , then y . >>

```
Attributes[Series] = {Protected}
```

```
Options[Series] = {Analytic → True, Assumptions →  $\$Assumptions$ }
```

Expand the rotation matrix into a Taylor series:

```
In[69]:= Series[XRotationMatrix[Subscript[ $\Delta\theta$ , x]},{Subscript[ $\Delta\theta$ , x],0,1}]//MatrixForm
```

```
Out[69]//MatrixForm=
```

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 + O(\Delta\theta_x^2) & -\Delta\theta_x + O(\Delta\theta_x^2) & 0 \\ 0 & \Delta\theta_x + O(\Delta\theta_x^2) & 1 + O(\Delta\theta_x^2) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Use **Normal**[] to chop off higher order terms:

```
In[71]:= Normal[Series[XRotationMatrix[Subscript[Δθ, x]],{Subscript[Δθ, x],0,1}]]//M
```

```
Out[71]//MatrixForm=
```

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & -\Delta\theta_x & 0 \\ 0 & \Delta\theta_x & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Translational matrix is:

```
In[70]:= TranslateMatrix[-Δx,-Δy,-Δz]//MatrixForm
```

```
Out[70]//MatrixForm=
```

$$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -\Delta x & -\Delta y & -\Delta z & 1 \end{pmatrix}$$

Now let's put all the rotation and translational components together.

```
Normal[Series[TranslateMatrix[-Δx,-Δy,-Δz].XRotationMatrix[Subscript[Δθ, x]
```

$$\begin{pmatrix} 0 & -\Delta\theta_z & \Delta\theta_y & 0 \\ \Delta\theta_x \Delta\theta_y + \Delta\theta_z & -\Delta\theta_x \Delta\theta_y \Delta\theta_z & -\Delta\theta_x \Delta\theta_z & -\Delta\theta_x \Delta\theta_y \Delta\theta_z \\ \Delta\theta_x \Delta\theta_z - \Delta\theta_y & \Delta\theta_x + \Delta\theta_y \Delta\theta_z & 0 & 0 \\ -\Delta x + (\Delta z - \Delta y \Delta\theta_x) \Delta\theta_y + (-\Delta y - \Delta z \Delta\theta_x) \Delta\theta_z & -\Delta y - \Delta z \Delta\theta_x + (\Delta x + (\Delta y \Delta\theta_x - \Delta z) \Delta\theta_y) \Delta\theta_z & -\Delta z + \Delta y \Delta\theta_x & 0 \end{pmatrix}$$

Set all second-order terms to zero:

$$\begin{pmatrix} 0 & -\theta_z & \theta_y & 0 \\ \theta_z & 0 & -\theta_x & 0 \\ -\theta_y & \theta_x & 0 & 0 \\ -\Delta x & -\Delta y & -\Delta z & 0 \end{pmatrix}$$

$$\{X, Y, Z, 1\} \cdot \begin{pmatrix} 0 & -\Delta\theta_z & \Delta\theta_y & 0 \\ \Delta\theta_z & 0 & -\Delta\theta_x & 0 \\ -\Delta\theta_y & \Delta\theta_x & 0 & 0 \\ -\Delta x & -\Delta y & -\Delta z & 0 \end{pmatrix}$$

$$\{-\Delta x - Z \Delta\theta_y + Y \Delta\theta_z, -\Delta y + Z \Delta\theta_x - X \Delta\theta_z, -\Delta z - Y \Delta\theta_x + X \Delta\theta_y, 0\}$$

$$\{\Delta X, \Delta Y, \Delta Z, 0\} \sim \{-\Delta x - Z \Delta\theta_y + Y \Delta\theta_z, -\Delta y + Z \Delta\theta_x - X \Delta\theta_z, -\Delta z - Y \Delta\theta_x + X \Delta\theta_y, 0\}$$

Dividing by Δt , we obtain the following relations:

$$\frac{\Delta X}{\Delta t} = \frac{\Delta \theta_z}{\Delta t} Y - \frac{\Delta \theta_y}{\Delta t} X - \frac{\Delta x}{\Delta t} = CY - BX - U$$

$$\frac{\Delta Y}{\Delta t} = -\frac{\Delta y}{\Delta t} - \frac{\Delta \theta_z}{\Delta t} X + \frac{\Delta \theta_x}{\Delta t} Z = -V - CX + AZ$$

$$\frac{\Delta Z}{\Delta t} = -\frac{\Delta z}{\Delta t} - \frac{\Delta \theta_x}{\Delta t} Y + \frac{\Delta \theta_z}{\Delta t} X = -W - AY + BX$$

This result can be neatly summarized using the vector cross product:

$$\mathbf{t} = \{U, V, W\}; \quad \boldsymbol{\omega} = \{A, B, C\}; \quad \mathbf{r} = \{X, Y, Z\};$$


$$- \text{Cross}[\boldsymbol{\omega}, \mathbf{r}] - \mathbf{t}$$

$$\{-U + CY - BZ, -V - CX + AZ, -W + BX - AY\}$$

(NOTE: Here the lower case $\Delta x, \Delta y, \Delta z$ represent changes in the 3D world coordinates $\{X, Y, Z\}$ due to the small translation, and should have 0 subscripts (as above) to distinguish them from the x and y used below. Below we use $\{x, y\}$ to represent the *projection* of $\{X, Y, Z\}$. I hope to fix this notation in the future.)

