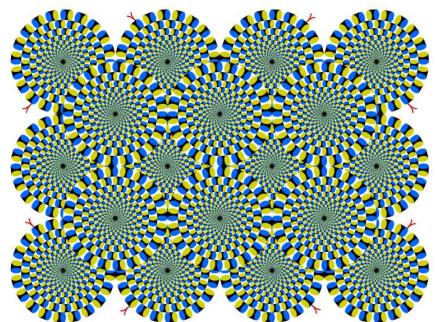


Motion Estimation

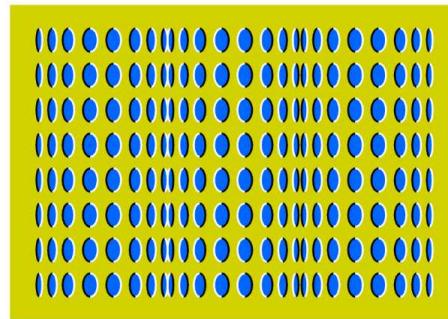
- Lots of uses
 - Track object behavior
 - Correct for camera jitter (stabilization)
 - Align images (mosaics)
 - 3D shape reconstruction
 - Special effects



Motion Illusion created by Akiyoshi Kitaoka



Motion Illusion created by Akiyoshi Kitaoka



ANALYSIS OF VISUAL MOTION

- GOAL: RECOVER RELATIVE MOTION BETWEEN OBSERVER AND ENVIRONMENT. I.E., INFER MOTION OF 3D OBJECTS FROM SEQUENCE OF 2D IMAGES
- GIVEN: $I(x, y, t)$
- MEASUREMENT PROBLEM:
COMPUTE INSTANTANEOUS VELOCITY, (u, v) , OF EVERY POINT AT EVERY TIME. I.E.,
RECOVER $V(x, y, t) = (u(x, y, t), v(x, y, t))$
 $= (\frac{dx}{dt}, \frac{dy}{dt})$
- INTERPRETATION PROBLEM:
USE VELOCITY INFO TO
INFER 3D STRUCTURE
HIGH-LEVEL MOTION DESCRIPTION
SUCH AS BEHAVIOR RECOGNITION

DIMENSIONS OF COMPLEXITY

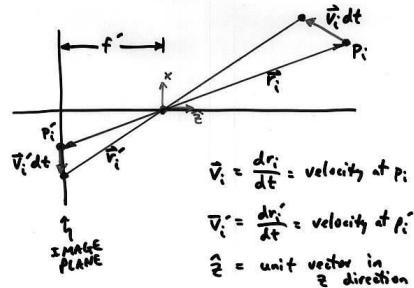
- OBJECT TYPE
 - POINTS
 - RIGID
 - NON-RIGID
- NUMBER OF MOVING OBJECTS
- TYPES OF OBJECT MOTION
 - 2D OR 3D TRANSLATION
 - 2D OR 3D ROTATION
 - ARTICULATED
- TYPES OF CAMERA MOTION
 - STATIONARY
 - TRANSLATION \perp OPTICAL AXIS
 - GENERAL TRANSLATION
 - ROTATION

HUMAN VISUAL SYSTEM

- 2 SEPARATE MOTION PROCESSES LIKELY AT WORK (BRADDICK):
 - SHORT-RANGE
 - * APPROX. CONTINUOUS MOTION
 - * SPATIAL DISPLACEMENTS $< 15^{\circ}$ arc
 - * TEMPORAL DISPLACEMENTS < 100 msec
 - LONG-RANGE
 - * LARGE SPATIAL AND/OR TEMPORAL DISPLACEMENTS
 - * FEATURE DETECTION + CORRESPONDENCE

MOTION FIELD / IMAGE FLOW FIELD

- MOTION FIELD = IMAGE VELOCITY OF A POINT THAT IS MOVING IN SCENE. I.E., 2D VECTOR FIELD THAT IS PERSPECTIVE PROJECTION ON TO IMAGE PLANE OF 3D VELOCITY FIELD OF MOVING SCENE



- BY PERSPECTIVE PROJ. EQ:

$$\begin{aligned}\frac{\vec{r}'_i}{f'} &= -\frac{\vec{r}_i}{\vec{r}_i \cdot \hat{z}} \\ \Rightarrow \vec{r}'_i &= -\frac{f' \vec{r}_i}{\vec{r}_i \cdot \hat{z}}\end{aligned}$$

- DIFFERENTIATING & SIMPLIFYING:

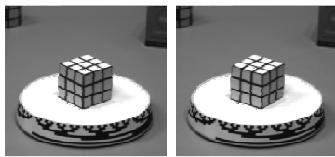
$$\vec{v}'_i = \frac{d\vec{r}'_i}{dt} = -\frac{f'(\vec{r}_i \times \vec{v}_i) \times \hat{z}}{(\vec{r}_i \cdot \hat{z})^2}$$

↑ MOTION FIELD

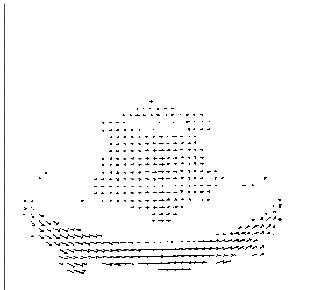
- HOW TO RECOVER \vec{v}' ?

OPTICAL FLOW

- OPTICAL FLOW (AKA IMAGE-INTENSITY FIELD) = MOTION OF THE BRIGHTNESS PATTERN IN IMAGE
- ASSUME MOTION FIELD \approx OPTICAL FLOW
- COUNTER EXAMPLES:
 - * ROTATING LAMBERTIAN SPHERE of CONSTANT ALBEDO, FIXED LIGHT SOURCE AND CAMERA
⇒ IMAGE CONSTANT OVER TIME
⇒ OPTICAL FLOW = 0
BUT MOTION FIELD ≠ 0
 - * FIXED LAMBERTIAN SPHERE of CONSTANT ALBEDO, FIXED CAMERA, MOVING LIGHT
⇒ OPTICAL FLOW ≠ 0
BUT MOTION FIELD = 0



Optical flow



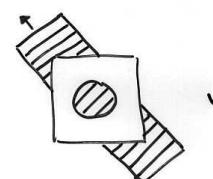
OPTICAL FLOW PROPERTIES

- ASSUMING GENERAL SCENE MOTION, THERE DOES NOT EXIST A UNIQUE OPTICAL FLOW FIELD THAT IS CONSISTENT WITH IMAGE SEQUENCE.
(BECAUSE OF BOTH GEOMETRIC

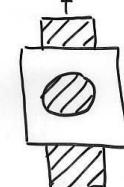
AND PHOTOMETRIC AMBIGUITY.)
⇒ NEED ADDITIONAL CONSTRAINTS
- GENERAL MOTION OF NON-RIGID OBJECTS ⇒ INFINITE NUMBER of INTERPRETATIONS THAT ARE CONSISTENT w/ DATA

- LOCAL CHANGES IN SPATIOTEMPORAL INTENSITY ARE NOT SUFFICIENT TO COMPLETELY DETERMINE OPTICAL FLOW

- EXAMPLES



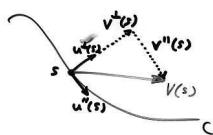
vs.



