Announcements

- Questions on the project?
- New turn-in info online
- Demos this Friday, 12-1:30

Motion Estimation

http://www.sandlotscience.com/Distortions/Breathing_objects.htm

http://www.sandlotscience.com/Ambiguous/barberpole.htm

Today's Readings

- Watt, 10.3-10.4 (handout)
- Trucco & Verri, 8.3 – 8.4 (skip 8.3.3, read only top half of p. 199) (handout)

Why estimate motion?

Lots of uses

- Track object behavior
- Correct for camera jitter (stabilization)
- Align images (mosaics)
- 3D shape reconstruction
- Special effects

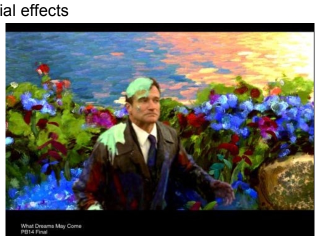

Optical flow

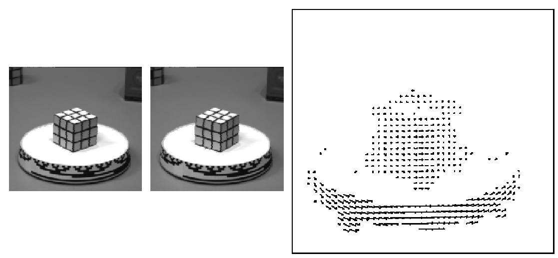

Problem definition: optical flow

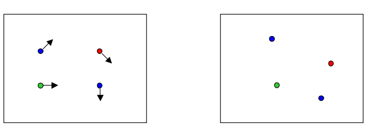

How to estimate pixel motion from image H to image I?

- Solve pixel correspondence problem
  - given a pixel in H, look for nearby pixels of the same color in I

Key assumptions

- **color constancy:** a point in H looks the same in I
  - For grayscale images, this is **brightness constancy**
- **small motion**: points do not move very far

This is called the **optical flow** problem

Optical flow constraints (grayscale images)

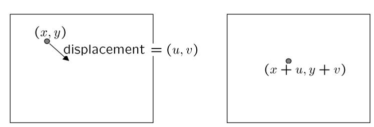

Let's look at these constraints more closely

- brightness constancy:

$$H(x, y) = I(x + u, y + v)$$

- small motion: (u and v are less than 1 pixel)
  - suppose we take the Taylor series expansion of I:

$$I(x+u, y+v) = I(x, y) + \frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v + \text{higher order terms}$$

$$\approx I(x, y) + \frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v$$

Optical flow equation

Combining these two equations

shorthand: $I_x = \frac{\partial I}{\partial x}$

$$
\begin{aligned}
0 &= I(x+u, y+v) - H(x,y) \\
&\approx I(x,y) + I_x u + I_y v - H(x,y) \\
&\approx (I(x,y) - H(x,y)) + I_x u + I_y v \\
&\approx I_t + I_x u + I_y v \\
&\approx I_t + \nabla I \cdot [u \ v]
\end{aligned}
$$

In the limit as u and v go to zero, this becomes exact

$$0 = I_t + \nabla I \cdot [\frac{\partial x}{\partial t} \ \frac{\partial y}{\partial t}]$$

Optical flow equation

$$0 = I_t + \nabla I \cdot [u \ v]$$

Q: how many unknowns and equations per pixel?

- A: u and v are unknown, 1 equation

Intuitively, what does this constraint mean?

- The component of the flow in the gradient direction is determined
- The component of the flow parallel to an edge is unknown

This explains the Barber Pole illusion
http://www.sandlotscience.com/Ambiguous/barberpole.htm

Aperture problem

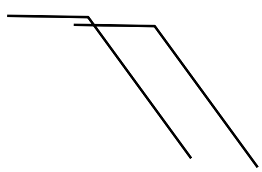

Aperture problem

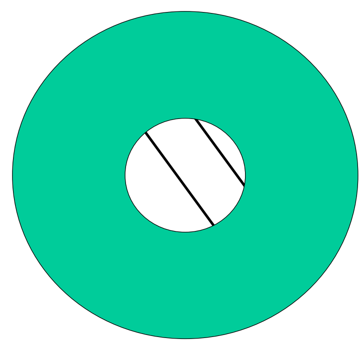

Solving the aperture problem

How to get more equations for a pixel?

