Wrap-up

https://cmu-ci-lab.github.io/16705

16-705 Fundamentals of Computer Vision
Spring 2026, Lecture 12

Course overview

Image features

Corner detection

Multi-scale detection

Haar-like

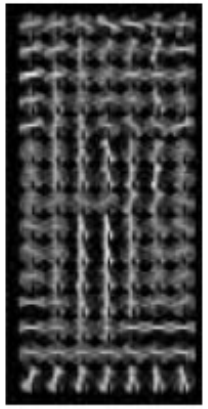

HOG

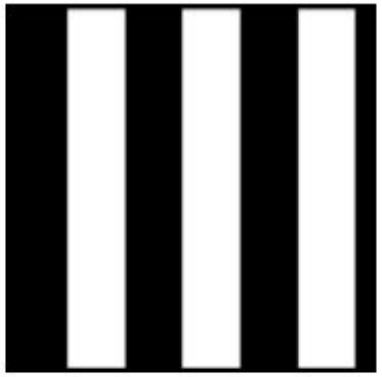

SURF

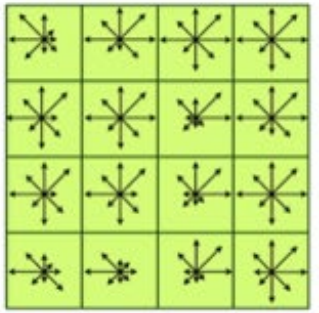

SIFT

Camera and multi-view geometry

$$x = \mathbf{P}\mathbf{X}$$

camera matrix

$\mathbf{P}$

pose estimation

$\mathbf{X}$

triangulation

$\mathbf{F}$

fundamental matrix

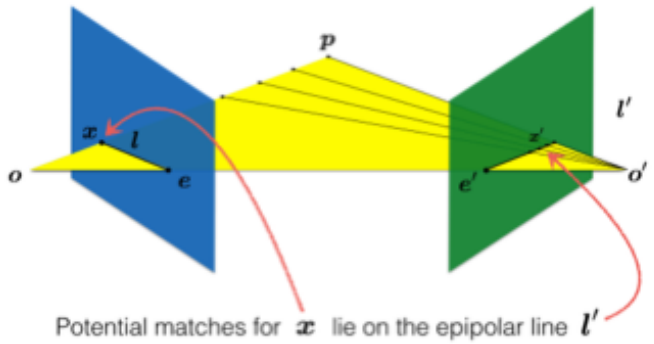

epipolar geometry

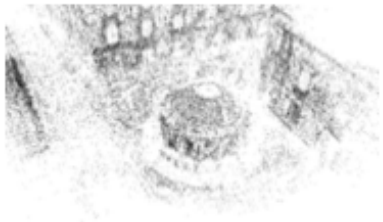

Reconstruction

Stereo

Stereo Rectification

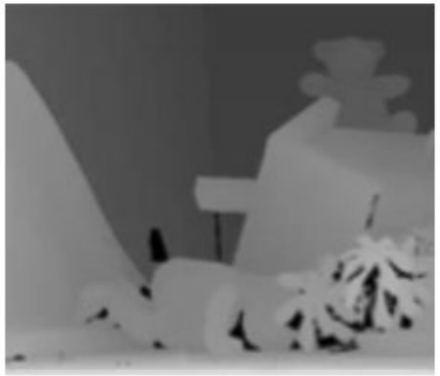

Block matching

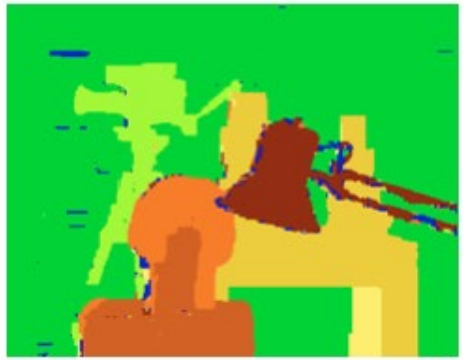