THEY LOOK THE SAME INSIDE THE CIRCULAR WINDOW!
I.E., DIFFERENT PHYSICAL MOTION,
BUT SAME MEASURABLE MOTION
IN FIXED WINDOW
APERTURE PROBLEM

APERTURE PROBLEM

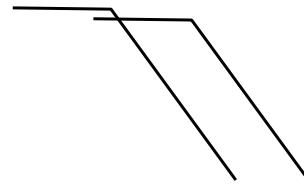
- GIVEN A STRAIGHT LINE IN MOTION AND A FIXED MEASUREMENT WINDOW, ONLY THE COMPONENT OF MOTION \perp TO LINE ORIENTATION IS RECOVERABLE



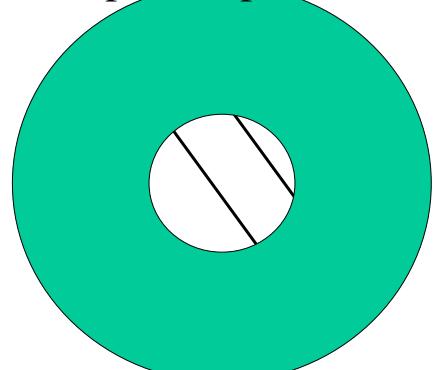
$$v(s) = v''(s) u''(s) + v^\perp(s) u^\perp(s)$$

$u^\perp(s), u''(s)$ are unit vectors
 $v^\perp(s)$ and are measurable
 $v''(s)$ computable
 $v''(s)$ not locally computable in general

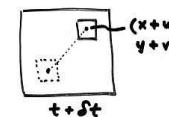
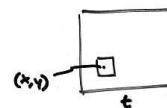
Aperture problem



Aperture problem



OPTICAL FLOW CONSTRAINT EQUATION



$$\text{Optical flow} = (u, v) \quad (\text{velocities})$$

$$\text{Displacement} = \begin{cases} \delta x = u \delta t \\ \delta y = v \delta t \end{cases}$$

Brightness Constancy Assumption:
Brightness pattern of patches is same in both images

$$\Rightarrow I(x + u \delta t, y + v \delta t, t + \delta t) = I(x, y, t)$$

- Expand LHS using Taylor series about pt (x, y, t) and keep only linear (1^{st} order) terms:

$$I(x, y, t) + \frac{\partial I}{\partial x} dx + \frac{\partial I}{\partial y} dy + \frac{\partial I}{\partial t} dt \approx I(x, y, t)$$

- Divide by dt and take limit $dt \rightarrow 0$:

$$\frac{dx}{dt} \frac{\partial I}{\partial x} + \frac{dy}{dt} \frac{\partial I}{\partial y} + \frac{\partial I}{\partial t} = 0$$

- Let $u = \frac{dx}{dt}$, $v = \frac{dy}{dt}$

$$\Rightarrow \frac{\partial I}{\partial x} u + \frac{\partial I}{\partial y} v + \frac{\partial I}{\partial t} = 0$$

Optical Flow Constraint Eq.

- or

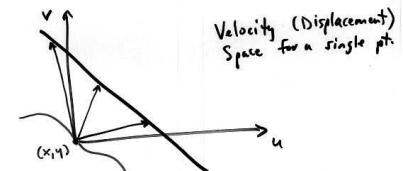
$$I_x u + I_y v + I_t = 0$$

Given: 2 images, $I_1(x, y, t_0) = I(x, y, t_0)$
 $I_2(x, y, t_0 + dt) = I(x, y, t_0 + dt)$

Compute: $\frac{\partial I}{\partial x}$, $\frac{\partial I}{\partial y}$, $\frac{\partial I}{\partial t}$ using gradient operators

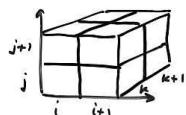
Solve for: 2 unknowns: u, v

Underconstrained: (u, v) must lie on a line in (u, v) space!



(u, v) cannot be found uniquely!

Computing $\frac{\partial I}{\partial x}$, $\frac{\partial I}{\partial y}$, $\frac{\partial I}{\partial t}$



$$\frac{\partial I}{\partial x} = I_x = \frac{1}{4} (I_{i+1,j,k} + I_{i+1,j,k+1} + I_{i,j+1,k} + I_{i,j+1,k+1}) - \frac{1}{4} (I_{i,j,k} + I_{i,j,k+1} + I_{i+1,j,k} + I_{i+1,j,k+1})$$

Similarly for $\frac{\partial I}{\partial y}$ and $\frac{\partial I}{\partial t}$

Computing Optical flow

- Formulate error in optical flow constraint

$$\epsilon_c = \iint_{\text{Image}} (I_x u + I_y v + I_t)^2 dx dy$$

where $I_x = \frac{\partial I}{\partial x}$, $I_y = \frac{\partial I}{\partial y}$, $I_t = \frac{\partial I}{\partial t}$

- Add additional constraint(s)

- Smoothness Constraint

- (Horn and Schunck, 1981)
- Nearby points in image undergo similar motion

⇒ Given n nearby points in domain D in image, assume $V = (u, v)$ is constant for all n points.
⇒ n linear equations in 2 unknowns

- Formulate error in image smoothness:

$$e_s = \iint_D \underbrace{\left(\frac{\partial u}{\partial x} \right)^2 + \left(\frac{\partial u}{\partial y} \right)^2}_{\|\nabla u\|^2} + \underbrace{\left(\frac{\partial v}{\partial x} \right)^2 + \left(\frac{\partial v}{\partial y} \right)^2}_{\|\nabla v\|^2} dx dy$$

- Find (u, v) at each image pixel that minimizes:

$$e = e_s + \lambda e_c$$

Optical Flow Algorithm (Discrete)

Consider image point (i, j)

- Departure of smoothness of optical flow at (i, j)

$$s_{ij} = \frac{1}{4} ((u_{i+1,j} - u_{ij})^2 + (v_{i+1,j} - v_{ij})^2 + (u_{i,j+1} - u_{ij})^2 + (v_{i,j+1} - v_{ij})^2)$$

- Error in optical flow constraint equation

$$c_{ij} = (I_x u_{ij} + I_y v_{ij} + I_t)^2$$

- Goal: Find the set $\{u_{ij}\}$ and $\{v_{ij}\}$ that minimize

$$e = \sum_i \sum_j (s_{ij} + \lambda c_{ij})$$

- Differentiating e wrt u_{kl} and v_{kl}

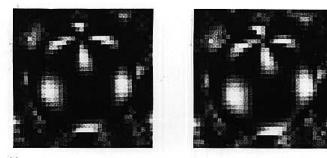
$$\frac{\partial e}{\partial u_{kl}} = 2(u_{kl} - \bar{u}_{kl}) + 2\lambda(I_x u_{kl} + I_y v_{kl} + I_t) I_x$$

$$\frac{\partial e}{\partial v_{kl}} = 2(v_{kl} - \bar{v}_{kl}) + 2\lambda(I_x u_{kl} + I_y v_{kl} + I_t) I_y$$

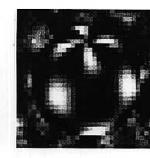
where \bar{u}_{kl} and \bar{v}_{kl} are averages of u and v around (k, l)