- Basic idea: impose additional constraints
  - most common is to assume that the flow field is smooth locally
  - one method: pretend the pixel's neighbors have the same (u,v)
    - » If we use a 5x5 window, that gives us 25 equations per pixel!

$$0 = I_t(\mathbf{p_i}) + \nabla I(\mathbf{p_i}) \cdot [u \ v]$$

$$\underset{25\times 2}{\underbrace{\begin{bmatrix} I_x(\mathbf{p_1}) & I_y(\mathbf{p_1}) \\ I_x(\mathbf{p_2}) & I_y(\mathbf{p_2}) \\ \vdots & \vdots \\ I_x(\mathbf{p_{25}}) & I_y(\mathbf{p_{25}}) \end{bmatrix}}_{A}} \underset{2\times 1}{\underbrace{\begin{bmatrix} u \\ v \end{bmatrix}}_{d}} = - \underset{25\times 1}{\underbrace{\begin{bmatrix} I_t(\mathbf{p_1}) \\ I_t(\mathbf{p_2}) \\ \vdots \\ I_t(\mathbf{p_{25}}) \end{bmatrix}}_{b}}$$

Lukas-Kanade flow

Prob: we have more equations than unknowns

$$\underset{25\times 2}{A} \ \underset{2\times 1}{d} = \underset{25\times 1}{b} \longrightarrow \text{minimize } \|Ad - b\|^2$$

Solution: solve least squares problem

- minimum least squares solution given by solution (in d) of:

$$\underset{2\times 2}{(A^T A)} \ \underset{2\times 1}{d} = \underset{2\times 1}{A^T b}$$

$$\underbrace{\begin{bmatrix} \sum I_x I_x & \sum I_x I_y \\ \sum I_x I_y & \sum I_y I_y \end{bmatrix}}_{A^T A} \begin{bmatrix} u \\ v \end{bmatrix} = - \underbrace{\begin{bmatrix} \sum I_x I_t \\ \sum I_y I_t \end{bmatrix}}_{A^T b}$$

- The summations are over all pixels in the K x K window
- This technique was first proposed by Lukas & Kanade (1981)
  - described in Trucco & Verri handout

Conditions for solvability

- Optimal (u, v) satisfies Lucas-Kanade equation

$$\underbrace{\begin{bmatrix} \sum I_x I_x & \sum I_x I_y \\ \sum I_x I_y & \sum I_y I_y \end{bmatrix}}_{A^T A} \begin{bmatrix} u \\ v \end{bmatrix} = - \underbrace{\begin{bmatrix} \sum I_x I_t \\ \sum I_y I_t \end{bmatrix}}_{A^T b}$$

When is This Solvable?

- **A^TA** should be invertible
- **A^TA** should not be too small due to noise
  - eigenvalues λ_1 and λ_2 of **A^TA** should not be too small
- **A^TA** should be well-conditioned
  - λ_1 / λ_2 should not be too large (λ_1 = larger eigenvalue)

Eigenvectors of A^TA

$$A^T A = \begin{bmatrix} \sum I_x I_x & \sum I_x I_y \\ \sum I_x I_y & \sum I_y I_y \end{bmatrix} = \sum \begin{bmatrix} I_x \\ I_y \end{bmatrix} [I_x \ I_y] = \sum \nabla I (\nabla I)^T$$

Suppose (x,y) is on an edge. What is A^TA? derive on board

- gradients along edge all point the same direction
- gradients away from edge have small magnitude

$$\left(\sum \nabla I (\nabla I)^T\right) \approx k \nabla I \nabla I^T$$

$$\left(\sum \nabla I (\nabla I)^T\right) \nabla I = k \|\nabla I\| \nabla I$$

- ∇I is an eigenvector with eigenvalue $k\|\nabla I\|$
- What's the other eigenvector of A^TA?
  - let N be perpendicular to ∇I

$$\left(\sum \nabla I (\nabla I)^T\right) N = 0$$

  - N is the second eigenvector with eigenvalue 0

The eigenvectors of A^TA relate to edge direction and magnitude

Edge

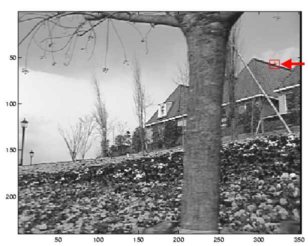
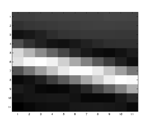
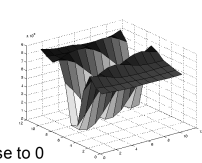

$\sum \nabla I (\nabla I)^T$

- gradients are all the same or close to 0
- large λ_1, small λ_2

Low texture region

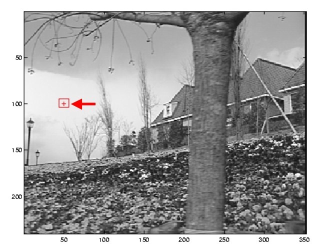
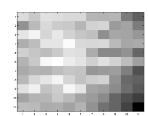
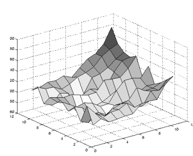