So far so good. We have described the velocity of P in world coordinates in terms of the rotational and translational velocity components of the moving coordinate system. What is happening in the image--i.e. to the motion field or optic flow?

■ **Next step: Relate P velocity and depth Z to the motion field, $\{u, z\}$**



(Another note: light travels in straight lines, and so should the projected line in the figure above!)

For convenience, we assume the focal length to be one. The perspective projection is:

$$(x, y) = \left(\frac{X}{Z}, \frac{Y}{Z} \right)$$

and the motion field in terms of Z, and the rates of change of X,Y, and Z are:

$$u = \frac{dx}{dt} = \dot{x} = \frac{\dot{X}}{Z} - \frac{X\dot{Z}}{Z^2}$$

$$v = \frac{dy}{dt} = \dot{y} = \frac{\dot{Y}}{Z} - \frac{Y\dot{Z}}{Z^2}$$

To simplify notation, we've used the "dot" convention to indicate the temporal derivatives of X,Y, and Z.

Note that this is the same result we obtained in Lecture 18:

$$\mathbf{v}_i = \frac{(\mathbf{r}_o \times \mathbf{v}_o) \times \mathbf{z}}{(\mathbf{r}_o \cdot \mathbf{z})^2}$$

```
Cross[Cross[{x, y, z}, {ẋ, ẏ, ż}], {0, 0, 1}] / (({x, y, z} . {0, 0, 1}) ^ 2
Drop[%, -1] // Expand
```

```
{(z ẋ - x ż) / z^2, (z ẏ - y ż) / z^2, 0}
```

```
{(ẋ / z - x ż / z^2), (ẏ / z - y ż / z^2)}
```

■ Main result for goal 1, the generative model:

Substituting the expressions for the rate of change of X,Y, and Z, we have:

$$u = \left(\frac{-U + Wx}{Z} \right) + (-B + Cy + Axy - Bx^2)$$

$$v = \left(\frac{-V + Wy}{Z} \right) + (-Cx + A + Ay^2 - Bxy)$$

Note that we have organized the terms on the right of each equation so that the first parts do not depend on A,B, or C--that is

$$\mathbf{u}_T = \left(\frac{-U + Wx}{Z} \right)$$


is a purely translational component, and the second term in brackets is purely rotational, and further does not depend on Z:

$$\mathbf{u}_R = (-B + Cy + Axy - Bx^2)$$

So in general, we can write:

$$\mathbf{u} = \mathbf{u}_T + \mathbf{u}_R \quad \mathbf{v} = \mathbf{v}_T + \mathbf{v}_R$$

The figure on the left below shows the flow field one would expect from a purely translational motion--there is a center of expansion (which could be off a finite image plane). The right panel shows the flow pattern of a rotational field.



The above equations can be summarized in the following matrix equation:

$$\mathbf{u}(x,y) = p(x,y)\mathbf{A}(x,y) \cdot \mathbf{t} + \mathbf{B}(x,y) \cdot \omega$$

where $p(x,y) = 1/Z(x,y)$.

We can verify that the matrix equation gives the solution we derived:

```
Clear[A, B, CC, U, V, W, X, Y, Z];
AA[x_, y_] := {{-f, 0, x}, {0, -f, y}}
BB[x_, y_] := {{(x*y)/f, -(f+x^2)/f, y}, {f+y^2/f, -(x*y)/f, -x}}
uu[x_, y_] := (1/Z)*AA[x, y].{U, V, W} + BB[x, y].{A, B, CC}
```

```
f = 1;
uu[x, y]
```

$$\left\{ \frac{Wx - U}{Z} + B(-x^2 - 1) + CCy + Axy, -CCx - Byx + A(y^2 + 1) + \frac{Wy - V}{Z} \right\}$$

Let's plot the motion field for a planar surface. The equation for a plane in camera coordinates can be written:

$$Z = aX + bY + Z_0$$

Substituting from $x = fX/Z$, $y = fY/Z$,

we get

$$Z(x,y) = \frac{f Z_0}{(f - ax - by)}$$

We can use this to define a function $p(x,y) = 1/Z$:

$$p[x_, y_] := (1/Z0) - (a/(f*Z0))*x - (b/(f*Z0)) y$$

```

Clear[A, B, CC, U, V, W, X, Y, Z, uu, p];

p[x_, y_] := (1/Z0) - (a/(f*Z0))*x - (b/(f*Z0)) y
a = 0; b = 1.0; f = 1.0; Z0 = -1;

AA[x_, y_] := {{-f, 0, x}, {0, -f, y}}
BB[x_, y_] := {{(x*y)/f, -(f+x^2)/f, y}, {f+y^2/f, -(x*y)/f, -x}}
uu2[x_, y_] := p[x, y] * AA[x, y] . {U, V, W} + BB[x, y] . {A, B, CC}

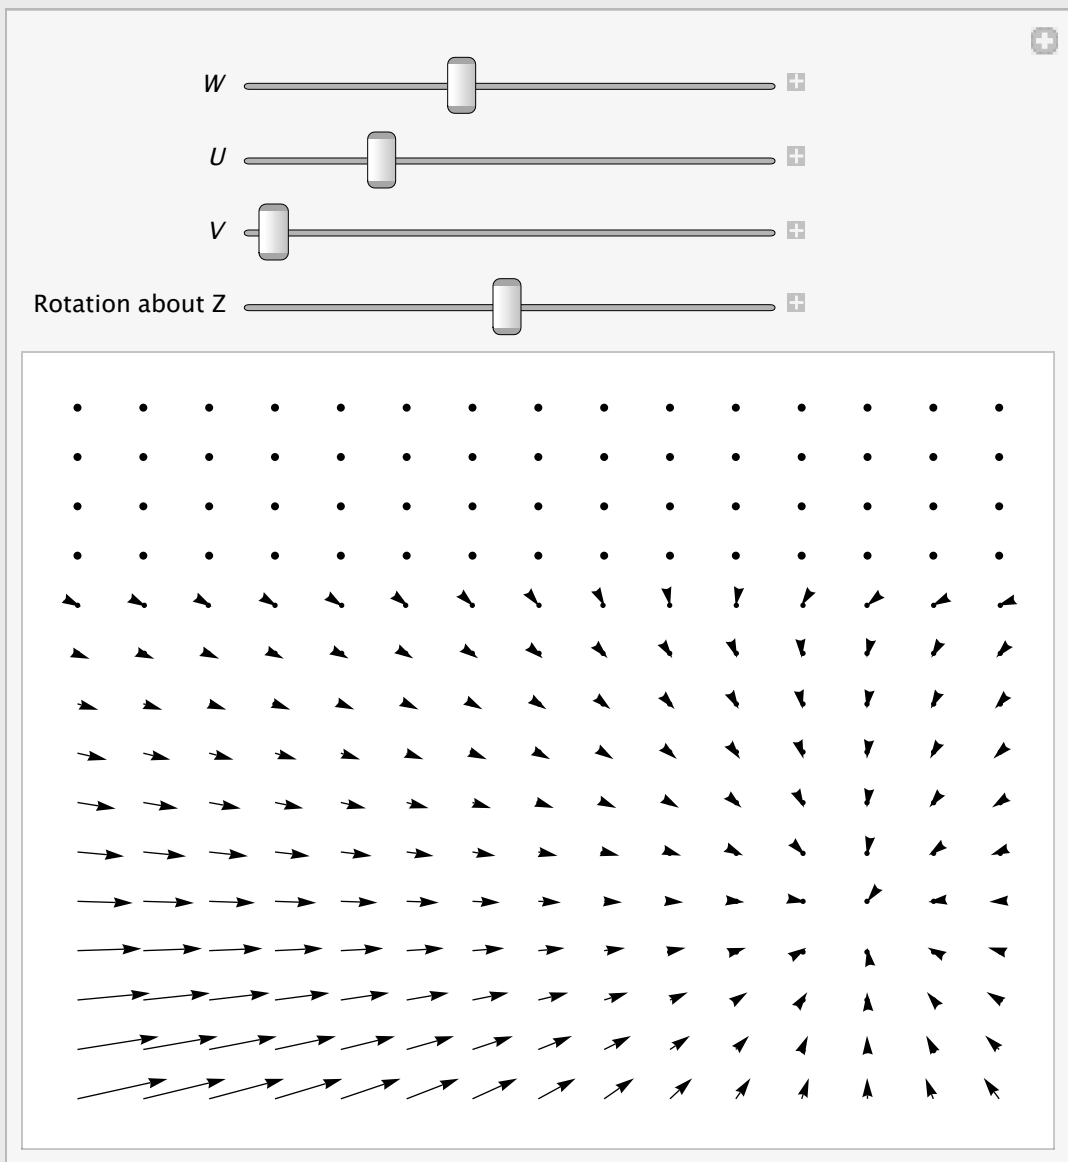
```

```

A = 0; B = 0; CC = -2;
U = 0; V = 0; W = 0;
Manipulate[
  VectorFieldPlot[If[y < 1, p[x, y] * AA[x, y] . {U, V, W} + BB[x, y] . {A, B, CC},
    {0, 0}], {x, -10, 10}, {y, -10, 5}, ScaleFactor -> 2],
  {{W, 20}, 0, 50}, {U, 0, 500}, {V, -100, 50},
  {{A, 0.0, "Rotation about X axis"}, -300, 300}]

```

Out[68]=



With the rotation about the X axis set to zero ($A=0$), observe how changing the camera translational vector elements (U,V,W) changes the field singularity corresponding to the focus of expansion. The focus of expansion indicates the direction of heading.