- Setting $\frac{\partial e}{\partial u_{kl}} = 0$ and $\frac{\partial e}{\partial v_{kl}} = 0$ we get:

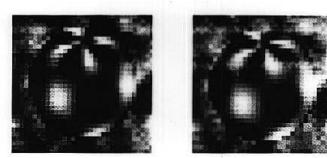
$$\left. \begin{aligned} u_{kl}^{n+1} &= \bar{u}_{kl} - \frac{I_x \bar{u}_{kl} + I_y \bar{v}_{kl} + I_t}{1 + \lambda(I_x^2 + I_y^2)} I_x \\ v_{kl}^{n+1} &= \bar{v}_{kl} - \frac{I_x \bar{u}_{kl} + I_y \bar{v}_{kl} + I_t}{1 + \lambda(I_x^2 + I_y^2)} I_y \end{aligned} \right\}$$



(a)



(b)



(c)

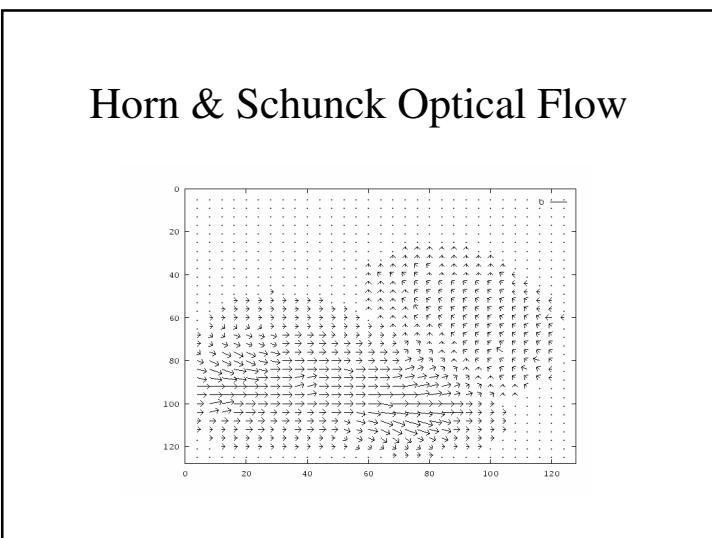
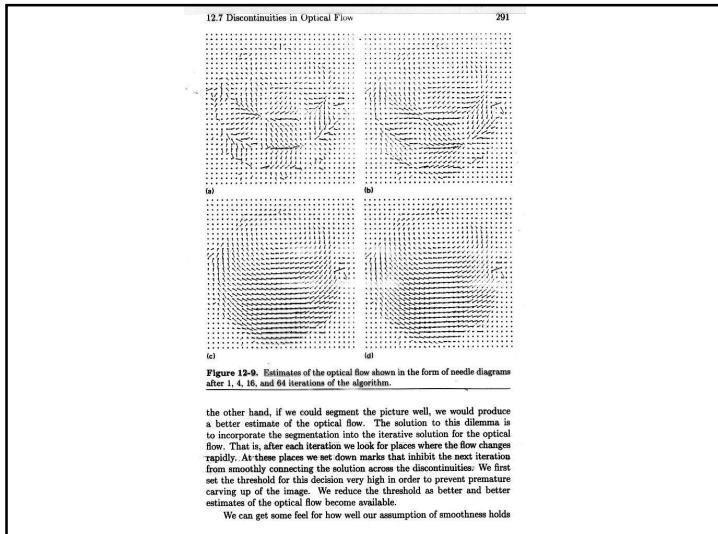


(d)

Figure 12-8. Four frames of a synthetic image sequence showing a sphere slowly rotating in front of a randomly patterned background.

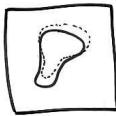
12.7 Discontinuities in Optical Flow

There will be discontinuities in the optical flow on the silhouettes, where one object obscures another. We must detect these places if we are to prevent the method presented above from trying to continue the solution smoothly from one region to the other. This seems like a chicken-and-egg problem: If we have a good estimate of the optical flow, we can look for places where it changes very rapidly in order to segment the picture. On

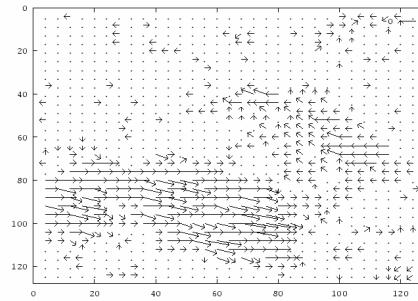


Results

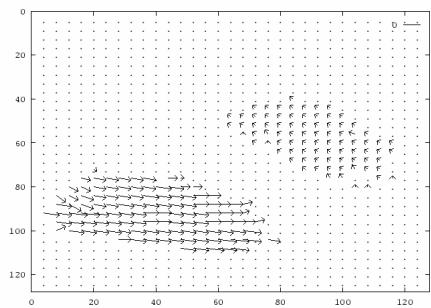
- * Works well for textured objects
- * Does not work well for homogeneous regions
- * Does not work well at points where optical flow is discontinuous. E.g. occluding contours
⇒ Mark these points and exclude them from the relaxation algorithm



Fleet & Jepson Optical Flow



Tian & Shah Optical Flow



Solving the Aperture Problem

- Basic idea: assume motion field is smooth
- Horn and Schunk: add smoothness term

$$\int \int (I_t + \nabla I \cdot [u \ v])^2 + \lambda^2 (\|\nabla u\|^2 + \|\nabla v\|^2) dx dy$$
- Lucas and 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 and Schunk

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

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

A
 25×2

d
 2×1

b
 25×1

Lucas-Kanade Flow

- Problem: more equations than unknowns

$$\begin{array}{c} A \ d = b \\ 25 \times 2 \quad 2 \times 1 \quad 25 \times 1 \end{array} \longrightarrow \text{minimize } \|Ad - b\|^2$$
- Solution: solve least squares problem
 - minimum least squares solution given by solution of:
$$(A^T A) d = A^T b$$

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

$A^T A$
 $A^T b$

 - The summations are over all pixels in the K x K window
 - This technique was first proposed by Lukas and Kanade (1981)

Conditions for Solvability

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

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

$A^T A$
 $A^T b$

When is this solvable?

