Feature Tracking and Optical Flow



Computer Vision Jia-Bin Huang, Virginia Tech

Many slides from D. Hoiem

Administrative Stuffs

- HW 1 due 11:59 PM Sept 19
 - Submission through Canvas
- Regular office hour by Jia-Bin
 - 3:00- 4:00 PM, Friday Sept 16
- Bonus office hour by Akrit
 - 10:30 AM 11: 30 AM, Monday Sept 19
- HW 1 Competition: Edge Detection
 - <u>Submission link</u>
 - <u>Leaderboard</u>

Previous class

- Interest point/keypoint/feature detectors
 - Harris: detects corners
 - DoG: detects peaks/troughs
- Interest point/keypoint/feature descriptors
 - SIFT (do read the paper)

Distinctive image features from scale-invariant keypoints

<u>DG Lowe</u> - International journal of computer vision, 2004 - Springer Abstract This paper presents a method for extracting distinctive invariant features from images that can be used to perform reliable matching between different views of an object or scene. The features are invariant to image scale and rotation, and are shown to provide ... Cited by 36572 Related articles All 214 versions Web of Science: 13301 Import into BibTeX

- Feature matching
 - Ratio distance = ||f₁ f₂ || / || f₁ f₂' ||
 - Remove 90% false matches, 5% of true matches in Lowe's study







This class: recovering motion

- Feature tracking
 - Extract visual features (corners, textured areas) and "track" them over multiple frames
- Optical flow
 - Recover image motion at each pixel from spatio-temporal image brightness variations

Two problems, one registration method

B. Lucas and T. Kanade. <u>An iterative image registration technique with an application to</u> <u>stereo vision.</u> In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1981.

Feature tracking

- Many problems, such as structure from motion require matching points
- If motion is small, tracking is an easy way to get them



盟

Feature tracking - Challenges

- Figure out which features can be tracked
- Efficiently track across frames
- Some points may change appearance over time (e.g., due to rotation, moving into shadows, etc.)
- Drift: small errors can accumulate as appearance model is updated
- Points may appear or disappear: need to be able to add/delete tracked points

Feature tracking



- Given two subsequent frames, estimate the point translation
- Key assumptions of Lucas-Kanade Tracker
 - **Brightness constancy:** projection of the same point looks the same in every frame
 - Small motion: points do not move very far
 - **Spatial coherence:** points move like their neighbors

The brightness constancy constraint



• Brightness Constancy Equation:

$$I(x, y, t) = I(x + u, y + v, t + 1)$$

Take Taylor expansion of I(x+u, y+v, t+1) at (x, y, t) to linearize the right side:

Image derivative along x Difference over frames

$$I(x+u, y+v, t+1) \approx I(x, y, t) + I_x \cdot u + I_y \cdot v + I_t$$

$$I(x+u, y+v, t+1) - I(x, y, t) = I_x \cdot u + I_y \cdot v + I_t$$
So:

$$I_x \cdot u + I_y \cdot v + I_t \approx 0 \quad \rightarrow \nabla I \cdot [u \ v]^T + I_t = 0$$

The brightness constancy constraint

Can we use this equation to recover image motion (u,v) at each pixel?

$$\nabla \mathbf{I} \cdot \begin{bmatrix} \mathbf{u} & \mathbf{v} \end{bmatrix}^{\mathrm{T}} + \mathbf{I}_{\mathrm{t}} = \mathbf{0}$$

• How many equations and unknowns per pixel?

•One equation (this is a scalar equation!), two unknowns (u,v)

The component of the motion perpendicular to the gradient (i.e., parallel to the edge) cannot be measured

If (u, v) satisfies the equation, so does (u+u', v+v') if $\nabla I \cdot [u' v']^T = 0$ (u', v')(u+u', v+v)

The aperture problem



The aperture problem





The barber pole illusion



http://en.wikipedia.org/wiki/Barberpole illusion

The barber pole illusion





http://en.wikipedia.org/wiki/Barberpole illusion

Solving the ambiguity...

B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 674–679, 1981.

- How to get more equations for a pixel?
- Spatial coherence constraint
 - Assume the pixel's neighbors have the same (u,v)
 - If we use a 5x5 window, that gives us 25 equations per pixel



Solving the ambiguity...

• Least squares problem:

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

$$A \quad d = b$$

25x2 2x1 25x1

Matching patches across images

• Overconstrained linear system

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

$$25 \times 2 = 2 \times 1 = 25 \times 1$$

Least squares solution for d given by $(A^T A) d = A^T b$

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

The summations are over all pixels in the K x K window

Conditions for solvability

Optimal (u, v) satisfies Lucas-Kanade equation

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

When is this solvable? I.e., what are good points to track?

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

Does this remind you of anything?

Criteria for Harris corner detector

Low-texture region



– small λ_1 , small λ_2

Edge



High-texture region



 $\sum \nabla I (\nabla I)^T$ - gradients are different, large magnitudes
- large λ_1 , large λ_2

The aperture problem resolved



The aperture problem resolved





Dealing with larger movements: Original (x,y) position Iterative refinement

Initialize (x',y') = (x,y)1.