Now while leaving the (U,V,W) fixed, introduce a non-zero rotation (e.g. $A > 0$). Notice how the focus of expansion moves. The focus no longer corresponds to the direction of heading.

Goal 2: Inference model: given (u,v), how can we obtain estimates of A,B,C,U,V,W,Z?

In general we can't obtain all seven unknowns (see Horn's book, chap. 17). One problem is that scaling Z by a constant, can be exactly compensated for by a reciprocal scaling of (U,V,W) yielding an unchanged motion field.

In other words, a given motion flow on your retina could be due to fast movement through a bigger world, or slower movement through a smaller world.

Horn discusses least squares solutions for the direction of camera motion, and for its rotational component. See also Heeger and Jepson (1990).

Although one can use Bayesian methods for the insufficiently constrained problem of estimating any or all of the seven unknowns, let's see how far one can get with simple algebra to get movement direction (not speed) and relative depth (not absolute depth). We follow the original work of Longuet-Higgins et al. (1980) for estimating the camera direction, and relative depth.

■ Pure translation: Obtaining direction of heading and relative depth

Suppose the rotational component is zero. Then measurements of the motion field will give us the translational components. These components constrain U,V,W, and Z at each point (x,y) in the conjugate image plane.


$$u^{\tau} = \frac{-U + xW}{Z}$$

$$v^{\tau} = \frac{-V + yW}{Z}$$

Combining these two equations to eliminate Z:

$$\left(y - \frac{V}{W}\right) = \left(x - \frac{U}{W}\right) \frac{v^{\tau}}{u^{\tau}}$$

This equation is a straight line whose slope is determined by the ratio of the vertical and horizontal components of the flow field (v^{τ}/u^{τ}), and which passes through the point (V/W, U/W). This point depends only on the camera's translational velocity, so other motion flow field lines with different ratios of vertical and horizontal components (v^{τ^2}/u^{τ^2}) of the flow also pass through this point. The point (V/W, U/W) is the *focus of expansion*.



Two motion field lines determine the focus of expansion, and thus the camera's translational direction, whose cosine is:

$$\frac{(U, V, W)}{\sqrt{U^2 + V^2 + W^2}} = \frac{(U/W, V/W, 1)}{\sqrt{U^2/W^2 + V^2/W^2 + 1}}$$

We can also obtain an estimate of the relative depth of points:


$$\frac{Z}{W} = \frac{x - U/W}{u^T} = \frac{y - V/W}{v^T}$$

■ Pure rotation

As we noted above, when the camera has zero translational velocity, the motion field does not depend on the depth structure of the scene. The field is a quadratic function of image position.

■ General motion field: Estimate both rotation and translational components using points of occlusion to get Z, the depth structure

What if we have a mix of translation and rotation? One solution suggested by Longuet-Higgins and Prazdny is to make use of *motion parallax*, where we have two 3D points that project to the same conjugate image point.




In general, these two points will have different motion field vectors at this image point. If we take the difference, we have:

$$u_1 - u_2 = (-U + xW)\left(\frac{1}{Z_1} - \frac{1}{Z_2}\right)$$

$$v_1 - v_2 = (-V + yW)\left(\frac{1}{Z_1} - \frac{1}{Z_2}\right)$$

$$\left(y - \frac{V}{W}\right) = \left(\frac{v_1 - v_2}{u_1 - u_2}\right)\left(x - \frac{U}{W}\right)$$

Again, finding the focus of expansion (x_0, y_0) , which involves finding at least two motion parallax pairs,



will give us the camera (or eye) direction


$$x_0 = \frac{U}{W}$$

$$y_0 = \frac{V}{W}$$

To find relative depth, we need to know A,B,C:

$$u(x_0, y) = (C + Ax_0)y - B(1 + x_0^2)$$

$$v(x, y_0) = -(C + By_0)x + A(1 + y_0^2)$$



These relations provide sufficient information to calculate A,B,C (from two or more points). A,B,C in turn determine u^R and v^R .

$$u = (x - x_0) \frac{W}{Z} + u^R; \quad v = (y - y_0) \frac{W}{Z} + v^R$$

With some rearrangement, we can obtain a formula for relative depth:

$$u - u^R = (x - x_0) \frac{W}{Z}$$

$$v - v^R = (y - y_0) \frac{W}{Z}$$

$$\frac{Z}{W} = \frac{x - x_0}{u - u^R} = \frac{y - y_0}{v - v^R}$$

Although we won't take the time to go over the results, a potentially important form of information for relative depth, camera motion, and time-to-contact comes from an analysis of the flow patterns generated by textured surfaces (Koen-

derink, J. J., & van Doorn, A. J., 1976) and the above cited article by Longuet-Higgins and Prazdny. The idea is to compute estimates of the rotation, dilation, and shear of the motion field.

Exercise: Time-to-contact

Problem: Show that the reciprocal of the temporal rate of expansion of an object heading directly towards you is equal to the time to contact. (Lee and Reddish, 1981).

Heading experiments

Structure from motion: Psychophysics

Warren and Hannon (1988) provided the first compelling evidence that the human visual system could compensate for eye rotation purely from optical information.

