

Announcements

- Project1 due Tuesday

Motion Estimation

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

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

Today's Readings

- Trucco & Verri, 8.3 – 8.4 (skip 8.3.3, read only top half of p. 199)
- *Supplemental:*
 - R. Bergen, P. Anandan, K.J. Hanna, and R. Hingorani, [Hierarchical model-based motion estimation](#), European Conf. on Computer Vision (ECCV), 1992
 - http://www.cs.washington.edu/education/courses/576/03sp/readings/bergen_eccv92.pdf
- Numerical Recipes (Newton-Raphson), 9.4 (first four pages)
 - http://www.ulib.org/webRoot/Books/Numerical_Recipes/bookpdf/c9-4.pdf



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: Q: what's the equation?

- 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

$$0 = I(x+u, y+v) - H(x, y) \quad \text{shorthand: } I_x = \frac{\partial I}{\partial x}$$

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

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

$$\approx I_t + I_x u + I_y v$$

$$\approx I_t + \nabla I \cdot [u \ v]$$

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

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

Optical flow equation

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

Q: how many unknowns and equations per pixel?


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

Basic idea: assume motion field is smooth

Horn & Schunk: add smoothness term

$$\iint (I_t + \nabla I \cdot [u \ v])^2 + \lambda^2 (\|\nabla u\|^2 + \|\nabla v\|^2) \, dx \, dy$$

Lukas & Kanade: assume locally constant motion

- pretend the pixel's neighbors have the same (u,v)
 - If we use a 5x5 window, that gives us 25 equations per pixel!

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

- works better in practice than Horn & Schunk

Many other methods exist. Here's an overview:

- Barron, J.L., Fleet, D.J., and Beauchemin, S. Performance of optical flow techniques, *International Journal of Computer Vision*, 12(1):43-77, 1994.
- <http://www.cs.washington.edu/education/courses/576/03sp/readings/barron92performance.pdf>

Lukas-Kanade flow

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!

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

$$\begin{bmatrix} I_x(p_1) & I_y(p_1) \\ I_x(p_2) & I_y(p_2) \\ \vdots & \vdots \\ I_x(p_{25}) & I_y(p_{25}) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = - \begin{bmatrix} I_t(p_1) \\ I_t(p_2) \\ \vdots \\ I_t(p_{25}) \end{bmatrix}$$

$\begin{matrix} A & d & b \\ 25 \times 2 & 2 \times 1 & 25 \times 1 \end{matrix}$

RGB version

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*3 equations per pixel!

$$0 = I_t(p_i)[0, 1, 2] + \nabla I(p_i)[0, 1, 2] \cdot [u \ v]$$

$$\begin{bmatrix} I_x(p_1)[0] & I_y(p_1)[0] \\ I_x(p_1)[1] & I_y(p_1)[1] \\ I_x(p_1)[2] & I_y(p_1)[2] \\ \vdots & \vdots \\ I_x(p_{25})[0] & I_y(p_{25})[0] \\ I_x(p_{25})[1] & I_y(p_{25})[1] \\ I_x(p_{25})[2] & I_y(p_{25})[2] \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = - \begin{bmatrix} I_t(p_1)[0] \\ I_t(p_1)[1] \\ I_t(p_1)[2] \\ \vdots \\ I_t(p_{25})[0] \\ I_t(p_{25})[1] \\ I_t(p_{25})[2] \end{bmatrix}$$

$$\begin{matrix} A & d & b \\ 75 \times 2 & 2 \times 1 & 75 \times 1 \end{matrix}$$

Lukas-Kanade flow

Prob: we have more equations than unknowns

$$\begin{matrix} A & d = t \\ 25 \times 2 & 2 \times 1 & 25 \times 1 \end{matrix} \longrightarrow \text{minimize } \|Ad - b\|^2$$

Solution: solve least squares problem

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

$$(A^T A) d = A^T b$$

$$\begin{matrix} 2 \times 2 & 2 \times 1 & 2 \times 1 \end{matrix}$$

$$\begin{bmatrix} \sum I_x I_x & \sum I_x I_y \\ \sum I_x I_y & \sum I_y I_y \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = - \begin{bmatrix} \sum I_x I_t \\ \sum I_y I_t \end{bmatrix}$$

$$\begin{matrix} A^T A & A^T b \end{matrix}$$

- 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 reading

Conditions for solvability

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

$$\begin{bmatrix} \sum I_x I_x & \sum I_x I_y \\ \sum I_x I_y & \sum I_y I_y \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = - \begin{bmatrix} \sum I_x I_t \\ \sum I_y I_t \end{bmatrix}$$

$$\begin{matrix} A^T A & A^T b \end{matrix}$$

When is This Solvable?

- $A^T A$ should be invertible
- $A^T A$ should not be too small due to noise
 - eigenvalues λ_1 and λ_2 of $A^T A$ should not be too small
- $A^T A$ should be well-conditioned
 - λ_1 / λ_2 should not be too large ($\lambda_1 =$ larger eigenvalue)

Eigenvectors of $A^T A$

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

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

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

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

$$(\sum \nabla I (\nabla I)^T) \nabla I = k \|\nabla I\|^2 \nabla I$$

- ∇I is an eigenvector with eigenvalue $k \|\nabla I\|^2$
- What's the other eigenvector of $A^T A$?