2. Compute (u,v) by

$$I_t = I(x', y', t+1) - I(x, y, t)$$

 $\begin{bmatrix} \sum I_x I_x & \sum I_x I_y \\ \sum I_x I_y & \sum I_y I_y \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = -\begin{bmatrix} \sum I_x I_t \\ \sum I_y I_t \end{bmatrix}$ 2nd moment matrix for feature patch in first image

displacement

- Shift window by (u, v): x' = x' + u; y' = y' + v;3.
- Recalculate I_{t} 4.
- Repeat steps 2-4 until small change 5.
 - Use interpolation for subpixel values



Shi-Tomasi feature tracker

- Find good features using eigenvalues of secondmoment matrix (e.g., Harris detector or threshold on the smallest eigenvalue)
 - Key idea: "good" features to track are the ones whose motion can be estimated reliably
 - Track from frame to frame with Lucas-Kanade
 - This amounts to assuming a translation model for frame-toframe feature movement
- Check consistency of tracks by *affine* registration to the first observed instance of the feature
 - Affine model is more accurate for larger displacements
 - Comparing to the first frame helps to minimize drift

J. Shi and C. Tomasi. <u>Good Features to Track</u>. CVPR 1994.

Tracking example



Figure 1: Three frame details from Woody Allen's *Manhattan*. The details are from the 1st, 11th, and 21st frames of a subsequence from the movie.



Figure 2: The traffic sign windows from frames 1,6,11,16,21 as tracked (top), and warped by the computed deformation matrices (bottom).

J. Shi and C. Tomasi. <u>Good Features to Track</u>. CVPR 1994.

Summary of KLT tracking

- Find a good point to track (harris corner)
- Use intensity second moment matrix and difference across frames to find displacement
- Iterate and use coarse-to-fine search to deal with larger movements
- When creating long tracks, check appearance of registered patch against appearance of initial patch to find points that have drifted

Implementation issues

- Window size
 - Small window more sensitive to noise and may miss larger motions (without pyramid)
 - Large window more likely to cross an occlusion boundary (and it's slower)
 - 15x15 to 31x31 seems typical
- Weighting the window
 - Common to apply weights so that center matters more (e.g., with Gaussian)

Why not just do local template matching?

- Slow (need to check more locations)
- Does not give subpixel alignment (or becomes much slower)
 - Even pixel alignment may not be good enough to prevent drift
- May be useful as a step in tracking if there are large movements

Optical flow



Picture courtesy of Selim Temizer - Learning and Intelligent Systems (LIS) Group, MIT

Motion and perceptual organization

 Even "impoverished" motion data can evoke a strong percept



G. Johansson, "Visual Perception of Biological Motion and a Model For Its Analysis", *Perception and Psychophysics 14, 201-211, 1973.*

Motion and perceptual organization

 Even "impoverished" motion data can evoke a strong percept



G. Johansson, "Visual Perception of Biological Motion and a Model For Its Analysis", *Perception and Psychophysics 14, 201-211, 1973.*

Uses of motion

- Estimating 3D structure
- Segmenting objects based on motion cues
- Learning and tracking dynamical models
- Recognizing events and activities
- Improving video quality (motion stabilization)

Motion field

• The motion field is the projection of the 3D scene motion into the image





What would the motion field of a non-rotating ball moving towards the camera look like?

Optical flow

- Definition: optical flow is the *apparent* motion of brightness patterns in the image
- Ideally, optical flow would be the same as the motion field
- Have to be careful: apparent motion can be caused by lighting changes without any actual motion
 - Think of a uniform rotating sphere under fixed lighting vs. a stationary sphere under moving illumination

Lucas-Kanade Optical Flow

- Same as Lucas-Kanade feature tracking, but for each pixel
 - As we saw, works better for textured pixels
- Operations can be done one frame at a time, rather than pixel by pixel
 - Efficient

Iterative Refinement

- Iterative Lukas-Kanade Algorithm
 - 1. Estimate displacement at each pixel by solving Lucas-Kanade equations
 - 2. Warp I(t) towards I(t+1) using the estimated flow field

- Basically, just interpolation

3. Repeat until convergence



Example



Multi-resolution registration





Optical Flow Results



Optical Flow Results



Errors in Lucas-Kanade

- The motion is large
 - Possible Fix: Keypoint matching
- A point does not move like its neighbors
 - Possible Fix: Region-based matching
- Brightness constancy does not hold
 - Possible Fix: Gradient constancy

State-of-the-art optical flow

- Start with something similar to Lucas-Kanade
- + gradient constancy
- + energy minimization with smoothing term
- + region matching
- + keypoint matching (long-range)





Region-based +Pixel-based +Keypoint-based Large displacement optical flow, Brox et al., CVPR 2009

Things to remember

- Major contributions from Lucas, Tomasi, Kanade
 - Tracking feature points
 - Optical flow
 - Stereo (later)
 - Structure from motion (later)
- Key ideas
 - By assuming brightness constancy, truncated Taylor expansion leads to simple and fast patch matching across frames
 - Coarse-to-fine registration

Next week

• HW 1 due Monday

• Object/image alignment