Royden, Banks & Crowell (1992) later pointed out the role of proprioceptive information in heading computation, especially for faster motions.

Structure from motion: Physiology

A possible neurophysiological basis for derivative measurements of flow (e.g. rotation, dilation, shear), see: (Saito, H.-A., Yukie, M., Tanaka, K., Hikosaka, K., Fukada, Y., & Iwai, E., 1986). For work relating to eye movement compensation in optic flow and heading, See Bradley et al. (1996). See Duffy (2000) for recent work.

Challenges to computing structure from motion

Multiple motions, transparency

In general, the environment has multiple moving objects--imagine driving a car. It can involve complex surface relationships, such as near bushes in front of a distant set of houses.

One approach to these problems is to assume the flow field arises from a discrete set of layers, and estimate flow parameters within each.

There has been considerable work in computer vision to solve the problem of tracking, e.g. multiple pedestrians.

We don't yet have a complete picture of how the human visual system copes with multiple motions when determining scene structure, direction of heading, or time to contact.

■ Empirical work

For more information, see work by Royden, Banks, Warren, Crowell and others.

There has been considerable interest in the neural basis of optic flow computations and their functions. For a review see: Wurtz, R. H. (1998). Optic flow: A brain region devoted to optic flow analysis? *Curr Biol*, 8(16), R554-556.

For ideas on neural coding, see: Zemel, R. S., & Sejnowski, T. J. (1998). A model for encoding multiple object motions and self-motion in area MST of primate visual cortex. *J Neurosci*, 18(1), 531-547.

Depth between objects

- **Depth from shadows** (<http://gandalf.psych.umn.edu/users/kersten/kersten-lab/demos/160x120nt.mov>)

Depth from viewer

- **Cue integration: Shadow displacement & size change for depth**

Frame of reference issues in cue integration.

Modularity for cue integration: Shadow displacement & size change for depth

Frame of reference issues in cue integration (Schrater, & Kersten, 2000).

Earlier we looked at a simple graph for cue integration and showed how an optimal estimate (for the Gaussian case), say for depth, was a weighted combination of the estimates for the individual cues. The weights were determined from the relative reliabilities of the cues.

But a close examination of the generative models that result in multiple cues can show a more complex set of dependencies.


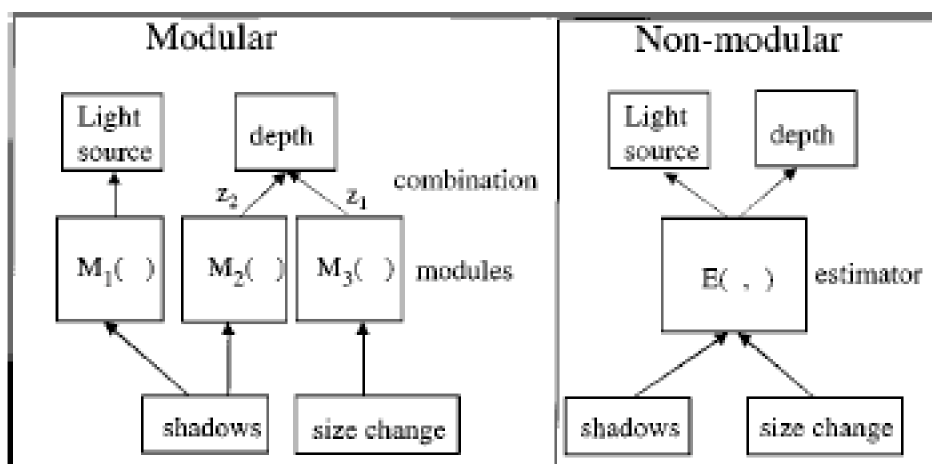



Figure 3. Whether independent data measures are singly connected to the estimated variable S_x determines whether or not estimation modules can be created for S_x . *Left* example of Bayesian modularity. Boxes show how the variables can be split to form two modules. *Right* example of a non-modular estimation.

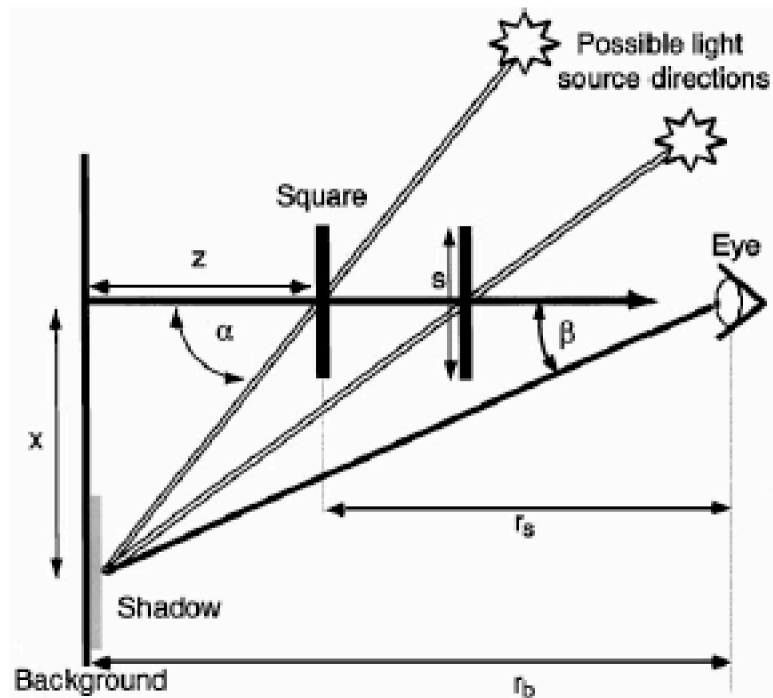
This has an impact on the architecture for optimal inverse inference--whether the algorithm can be broken into distinct modules or not. The non-modular case below is an example of what Clark and Yuille called "strong fusion". This is related to the notion of "cooperative computation" discussed earlier in Lecture 23 on perceptual integration.



