lecture 24

review for final exam

Tips for final exam

Technical things you do NOT need to review from lectures 1-8:

- blur (thin lens equation)
- slant & tilt
- radiometry definitions (radiance, irradiance, BRDF, camera response, color, vignetting, dynamic range)

General Principle

- course had 23 lectures

∴ $\sim \frac{100}{23} \simeq 4\frac{1}{2}$ point per lecture

- I have already used 70 % (10 midterm + 60 assignments)

∴ I will try to use the last 30 % to cover "other" material

Image Processing (lectures 8-15)

- ~~edge detection (Canny)~~ A2
- ~~convolution~~ A2, A3
- ~~scale space constructions~~ A3
- ~~blob detection~~ A3
- image registration A 3 + final
- SIFT final
- vanishing points final
- least squares final
- Hough and RANSAC final

3D Scene Inference (lectures 16-23)

- shape from shading final
- ~~shape from texture, depth from defocus~~ A3
- SVD final ~~non-linear least squares~~
- ~~camera calibration~~ A4
- homographies A4 (also final)
- structure from motion: factorization final
- stereo 1: epipolar geometry final
- stereo 2: correspondence final

Vision as "inverse optics"

(given an image, infer the 3D scene that produced it)

Key difficulty: images often do not fully specify the 3D scene. (non-invertibility)

Given a, b, c

and a model

$au + bv + c = 0$,

what is u, v ?

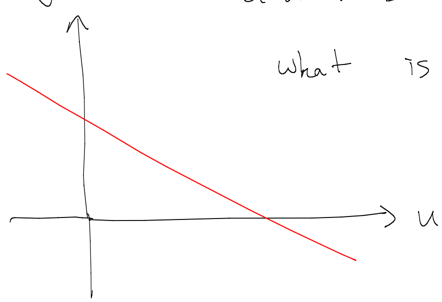

Solution: use more constraints / data

Given $\{(a,b,c)_i\}$ and a model

$a_i u + b_i v + c_i = 0$,

what is u, v ?

Example: Image Registration

$I(x,y)$ $J(x,y)$

$$\frac{\partial I}{\partial x} h_x + \frac{\partial I}{\partial y} h_y + I - J = 0$$

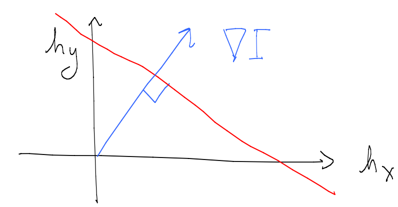

To estimate $(h_x \; h_y)$ we need more constraints.

$$\frac{\partial I}{\partial x} h_x + \frac{\partial I}{\partial y} h_y + I - J = 0$$

@ (x_1, y_1) and (x_2, y_2)

gives solution if ∇I directions differ

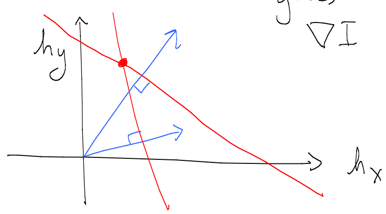

Fitting a model with no noise

- $A\vec{x} = b \Rightarrow$ solve for $\vec{x} \neq 0$
- $A\vec{x} = 0 \Rightarrow$ find null space (SVD)

Other Examples

- vanishing points $(x_v \; y_v)$ — $Ax = b$
- camera calibration (P) — $Ax = b$
- photometric stereo $(\rho \vec{n})$ — $Ax = b$
- structure from motion (RS) — $Ax = 0$
- homographies (H) — $Ax = 0$
- fundamental matrix (F) — $Ax = 0$

Finding a null vector using SVD

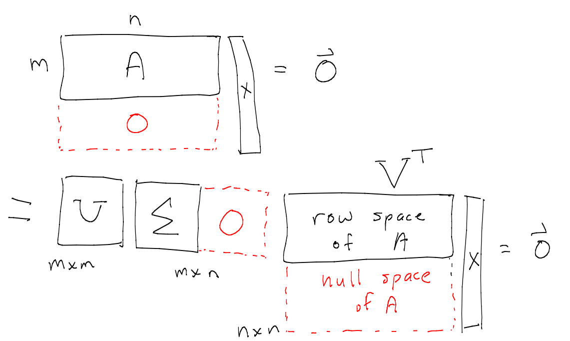

MATLAB's SVD gives $V_{n \times n}$.

Vision as "inverse optics":

given an image, infer the parameters of a model that explains the image

Key difficulties:

- image data may not fully constrain the solution (use more data, constraints)
- images have noise $\Rightarrow$ least squares
- not all data fits the model $\Rightarrow$ ignore outliers (Hough, RANSAC)

Examples

- vanishing points $(x_v \; y_v)$
- camera calibration (P)
- photometric stereo $(\rho \vec{n})$
- structure from motion (RS)
- homographies (H)
- fundamental matrix (F)

$$A\vec{x} = 0$$
$$A\vec{x} = b$$

$$A = U \Sigma V^T$$

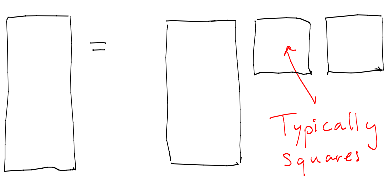

Typically a least squares model assumes certain singular values are 0.