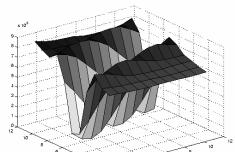
- $A^T A$ should be invertible
- $A^T A$ should not be too small due to noise
 - eigenvalues λ_1 and λ_2 of $A^T A$ should not be too small
- $A^T A$ should be well-conditioned
 - λ_1 / λ_2 should not be too large (λ_1 = larger eigenvalue)

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 $A^T A$?
 - 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

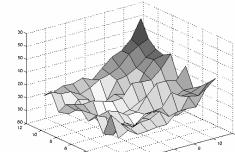
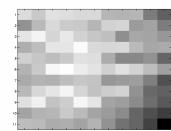
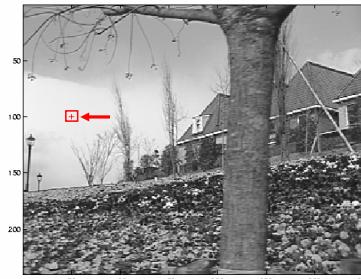
Edge



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

- large gradients, all the same
- large λ_1 , small λ_2

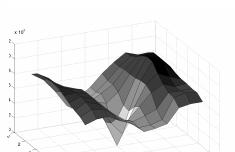
Low Texture Region



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

- gradients have small magnitude
- small λ_1 , small λ_2

High Texture 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 Lucas-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**
 - Also known as **Newton-Raphson** method
 - Lucas-Kanade method does one iteration of Newton's method
 - Better results are obtained with more iterations

Iterative Refinement

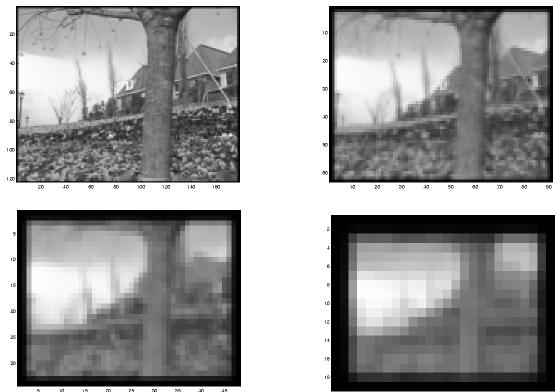
- Iterative Lucas-Kanade Algorithm
 1. Estimate velocity at each pixel by solving Lucas-Kanade equations
 2. Warp H towards I using the estimated flow field
 - *use image warping techniques*
 3. Repeat until convergence

Revisiting the Small Motion Assumption

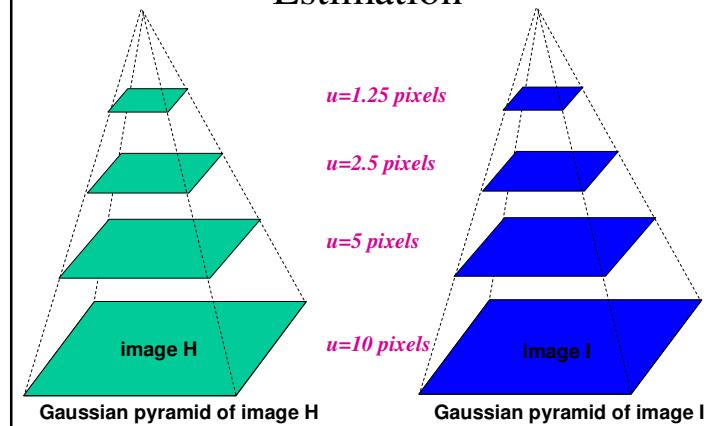


- When is the motion small enough?
 - Not if it's much larger than one pixel (2^{nd} order terms dominate)
 - How might we solve this problem?

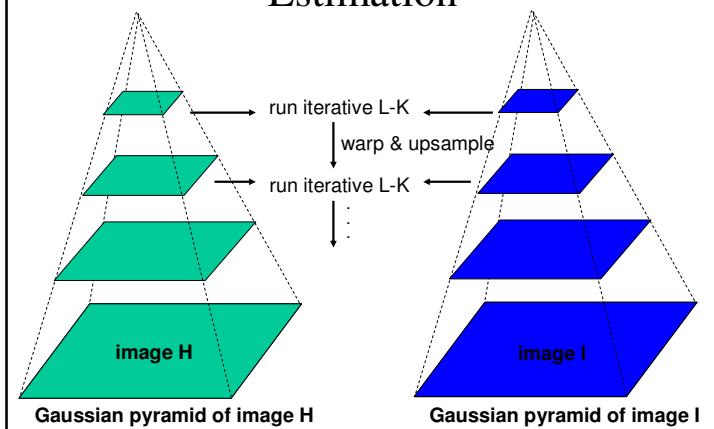
Reduce the Resolution



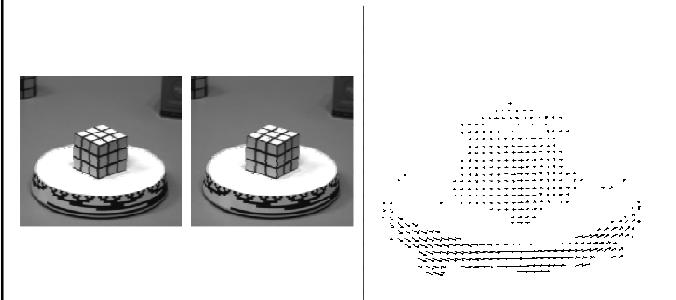
Coarse-to-Fine Optical Flow Estimation



Coarse-to-Fine Optical Flow Estimation



Optical Flow Result



Spatiotemporal (x-y-t) Volumes

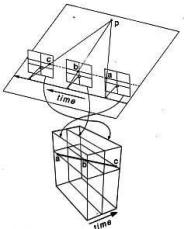


Figure 8.9 Geometry of image acquisition by a camera that is capturing images at equidistant intervals while the camera is translating horizontally in a direction parallel to the camera's image plane. Three positions of the camera's center of projection and its image plane are illustrated, and in each instance, the image position of the scene point p is indicated. The figure illustrates that a fixed scene point sweeps out a straight horizontal trajectory in the spatiotemporal stack of images that are captured by a camera at equidistant intervals while the camera is translating horizontally along a straight line parallel to its image plane. (From [Bolles, Baker, and Marimont 1987] with permission.)

8.1.2 Estimation Based on Spatiotemporal Coherence

Approaches to motion-field estimation that we are describing as based on spatiotemporal coherence apparently originated with Yamamoto [Yamamoto

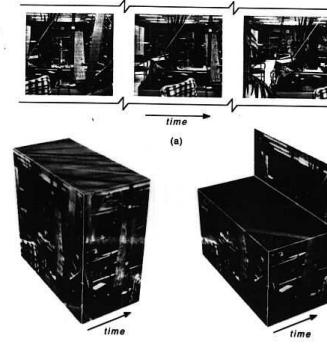


Figure 8.10 Spatiotemporal image data acquired by a camera that is capturing images at equidistant intervals while the camera is translating horizontally (from right to left) in a direction parallel to the camera's image plane. (a) Three frames of a 125-image sequence. (b) Oblique view of the image sequence in (a) after the sequence has been converted into a 3D spatiotemporal image data block. (c) Same as (b), except now the image data block is horizontally sliced; the top of the last image is shown in conjunction with the bottom of the first image to illustrate the genesis of horizontal spatiotemporal streaks in the image data block. (Adapted from [Bolles, Baker, and Marimont 1987] with permission.)