Let's take a look at a specific case involving size and shadow position as cues for an object's 3D position in space.



The figure below shows some of the relationships between the data (shadow position β , size of the target square is a --not shown), and unknown parameters to be estimated (z , r_s) of interest, (the unit-less parameter, z/r_b is not shown), and unknowns to be integrated out (α , s , r_b --depending on the task).



When we perceive a change in depth, *what variable does perceived depth correspond to?* Here are three possibilities: relative (unit-less) distance z/r_b , depth from the observer, r_s , and distance from the background z .

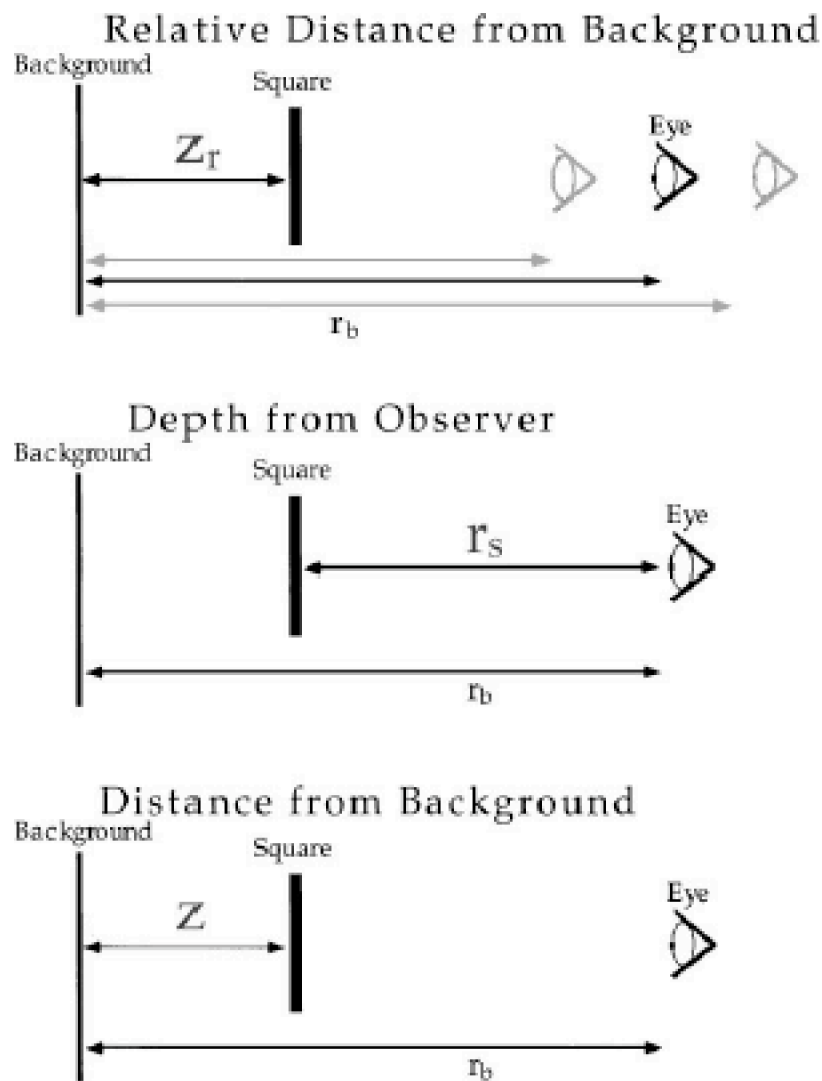


Figure 6. Diagram illustrating the depth variables to be estimated. The variable $z_r = z/r_b$ can't be shown directly, because it is an equivalence class of z and r_b distances.

Paul Schrater worked through the math and showed that these different assumptions about depth representation produced different generative models for producing the image size \mathbf{a} , and shadow position, β .