This gives the solution.

Sometimes, though, adding more data (and ignoring outliers) and setting singular values to 0 still doesn't fully constrain the solution.

Factorization Method for Structure from Motion

$$\begin{bmatrix} x_{ki} \\ y_{ki} \end{bmatrix}_{2F \times N} = \tilde{R}_{2F \times 3} \; S_{3 \times N}$$

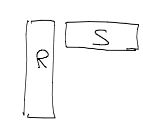

$$= \left(\tilde{R}_{2F \times 3} \, Q_{3 \times 3}\right)\left(Q^{-1}_{3 \times 3} \, S_{3 \times N}\right)$$

Projective Reconstruction Ambiguity

$$\begin{pmatrix} w_1 x_1 \\ w_1 y_1 \\ w_1 \end{pmatrix} = P_1 \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = P_1 M M^{-1} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

$$\begin{pmatrix} w_2 x_2 \\ w_2 y_2 \\ w_2 \end{pmatrix} = P_2 \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = P_2 M M^{-1} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

e.g. M could be translation, rotation, shear, or much more complex.

"Photometric Stereo"

(see shape from shading)

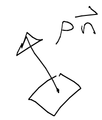

$$\begin{bmatrix} I_1 \\ I_2 \\ I_3 \end{bmatrix} = \underbrace{\begin{bmatrix} - L_1 - \\ - L_2 - \\ - L_3 - \end{bmatrix}}_{3\times 3} \underbrace{\begin{bmatrix} \rho \vec{n} \end{bmatrix}}_{3\times 1}$$

$$= \underbrace{\begin{bmatrix} L_1 \\ L_2 \\ L_3 \end{bmatrix}}_{3\times 3} Q\, Q^{-1} \underbrace{\begin{bmatrix} \rho \vec{n} \end{bmatrix}}_{3\times 3}$$

Shape from shading (see Exercises 2)

$$I(x) = \vec{n} \cdot \vec{L}$$

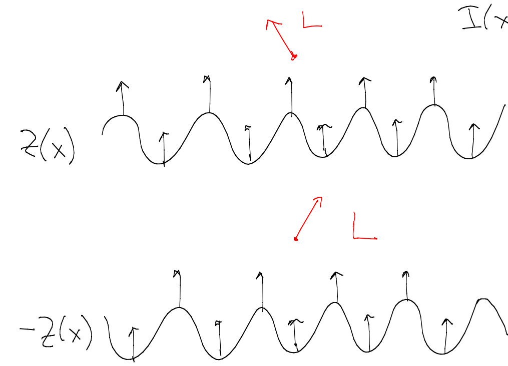

Hollow Mask Illusion

(see youtube video)

Non-linear constraints

- shape from shading often assumes surface has constant reflectance
- image registration $I(x+h) \approx J(x)$ and stereo correspondence
- orthogonal $x\ y$ camera axes in structure from motion (or skew $s=0$ in camera calibration)
- ....

$\Rightarrow$ often one poses a non-linear, outlier-robust, least squares problem

Other similarities between cues / problems that arise from geometry

Euclidean

$aX + bY + cZ = d$

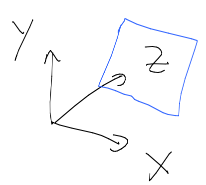

- shading on a sunny day
- structure from motion (factorization)

Projective

$ax + by + c = \frac{d}{Z}$

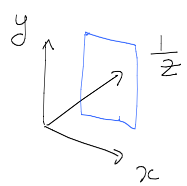

- vanishing points
- texture size & density
- depth from focus
- stereo disparity

Shading and texture

- surface normal & foreshortening
  - shading $(n \cdot \ell)$
  - texture $(n \cdot z)$

- distance
  - shading : vignetting, $1/z^2$
  - texture : size, density

flash camera & slanted plane

Think about shading (illumination) and texture information here.

Depth from focus and stereo

Stereo

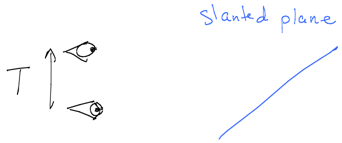

disparity varies $\sim \frac{T}{Z}$

Focus

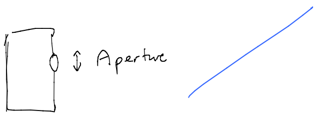

blur varies $\sim \frac{A}{Z}$

Homographies and Rectification

$$\begin{bmatrix} w\tilde{x} \\ w\tilde{y} \\ w \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Special cases

- $h_{31} = h_{32} = 0, \quad h_{33} \neq 0$
- $h_{31} x + h_{32} y + h_{33} = 0$

Image Rectification

Suppose I want a homography that maps $(x, y) = (0, 0)$ to the point at infinity in x axis direction

(Require the three rows are linearly independent)

$$\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} * & * & 1 \\ * & * & 0 \\ a & b & 0 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

More Computer Vision ?

Winter 2011

- ECSE 626 Statistical Computer Vision
- COMP 766 Shape Analysis in Computer Vision

Fall 2011

- COMP 646 Computational Perception

COURSE EVALUATIONS

Deadline Sunday Dec. 5