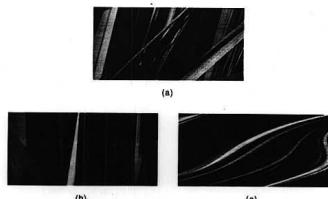


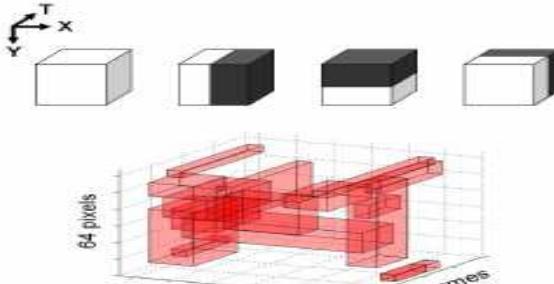
Figure 8.11 Slices of spatiotemporal image data blocks illustrating the image motions of scene points under various camera motions. In each instance, the individual images in the data block from which the slices were taken are at equidistant displacements of the camera. (a) This horizontal slice from Figure 8.10(c) illustrates the linear (horizontal) spatiotemporal image trajectories of scene points under linear (horizontal) translation of the camera parallel to its image plane. (b) This vertical slice from Figure 8.10(c) illustrates the linear (vertical) spatiotemporal image trajectories of scene points under linear translation of the camera in a direction not parallel to its image plane—the planar slice here is not horizontal, but rather is oriented to contain the focus of expansion, which sweeps out a straight line analogous to the individual image frames. (c) This slice, which is nonplanar in the direction of illustration, illustrates the nonlinear (and therefore, nonlinear) spatiotemporal image trajectories of scene points under motion of the camera that includes both linear translation and simultaneous rotation. (From [Bolles, Baker, and Marimont 1987] with permission.)

Now, the spatiotemporal image of a scene point is a coherent structure of one higher dimension: it is a curve. The spatiotemporal image of a scene curve is a surface, and the spatiotemporal image of a scene surface is a volume. Whereas we can easily determine the image velocity of a scene point that maps onto a locally discriminable spatiotemporal image curve by projecting the tangent to this curve (i.e., its time derivative) onto the image plane, the image velocity of a scene point whose spatiotemporal image can be determined only to lie on some discriminable spatiotemporal surface has a

Visual Event Detection using Volumetric Features

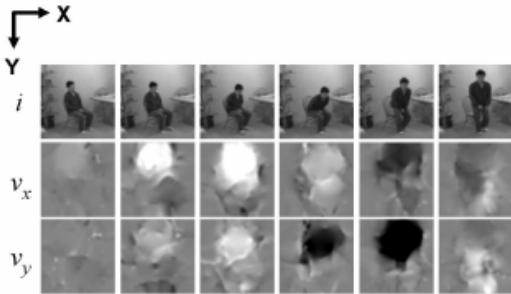
- Y. Ke, R. Sukthankar, and M. Hebert, CMU, CVPR 2005
- Goal: Detect motion events and classify actions such as *stand-up*, *sit-down*, *close-laptop*, and *grab-cup*
- Use $x-y-t$ features of optical flow
 - Sum of u values in a cube
 - Difference of sum of v values in one cube and v values in an adjacent cube

3D Volumetric Features



Approximately 1 million features computed

Optical Flow Features



Optical flow of *stand-up* action (light means positive direction)

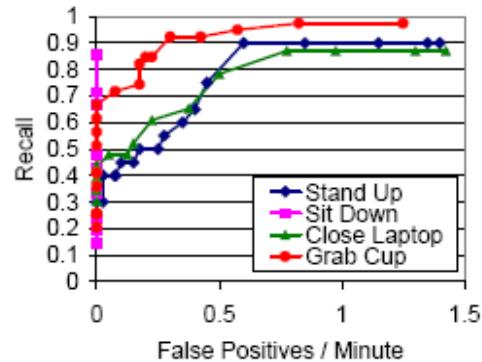
Classifier

- Cascade of binary classifiers that vote on the classification of the volume
- Given a set of positive and negative examples at a node, each feature and its optimal threshold is computed. Iteratively add filters at each node until a target detection rate (e.g., 100%) or false positive rate (e.g., 20%) is achieved
- Output of the node is the majority vote of the individual filters

Action Detection

- 78% - 92% detection rate on 4 action types: *sit-down*, *stand-up*, *close-laptop*, *grab-cup*
- 0 – 0.6 false positives per minute
- Note: while lengths of actions vary, the first frames are all aligned to a standard starting position for each action
- Classifier learns that beginning of video is more discriminative than end because of variable length
- Relatively robust to viewpoint (< 45 degrees) and scale (< 3x)

Results



Structure-from-Motion

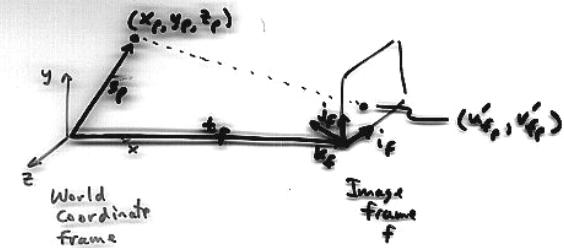
- Determining the 3-D structure of the world, and the motion of a camera (i.e., its extrinsic parameters) using a sequence of images taken by a moving camera
 - Equivalently, we can think of the world as moving and the camera as fixed
- Like stereo, but the position of the camera isn't known (and it's more natural to use many images with little motion between them, not just two with a lot of motion) and we have a long sequence of images, not just 2 images
 - We may or may not assume we know the intrinsic parameters of the camera, e.g., its focal length

FACTORIZATION METHOD

- Tomasi and Kanade, 1990
- GOAL: Compute both shape and motion. Shape computed in object-centered coordinate frame to avoid noise problems when small objects are far from camera
- ASSUMPTIONS:
 - Fixed scene
 - Moving Camera
 - Orthographic projection
 - Batch processing
 - No motion model
 - No shape model

- Input: From a sequence of F images, track P feature points over all frames (\Rightarrow must be visible entire time)

$$\{(u'_{fp}, v'_{fp}) \mid f=1, \dots, F; p=1, \dots, P\}$$



- Scene pt $s_p = (x_p, y_p, z_p)^T$ orthographically projects to (u'_{f_p}, v'_{f_p}) where

$$u'_{f_p} = i_f^T (s_p - t_f)$$

$$v'_{f_p} = j_f^T (s_p - t_f)$$

where i_f^T = unit vector in camera frame x-direction
 j_f^T = unit vector in camera frame y-direction
 $t_f = (a_f, b_f, c_f)^T$ translation vector from world coord origin to camera origin

- Let world coord. origin be at centroid of scene pts.

$$\Rightarrow \frac{1}{P} \sum_{p=1}^P s_p = 0$$

- Define image coords relative to centroid of image pts:

$$u_{f_p} = u'_{f_p} - \bar{u}_f$$

$$v_{f_p} = v'_{f_p} - \bar{v}_f$$

$$\text{where } \bar{u}_f = \frac{1}{P} \sum_{p=1}^P u'_{f_p}, \quad \bar{v}_f = \frac{1}{P} \sum_{p=1}^P v'_{f_p}$$

Hence,

$$\begin{aligned} u_{f_p} &= u'_{f_p} - \bar{u}_f \\ &= i_f^T (s_p - t_f) - \frac{1}{P} \sum_i i_f^T (s_p - t_f) \\ &= i_f^T ((s_p - t_f) - \frac{1}{P} \sum_i (s_p - t_f)) \\ &= i_f^T (s_p - \underbrace{\frac{1}{P} \sum_i t_f}_{=0}) \\ &= i_f^T s_p \end{aligned}$$

Use the fact that the centroid is preserved under orthographic projection.
I.e., projection of centroid = centroid of projections

- Similarly, $v_{f_p} = j_f^T s_p$

- Interpretation:

$[i_f, j_f]$ represents "motion" of fth frame
I.e., rotation of camera (6 unknowns)

$s_p = (x_p, y_p, z_p)$ represents "shape"
I.e., position of pth point in world coord. frame (3 unknowns)

\Rightarrow Each point in each frame produces 2 equations in 9 unknowns

$\Rightarrow P$ points in F frames produce $2FP$ equations

- Let $U = \begin{bmatrix} u_{11} & \dots & u_{1P} \\ \vdots & \ddots & \vdots \\ u_{F1} & \dots & u_{FP} \end{bmatrix}$ on frame f pts

pt 1 tracked over F frames

Similarly define V .