$\sum \nabla I (\nabla I)^T$

- gradients have small magnitude
- small λ_1, small λ_2

High textured region


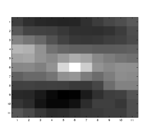
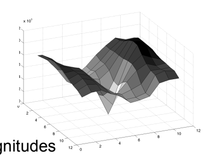

$\sum \nabla I (\nabla I)^T$

- gradients are different, large magnitudes
- large λ_1, large λ_2

Observation

This is a two image problem BUT

- Can measure sensitivity by just looking at one of the images!
- This tells us which pixels are easy to track, which are hard
  - very useful later on when we do feature tracking...

Errors in Lukas-Kanade

What are the potential causes of errors in this procedure?

- Suppose A^TA is easily invertible
- Suppose there is not much noise in the image

When our assumptions are violated

- Brightness constancy is **not** satisfied
- The motion is **not** small
- A point does **not** move like its neighbors
  - window size is too large
  - what is the ideal window size?

Improving accuracy

Recall our small motion assumption

$$0 = I(x+u, y+v) - H(x,y)$$

$$\approx I(x,y) + I_x u + I_y v - H(x,y)$$

This is not exact

- To do better, we need to add higher order terms back in:

$$= I(x,y) + I_x u + I_y v + \text{higher order terms} - H(x,y)$$

This is a polynomial root finding problem

- Can solve using **Newton's method** (1D case on board)
  - Also known as **Newton-Raphson** method
  - Today's reading (first four pages)
    - » http://www.ulib.org/webRoot/Books/Numerical_Recipes/bookcpdf/c9-4.pdf
- Lukas-Kanade method does one iteration of Newton's method
  - Better results are obtained via more iterations

Iterative Refinement

Iterative Lukas-Kanade Algorithm

1. Estimate velocity at each pixel by solving Lucas-Kanade equations
2. Warp H towards I using the estimated flow field
   - *use image warping techniques*
3. Repeat until convergence

Revisiting the small motion assumption

Is this motion small enough?

- Probably not—it's much larger than one pixel
- How might we solve this problem?

Reduce the resolution!

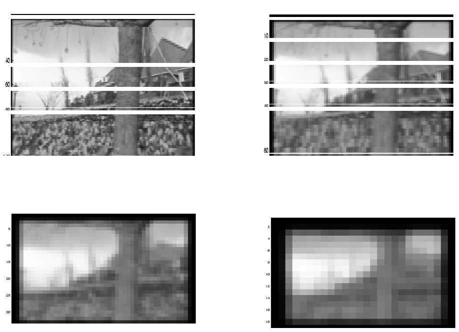

Coarse-to-fine optical flow estimation

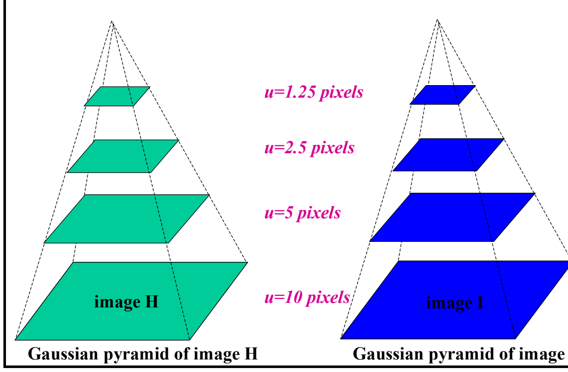

Coarse-to-fine optical flow estimation

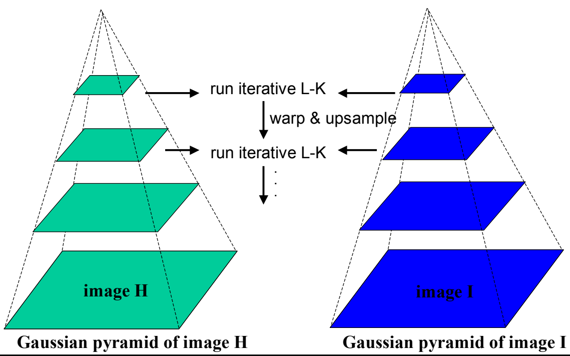

Optical flow result

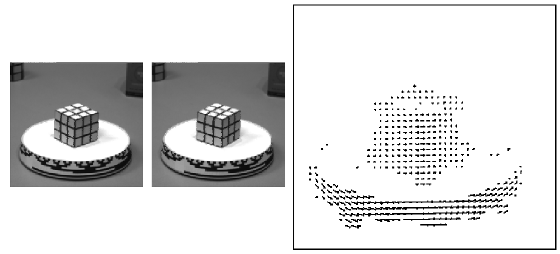

Optical flow demo

- http://extra.cmis.csiro.au/cgi-bin/motionfast.cgi

Motion tracking

Suppose we have more than two images

- How to track a point through all of the images?
  - In principle, we could estimate motion between each pair of consecutive frames
  - Given point in first frame, follow arrows to trace out it's path
  - Problem: DRIFT
    - » small errors will tend to grow and grow over time—the point will drift way off course