Figure 7. Bayes nets for the three depth representations. a) Bayes net for relative distance to the background. This task involves estimating object relations (world centered), and requires the least prior knowledge. b) Bayes net for distance to observer. Notice that the use of the shadow information requires integrating across two variables, hence the shadow cue should have more uncertainty for this task. c) Bayes net for metric depth from the background. Estimating the distance from the background, z , is complicated by the image size and shadow position measurements also being jointly dependent on the observer’s distance to the background.

Fisher information is the asymptotic variance of the estimator, so can be used to calculate a weighted linear combination (an optimal estimator for the modular case).

Table 1. Table of MAP estimates and Fisher information values for the three depth estimations. For the representations which admit modular estimates, the estimates are shown sep the shadow and image size cues.

Task	Est from shadow	Est from size	Shadow Fisher info	Size Fi
Relative z	$z_r = \tan(\hat{\beta})$	$z_r = 1 - \frac{\mu_{z_r}}{\hat{a}}$	$\frac{1}{\sqrt{2} \tan(\hat{\beta})^2}$	$\frac{2i}{\mu_{z_r}^2 (\sigma_{z_r}^2)}$
Dist. from obs.	$r_s = \mu_{r_s} (1 - \tan(\hat{\beta}))$	$r_s = \frac{\mu_s}{\hat{a}}$	$\frac{1}{\mu_{r_s}^2 \tan(\hat{\beta})^2}$	$\frac{2i}{\mu_s^2 (\sigma_s^2)}$
Absolute z	$z = \frac{\mu_s \tan(\hat{\beta})}{\hat{a} (1 - \tan(\hat{\beta}))}$		$\frac{2\hat{a}^2 (1 - \tan(\hat{\beta}))^4}{\mu_s^2 \tan(\hat{\beta})^2}$	

The shadow cue is most reliable when the target object is close to the background. But the size cue is most reliable when the target is close to the viewer.

There have been no systematic experimental studies of this general theoretical prediction.

■ Bottom line

Optimal estimators for depth depend critically on the representation of depth

Different representations result in different generative models, and thus different modular structures for optimal inference

Human judgments of closeness may be better predicted by a model that represents depth from the observer, rather than relative depth from the background, in either absolute (e.g. metric) units, or relative units. More experimental work is needed.

Appendix

Initialization cell for homogeneous coordinate functions.

```
In[1]:= BeginPackage ["Geometry`Homogeneous`"]
XRotationMatrix::"usage" =
  "XRotationMatrix[phi] gives the matrix for rotation about
  x-axis by phi degrees in radians"
YRotationMatrix::"usage" =
  "YRotationMatrix[phi] gives the matrix for rotation about
  y-axis by phi degrees in radians"
ZRotationMatrix::"usage" =
  "ZRotationMatrix[phi] gives the matrix for rotation about
  z-axis by phi degrees in radians"
ScaleMatrix::"usage" =
  "ScaleMatrix[sx,sy,sz] gives the matrix to scale a vector by
  sx,sy, and sz in the x, y and z directions, respectively."
TranslateMatrix::"usage" =
  "TranslateMatrix[x,y,z] gives the matrix to translate coordinates
  by x,y,z."
ThreeDToHomogeneous::"usage" =
  "ThreeDToHomogeneous[sx,sy,sz] converts 3D coordinates to 4D
  homogeneous coordinates."
HomogeneousToThreeD::"usage" =
  "HomogeneousToThreeD[4Dvector] converts 4D homogeneous
  coordinates to 3D coordinates."
ZProjectMatrix::"usage" =
  "ZProjectMatrix[focal] gives the 4x4 projection matrix to map
  a vector through the origin to an image plane at focal
  distance from the origin along the z-axis."
ZOrthographic::"usage" =
  "ZOrthographic[vector] projects vector on to the x-y plane."
Begin["`private`"]
XRotationMatrix[theta_] :=
```

```

    {{1, 0, 0, 0}, {0, Cos[theta], -Sin[theta], 0},
     {0, Sin[theta], Cos[theta], 0}, {0, 0, 0, 1}};
YRotationMatrix[theta_] :=
    {{Cos[theta], 0, Sin[theta], 0}, {0, 1, 0, 0},
     {-Sin[theta], 0, Cos[theta], 0}, {0, 0, 0, 1}};
ZRotationMatrix[theta_] :=
    {{Cos[theta], -Sin[theta], 0, 0}, {Sin[theta], Cos[theta], 0, 0},
     {0, 0, 1, 0}, {0, 0, 0, 1}};
ScaleMatrix[sx_, sy_, sz_] :=
    {{sx, 0, 0, 0}, {0, sy, 0, 0}, {0, 0, sz, 0}, {0, 0, 0, 1}};
(*TranslateMatrix[x_, y_, z_] :=
    {{1, 0, 0, x}, {0, 1, 0, y}, {0, 0, 1, z}, {0, 0, 0, 1}};*)
TranslateMatrix[x_, y_, z_] :=
    {{1, 0, 0, 0}, {0, 1, 0, 0}, {0, 0, 1, 0}, {x, y, z, 1}};
ThreeDToHomogeneous[vec_] := Append[vec, 1];
HomogeneousToThreeD[vec_] := Drop[ $\frac{\text{vec}}{\text{vec}[[4]}}$ , -1];
ZProjectMatrix[focal_] :=
    {{1, 0, 0, 0}, {0, 1, 0, 0}, {0, 0, 1, 0}, {0, 0, N[ $\frac{1}{\text{focal}}$ ], 0}};
ZOrthographic[vec_] := Take[vec, 2];
End[]
EndPackage[]