- Define "Measurement Matrix"

$$W = \begin{bmatrix} U \\ V \end{bmatrix} \quad \begin{array}{l} 2F \times P \text{ matrix} \\ \text{of "registered" image pts} \end{array}$$

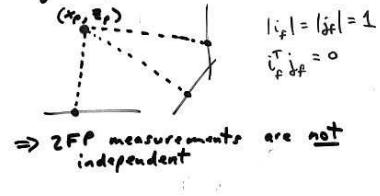
- Let $M = \begin{bmatrix} i_1^T \\ \vdots \\ i_F^T \end{bmatrix}$ on orientation of frame f's x-axis w.r.t world coord

2F x 3 "motion" matrix

$$S = [s_1 \dots s_P] \quad \begin{array}{l} 3 \times P \\ \text{"shape" matrix} \end{array}$$

- So, $W = MS$
given W , solve for M and S

- Problem is constrained because each pt projects in each frame along a "pencil" of projector lines:



- RANK Theorem:
Without noise, W is at most rank 3.
(since M is $2F \times 3$ and S is $3 \times P$)

- Noise corrupts W , so $\text{rank}(W) \neq 3$

By SVD decompose W :

$$W = L \Sigma R$$

$$\approx \begin{bmatrix} P \\ 2F \end{bmatrix} \begin{bmatrix} P \\ P \\ P \end{bmatrix} \begin{bmatrix} P \\ P \\ P \end{bmatrix}$$

Partitioning we get
 $L \Sigma R = L' \Sigma' R' + L'' \Sigma'' R''$

$$\text{where } L = \begin{bmatrix} L' & L'' \\ 0 & P-3 \end{bmatrix}_{2F}$$

$$\Sigma = \begin{bmatrix} \Sigma' & 0 \\ 0 & \Sigma'' \end{bmatrix}_{P-3}$$

$$R = \begin{bmatrix} R' \\ R'' \end{bmatrix}_{P-3}$$

8

By rank theorem, ideal matrix W^* has rank $\leq 3 \Rightarrow \sigma_4 = \dots = \sigma_p = 0$ for W^*

W = noisy version of W^*
 $\Rightarrow L'' \Sigma'' R''$ due to noise only
and $L' \Sigma' R'$ represents best approximation of W^*

Factorization Algorithm

1. Compute W
2. Compute $\text{SVD}(W) = L \Sigma R$
3. $\hat{M} = L' \Sigma'^{\frac{1}{2}}$ after partitioning
 $\hat{S} = \Sigma'^{\frac{1}{2}} R'$
(Hence set to 0 all but 3 largest singular vals)
4. Find 3×1 matrix A such that
 $M = \hat{M} A^T$ Nonlinear optimization problem
and $S = A \hat{S}$
5. Solve for M and S
 $\begin{cases} M = MA^T \\ S = A \hat{S} \end{cases}$

Results

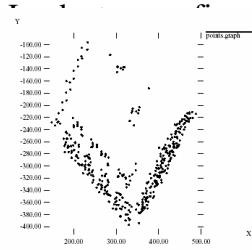


Figure 4.5: A view of the computed shape from approximately above the building (compare with figure 4.6).

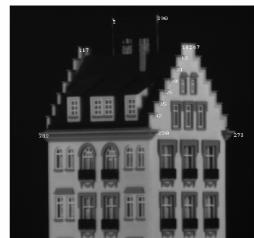


Figure 4.7: For a quantitative evaluation, distances between the features shown in the picture were measured on the actual model, and compared with the computed results. The comparison is shown in figure 4.8.

Extensions

- Paraperspective
 - [Poelman & Kanade, PAMI 97]
- Sequential Factorization
 - [Morita & Kanade, PAMI 97]
- Factorization under perspective
 - [Christy & Horaud, PAMI 96]
 - [Sturm & Triggs, ECCV 96]
- Factorization with Uncertainty
 - [Anandan & Irani, IJCV 2002]

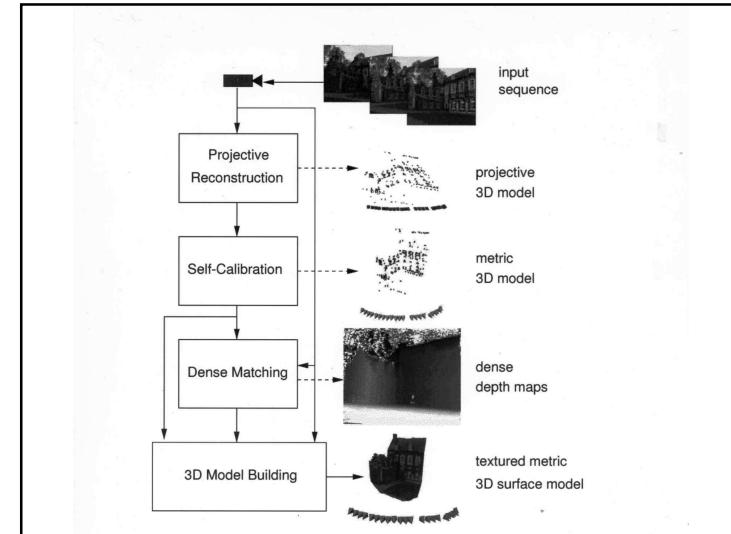
Stratified Approach to Structure from Motion