Energy minimization

Image formation and physics

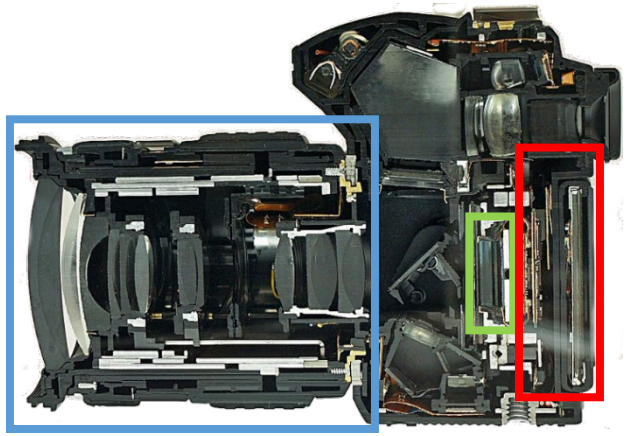

Image processing pipeline

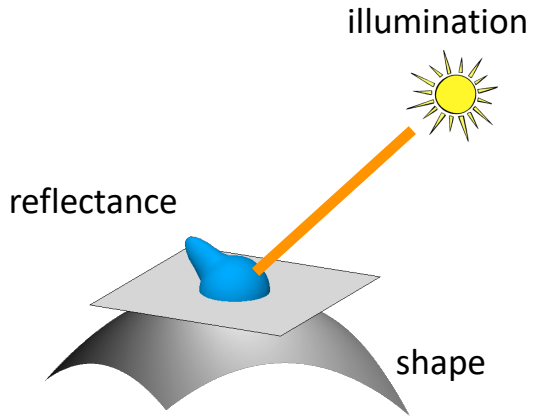

Image formation and rendering

Object recognition

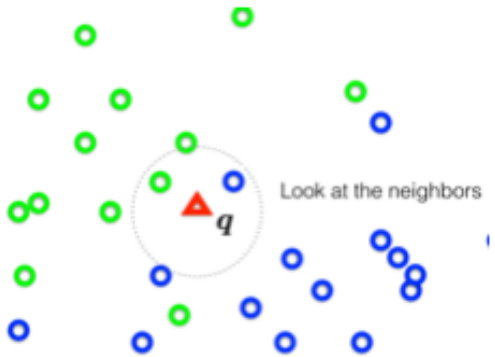

Nearest Neighbor

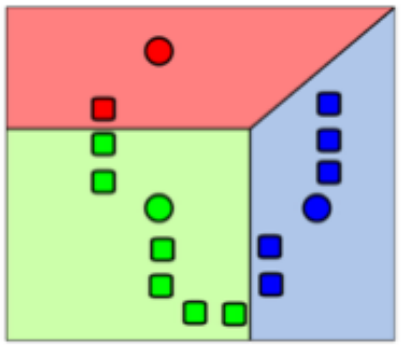

K-means

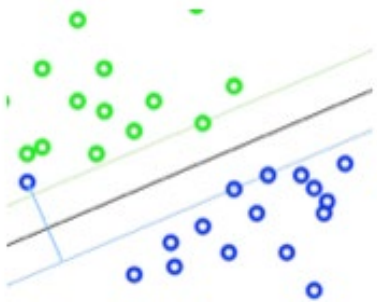

SVM

Neural networks

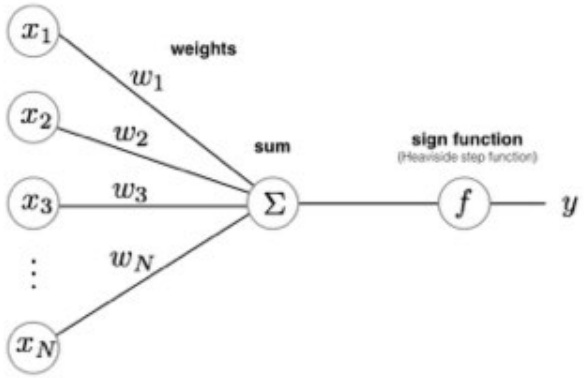

Perceptron

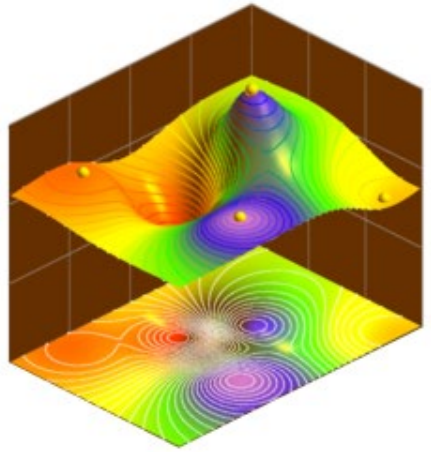

Gradient Decent

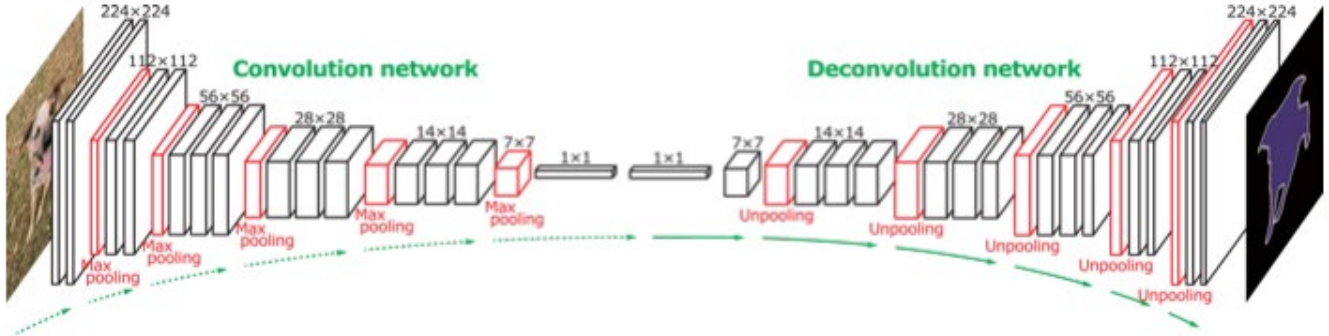

Convolutional Neural Networks

Optical flow and alignment

$$\begin{bmatrix} I_x(\boldsymbol{p}_1) & I_y(\boldsymbol{p}_1) \\ I_x(\boldsymbol{p}_2) & I_y(\boldsymbol{p}_2) \\ \vdots & \vdots \\ I_x(\boldsymbol{p}_{25}) & I_y(\boldsymbol{p}_{25}) \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix} = - \begin{bmatrix} I_t(\boldsymbol{p}_1) \\ I_t(\boldsymbol{p}_2) \\ \vdots \\ I_t(\boldsymbol{p}_{25}) \end{bmatrix}$$

Constant Flow

$$\min_{\boldsymbol{u},\boldsymbol{v}} \sum_{ij} \left\{ E_d(i,j) + \lambda E_s(i,j) \right\}$$

Horn-Schunck

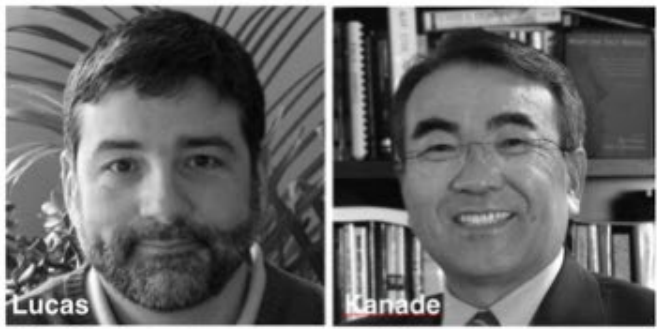

Lucas-Kanade
(Forward additive)

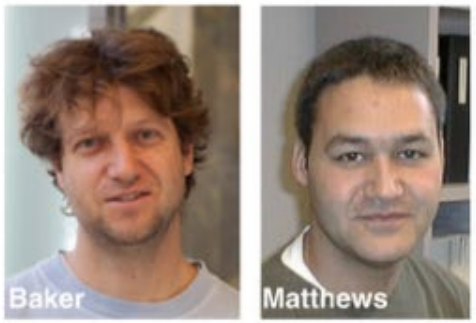

Baker-Matthews
(Inverse Compositional)

Tracking in videos

KLT

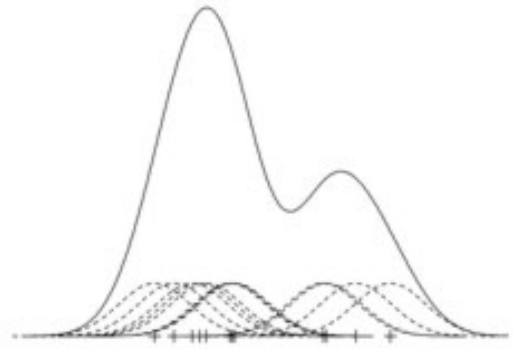

Mean shift

Neural rendering








Questions?

Do you plan on taking any other vision courses?

Any topics you wanted to learn more about?

Any topics you wanted to learn less about?

How were the assignments?