Feature Tracking

- Choose only the points ("features") that are easily tracked
- How to find these features?
  - windows where $\sum \nabla I (\nabla I)^T$ has two large eigenvalues
- Called the Harris Corner Detector

Feature Detection

Tracking features

Feature tracking

- Compute optical flow for that feature for each consecutive H, I

Complications:

- Occlusions—feature may disappear
  - need mechanism for deleting, adding new features
- Changes in shape, orientation
  - allow the feature to deform
- Changes in color
- Large motions
  - will pyramid techniques work for feature tracking?

Handling large motions

L-K requires small motion

- If the motion is much more than a pixel, use discrete **search** instead

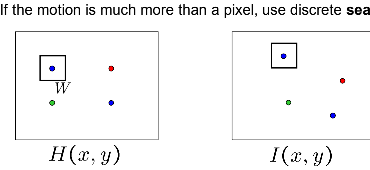

- Given window W in H, find best matching window in I
- Minimize sum squared difference (SSD) of pixels in window

$$min_{(u,v)} \left\{ \sum_{(x,y)\in W} |I(x+u, y+v) - H(x,y)|^2 \right\}$$

- Solve by doing a search over a specified range of (u,v) values
  - this (u,v) range defines the **search window**

Tracking Over Many Frames

Feature tracking with m frames

1. Select features in first frame
2. Given feature in frame i, compute position in i+1
3. Select more features if needed
4. i = i + 1
5. If i < m, go to step 2

Issues

- Discrete search vs. Lucas Kanade?
  - depends on expected magnitude of motion
  - discrete search is more flexible
- How often to update feature template?
  - update often enough to compensate for distortion
  - updating too often causes drift
- How big should search window be?
  - too small: lost features. Too large: slow

Incorporating Dynamics

Idea

- Can get better performance if we know something about the way points move
- Most approaches assume constant velocity

$$\dot{\mathbf{x}}_{i+1} = \dot{\mathbf{x}}_i$$
$$\mathbf{x}_{i+1} = 2\mathbf{x}_i - \mathbf{x}_{i-1}$$

or constant acceleration

$$\ddot{\mathbf{x}}_{i+1} = \ddot{\mathbf{x}}_i$$
$$\mathbf{x}_{i+1} = 3\mathbf{x}_i - 3\mathbf{x}_{i-1} + \mathbf{x}_{i-2}$$

- Use above to predict position in next frame, initialize search

Feature tracking demo

Oxford video

http://www.toulouse.ca/?/CamTracker/?/CamTracker/FeatureTracking.html

MPEG—application of feature tracking

- http://www.pixeltools.com/pixweb2.html

Image alignment

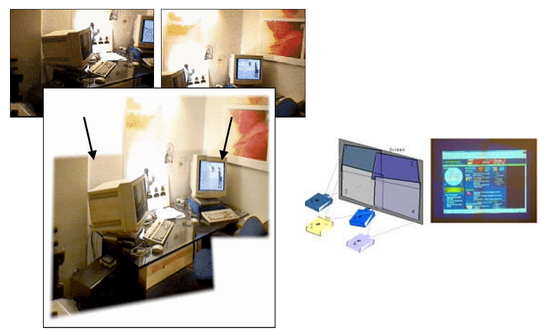

Goal: estimate single (u,v) translation for entire image

- Easier subcase: solvable by pyramid-based Lukas-Kanade

Summary

Things to take away from this lecture

- Optical flow problem definition
- Aperture problem and how it arises
- Assumptions
  - Brightness constancy, small motion, smoothness
- Derivation of optical flow constraint equation
- Lukas-Kanade equation
  - Derivation
  - Conditions for solvability
  - meanings of eigenvalues and eigenvectors
- Iterative refinement
  - Newton's method
  - Coarse-to-fine flow estimation
- Feature tracking
  - Harris feature detector
  - L-K vs. discrete search method
  - Tracking over many frames
  - Prediction using dynamics
- Applications
  - MPEG video compression
  - Image alignment