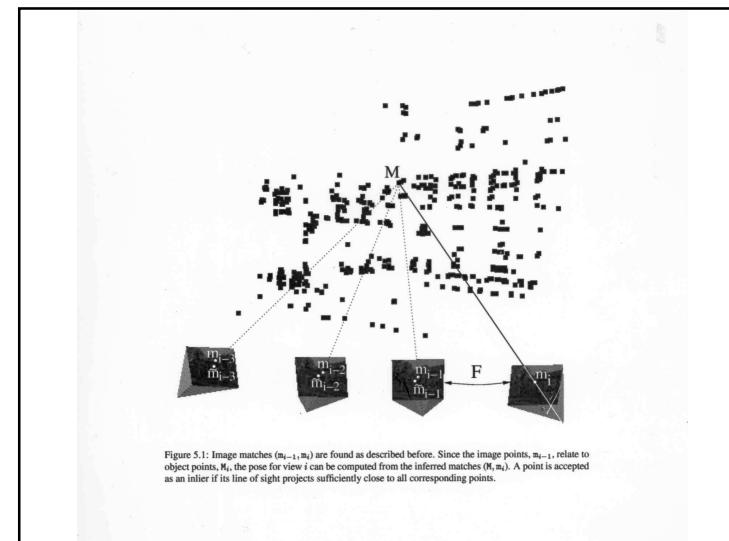
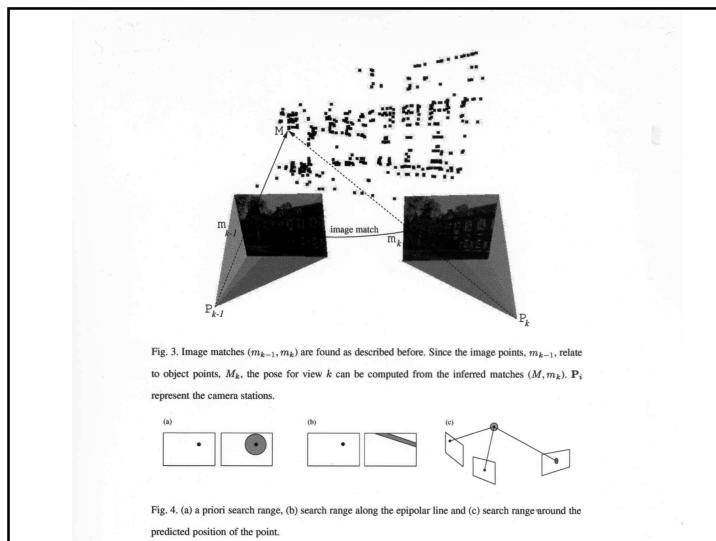
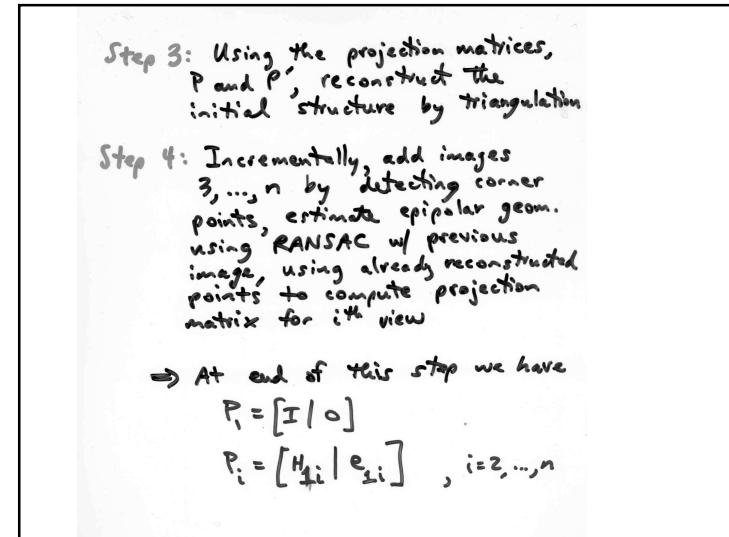
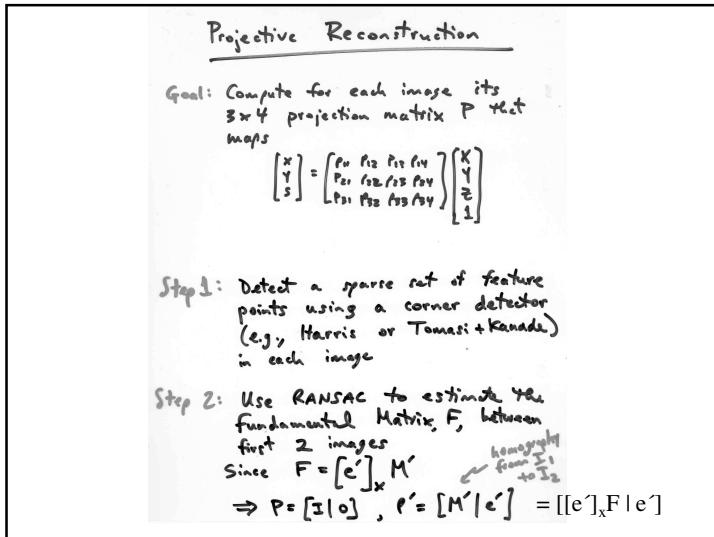
Given: Sequence of images from an uncalibrated camera moving arbitrarily around a static 3D scene

Goal: 3D model of scene

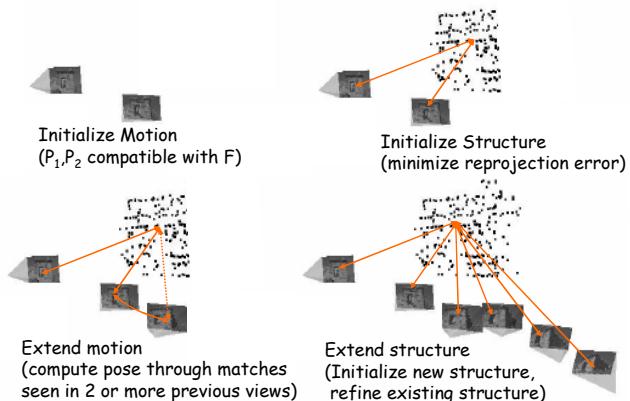
Approach:

1. Compute Projective reconstruction
2. Compute Affine reconstruction
3. Compute Metric reconstruction





Sequential Structure and Motion Computation



Sequential structure and motion recovery

- Initialize structure and motion from two views
- For each additional view
 - Determine pose
 - Refine and extend structure
- Determine correspondences robustly by jointly estimating matches and epipolar geometry

Metric Reconstruction

Goal: Euclidean reconstruction up to a scale factor (called Metric reconstruction)

Step 5: Self-Calibration
→ Compute the intrinsic camera parameters
E.g. add constraints on intrinsic parameters such as
1) no "skew"
2) square pixels

Note: these constraints do not require that other intrinsic parameters not vary over time
e.g. zoom can change

Result is a new set of projection matrices of form:

$$P_i = K_i [R_i | -R_i t_i]$$

$$\text{where } K_i = \begin{bmatrix} a_u & s & u_0 \\ 0 & a_v & v_0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$\begin{aligned} a_u &= f k_u & \frac{du}{dx} &= \text{aspect ratio of pixel} \\ a_v &= f k_v & (u_0, v_0) &= \text{principal point} \\ s &= \text{image skew} \end{aligned}$$

One common method for self-calib:
Compute the "absolute conic",
a virtual conic corresponding to
the unique degenerate quadric
of planes that is invariant
under all rigid transformations

Step 6: Compute dense correspondence

First, rectify images

Then detect correspondences
using Epipolar constraint

2) uniqueness constraint

3) ordering constraint
(monotonicity)

Step 7: Depth map fusion and uncertainty reduction

→ Find chains of correspondences
across subsequences of images

→ Intersect depth estimates

Pollefeys' Result



Object Tracking

- 2D or 3D motion of known object(s)
- Recent survey: “Monocular model-based 3D tracking of rigid objects: A survey” available at <http://www.nowpublishers.com/>