```

Out[1]= Geometry`Homogeneous`

Out[2]= XRotationMatrix[phi] gives the matrix for rotation about x-axis by phi degrees in radians

Out[3]= YRotationMatrix[phi] gives the matrix for rotation about y-axis by phi degrees in radians

Out[4]= ZRotationMatrix[phi] gives the matrix for rotation about z-axis by phi degrees in radians

Out[5]= ScaleMatrix[sx,sy,sz] gives the matrix to scale
a vector by sx,sy, and sz in the x, y and z directions, respectively.

Out[6]= TranslateMatrix[x,y,z] gives the matrix to translate coordinates by x,y,z.

Out[7]= ThreeDToHomogeneous[sx,sy,sz] converts 3D coordinates to 4D homogeneous coordinates.

Out[8]= HomogeneousToThreeD[4Dvector] converts 4D homogeneous coordinates to 3D coordinates.

Out[9]= ZProjectMatrix[focal] gives the 4x4 projection matrix to map a vector through the origin to an image plane at focal distance from the origin along the z-axis.

Out[10]= ZOrthographic[vector] projects vector on to the x-y plane.

Out[11]= Geometry`Homogeneous`private`

Out[21]= Geometry`Homogeneous`private`

References

- Bradley, D. C., Maxwell, M., Andersen, R. A., Banks, M. S., & Shenoy, K. V. (1996). Mechanisms of heading perception in primate visual cortex [see comments]. *Science*, 273(5281), 1544-7.
- Crowell, J. A., & Banks, M. S. (1996). Ideal observer for heading judgments. *Vision Res*, 36(3), 471-90.
- d'Avossa, G., & Kersten, D. (1996). Evidence in human subjects for independent coding of azimuth and elevation for direction of heading from optic flow. *Vision Research*, 36(18), 2915-2924.
- Duffy, C. J. (2000). Optic flow analysis for self-movement perception. *Int Rev Neurobiol*, 44, 199-218.
- Gibson. (1957). Optical motions and transformations as stimuli for visual perception. *Psychological Review*, 64, 288-295.
- Heeger, D. J., & Jepson, A. (1990). Visual perception of three-dimensional motion. *Neural Computation*, 2(2), 129-137.
- Jepson, A., & Black, M. J. (1993). *Mixture models for optical flow computation*. Paper presented at the Proc. IEEE Conf. Comput. Vision Pattern Recog., New York.
- (link)
- Koenderink, J. J., & van Doorn, A. J. (1976). Local Structure of Movement Parallax of the Plane. *J. Opt. Soc. Am.*, 66, 717-723.
- Lappe, M., & Rauschecker, J. P. (1993). A neural network for the processing of optic flow from ego-motion in man and higher mammals. *Neural Computation*, 5, 374-391.
- Lee, D. N., & Reddish, P. E. (1981). Plummeting gannets: a paradigm of ecological optics. *Nature*, 293(5830), 293-294.
- Longuet-Higgins, H. C., & Prazdny, K. (1980). The Interpretation of a Moving Retinal Image. *Proceedings of the Royal Society of London B*, 208, 385-397.
- Martin, W. N., & Aggarwal, J. K. (1988). *Motion understanding: robot and human vision*. Boston: Kluwer Academic Publishers,
- Perrone, J. A. (1992). Model for the computation of self-motion in biological systems. *Journal of the Optical Society of America*, A, 9(2), 177-194.
- Royden, C. S., Banks, M. S., & Crowell, J. A. (1992). The perception of heading during eye movements [see comments]. *Nature*, 360(6404), 583-5.
- Saito, H.-A., Yukie, M., Tanaka, K., Hikosaka, K., Fukada, Y., & Iwai, E. (1986). Integration of Direction Signals of Image Motion in the Superior Temporal Sulcus of the Macaque Monkey. *The Journal of Neuroscience*, 6(1), 145-157.
- Schrater PR, Kersten D (2000) How optimal depth cue integration depends on the task. *International Journal of Computer Vision* 40:73-91.
- Ullman, S. (1979). The Interpretation of Structure from Motion. *Proceedings of the Royal Society London B*, 203, 405-426.)

Wang, J. Y. A., & Adelson, E. H. (1994). Representing moving images with layers. *IEEE Transactions on Image Processing Special Issue: Image Sequence Compression*, 3(5), 625-638.

Warren, W. H., & Hannon, D. J. (1988). Direction of Self-Motion is Perceived from Optical Flow. *Nature*, 336((10)), 162-163.

Zemel, R. S., & Sejnowski, T. J. (1998). A model for encoding multiple object motions and self-motion in area MST of primate visual cortex. *J Neurosci*, 18(1), 531-547.