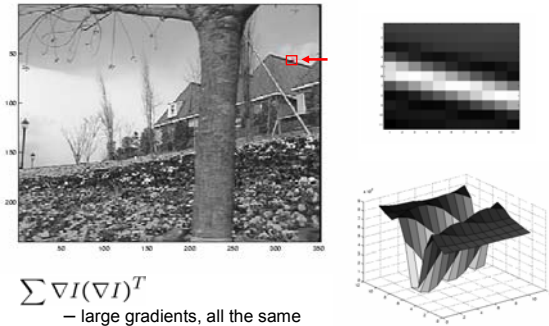
- let N be perpendicular to ∇I

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

- N is the second eigenvector with eigenvalue 0

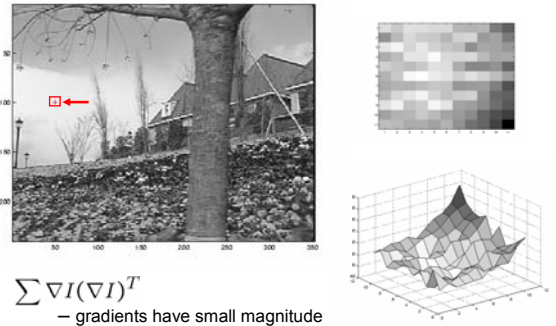
The eigenvectors of $A^T A$ relate to edge direction and magnitude

Edge



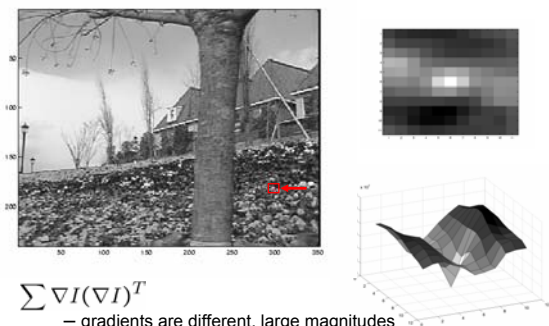
$\sum \nabla I (\nabla I)^T$
– large gradients, all the same
– large λ_1 , small λ_2

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 $A^T A$ 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 (2nd order terms dominate)
- How might we solve this problem?

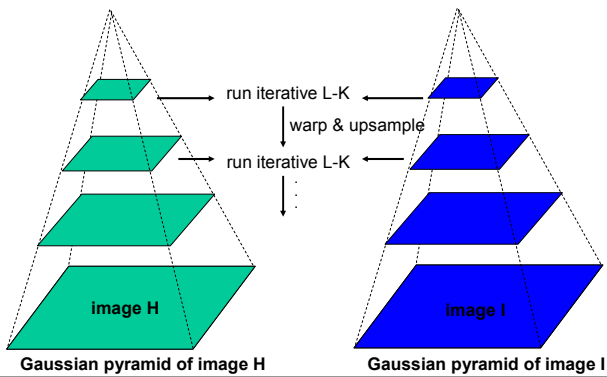
Reduce the resolution!



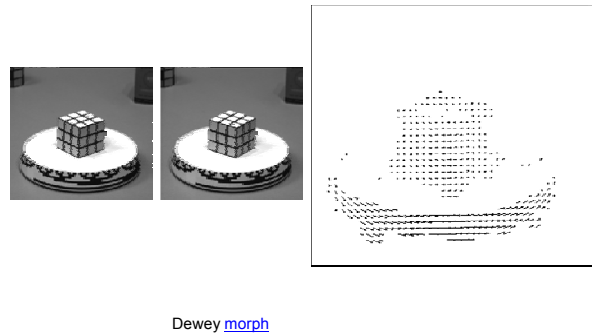
Coarse-to-fine optical flow estimation



Coarse-to-fine optical flow estimation



Optical flow result




Dewey [morph](#)


Robust methods

L-K minimizes a sum-of-squares error metric

- least squares techniques overly sensitive to outliers




Error metrics




quadratic

$$\rho(x) = x^2$$



truncated quadratic

$$\rho_{\alpha,\lambda}(x) = \begin{cases} \lambda x^2 & \text{if } |x| < \frac{\sqrt{\alpha}}{\sqrt{\lambda}} \\ \alpha & \text{otherwise} \end{cases}$$



lorentzian

$$\rho_{\sigma}(x) = \log\left(1 + \frac{1}{2}\left(\frac{x}{\sigma}\right)^2\right)$$


Robust optical flow

Robust Horn & Schunk

$$\iint \rho(I_t + \nabla I \cdot [u \ v]) + \lambda^2 \rho(\|\nabla u\|^2 + \|\nabla v\|^2) \, dx \, dy$$

Robust Lukas & Kanade

$$\sum_{(x,y) \in W} \rho(I_t + \nabla I \cdot [u \ v])$$



first image

quadratic flow

lorentzian flow

detected outliers

Reference

- Black, M. J. and Anandan, P., A framework for the robust estimation of optical flow, *Fourth International Conf. on Computer Vision (ICCV)*, 1993, pp. 231-236
<http://www.cs.washington.edu/education/courses/576/03sp/readings/black93.pdf>

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 its 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


When will this go wrong?

- 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 feature 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
- Compare feature in frame i to i+1 or frame 1 to i+1?
 - affects tendency to 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

$$\begin{aligned} \dot{x}_{i+1} &= \dot{x}_i \\ x_{i+1} &= 2x_i - x_{i-1} \end{aligned}$$

or constant acceleration

$$\begin{aligned} \ddot{x}_{i+1} &= \ddot{x}_i \\ x_{i+1} &= 3x_i - 3x_{i-1} + x_{i-2} \end{aligned}$$

- 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