XVision Tracker

- G. Hager and K. Toyama, CVIU, 1998
- GOAL: REAL-TIME TRACKING USING PC
- MAIN IDEAS:
 - 1) HIERARCHICAL
 - SMALL SET OF PRIMITIVE FEATURE TYPES
 - COMPLEX OBJECTS BUILT FROM SIMPLER FEATURES
 - 2) STATE VECTOR DESCRIPTION
 - EACH PRIMITIVE FEATURE DESCRIBED AS A STATE VECTOR
 - ⇒ SMALL # PARAMETERS THAT SPECIFY POSITION AND APPEARANCE
 - 3) WINDOW-BASED, REGIONS-OF-INTEREST

- 4) TRACKING = INVERSE ANIMATION
- IMAGE WARPING TRANSFORMS RAW IMAGE WINDOW INTO CANONICAL CONFIG.
 - \Rightarrow "RECTIFY" FEATURES, THEN DETECT
- 5) MOTION MODEL = TRANSFORMATION FOR WARPING WINDOWS
- 2D RIGID (3 DOF)
 - 2D AFFINE (6 DOF)

PRIMITIVE FEATURE : EDGE (LINE SEGMENT)

- STATE VECTOR AT TIME t :

$$L_t = (x_t, y_t, \theta_t, r_t)^T$$

"appearance"
i.e., ideal filter response

- STATE VECTOR UPDATE:

$$L_{t+\delta t} = L_t + \text{Edge}(\text{Warp}_{\text{rotation}}(I(t+\delta t); x_t, y_t, \theta_t); L_t)$$

- \Rightarrow 1) Estimate orientation at $t+\delta t$
- 2) Apply rotational warp to obtain vertical edge
- 3) Detect position of vertical edge

More specifically,

- i) Warp window at several (3) orientations, $+90^\circ, 0^\circ, -90^\circ$ wrt orientation at time t .

- Equiv.
j) Sum columns
i) Compute difference b/w adjacent column sums

- k) Mark diff = edge strength

- l) Quadratic interpolation to estimate orientation

Result = Estimate of translation $\overset{\text{class}}{\delta \text{trans}}$,
in direction \perp to edge,
and orientation change, $\delta\theta$

- m) Update State Vector:

$$\begin{bmatrix} x_t + \delta t \\ y_t + \delta t \\ \theta_t + \delta t \end{bmatrix} = \begin{bmatrix} x_t \\ y_t \\ \theta_t \end{bmatrix} + \begin{bmatrix} -\delta \text{trans} \sin(\theta_t + \delta\theta) \\ \delta \text{trans} \cos(\theta_t + \delta\theta) \\ \delta\theta \end{bmatrix}$$

Figure 2. Close-up of tracking windows at two time points. Left, time t_0 , where the edge tracking algorithm has computed the correct warp parameters to make an edge appear (the "setpoint"). Right, the edge acquired at time t_{n+1} . The warp parameters from t_n were used to acquire the image, but the underlying edge has changed orientation. Figure 3 shows how the new orientation is computed.

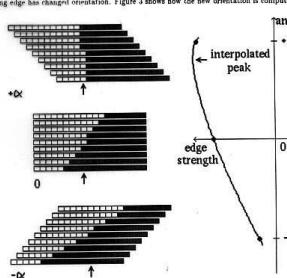


Figure 3. Schematic for computing edge orientation. The diagrams on the left show a window of pixels at three different "orientations". The middle figure displays the edge after a warped acquisition (Figure 2, right). The top and bottom figures show the effect of shifting rows to simulate orientational offset. Summing the columns for each figure and taking differences between adjacent sets of columns gives estimates for edge strength and orientation. The bottom figure shows the effect of shearing the image to simulate orientation offset in the window. At the right, these values are plotted with angle offset on the vertical axis and edge strengths on the horizontal axis. The three data points are fit to a quadratic, whose peak offers an estimate for the best angular offset of the actual edge. (Both the orientation of the edge in the middle figure and the extent of the shearing in the top and bottom figures have been exaggerated for illustrative purposes.)

PRIMITIVE FEATURE: REGION

- Given 2D Template Window corresponding to approx. planar scene surface, and motion model = affine

State Vector Description:

$$S_t = (x_t, y_t, \theta_t, s_x, s_y, \gamma_t, r_t)$$

pos, orient scale shear degrees of
motion w/ template

State Vector Update

$$S_{t+st} = S_t + SSD(w_{affine}, I(t+st); A_t, d_t)$$

\Rightarrow i) Apply estimated affine warp $\begin{cases} \text{translate} \\ \text{rotate, scale, shear} \end{cases}$
 ii) Compute degree of match using SSD
 iii) Brightness constancy assumption

- To estimate best affine transformation, solve optimization problem:

$$\min \sum_{\bar{x} \in W} (I(A\bar{x} + \bar{d}, t) - I(\bar{x}, t_0))^2 w(\bar{x})$$

where $\bar{x} = (x, y)^T$ = pixel in window W

$\bar{d} = (u, v)^T$ = translation vector

$$A = \begin{bmatrix} s_x & \gamma \\ 0 & s_y \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}$$

rotation,
scale,
shear

$w(\bar{x})$ positive weighting function

- Given solution at t , (A_t, \bar{d}_t) , and assuming dt small, solve approx. optimization problem:

$$\min \sum_{\bar{x} \in W} (\nabla I(\bar{x}, t) \cdot (A' \bar{x} + \bar{d}') + (J(\bar{x}, t) - I(\bar{x}, t_0))^2 w(\bar{x}))$$

where
 A' = change in A during dt
 \bar{d}' = change in \bar{d} during dt
 $J(\bar{x}, t) = I(A_t \bar{x} + \bar{d}_t, t + dt)$
 ~ warped window

Value for 6 unknowns in A' and \bar{d}'

- To improve efficiency of solving,
 - first solve for translation, rotation
 - compute scale and shear change

- To cope with illumination change normalize windows to have 1st moment = 0
 2nd moment = 1

- To cope with motion > fraction of pixel select resolution adaptively:
 if $|\bar{d}'| >$ threshold,

then halve resolution at $t+dt$
 (implement using warping
 + interpolation)

- else if $|\bar{d}'| <$ threshold
then double resolution at $t+dt$

Size	20 × 20	40 × 40	60 × 60	80 × 80	100 × 100
Rotational Warping	0.11	0.49	0.94	1.77	2.83
Scale and Shear ²	0.39	1.32	2.01	3.11	4.83
Resampling by 2	0.04	0.14	0.54	1.02	1.67
Resampling by 4	0.04	0.13	0.30	0.59	0.97
Affine	0.50	1.92	4.39	7.88	12.52
Affine by 2	0.56	2.15	4.93	8.90	14.19
Affine by 4	0.54	2.05	4.69	8.47	13.47

Figure 1. The time in milliseconds consumed image warping for various scene regions. The first two lines show the time for each of the warping stages. The third and fourth lines show the time taken for reducing image resolution by a factor of 2 and 4. The final lines show the time needed for affine warping at various scales based on the component times.

Edge Tracking

20 × 20	0.39 msec	on 175 MHz
40 × 40	1.17	R4400 processor
60 × 40	2.09	

Region (SSD) Tracking

	Rigid	Affine
40 × 40	5.6 / 1.5 *	8.4 / 3.7
60 × 60	9.5 / 3.2	15.6 / 8.1
80 × 80	17.7 / 6.1	28.5 / 14.4

* 1/2 resolution / 1/4 resolution

