

Soliton

Vision for a Better World

Computer Vision with Machine Learning

PSG Tech - 2018





Dhivakar

Computer Vision and
Machine Learning,
Soliton Technologies



Senthil

Computer Vision and
Machine Learning,
Soliton Technologies



Measurement & Automation

Labview, Test Stand



Embedded Systems

FPGA, Embedded Systems, DSP, Device Driver



Computer Vision

Computer Vision, Machine Vision, Machine Learning



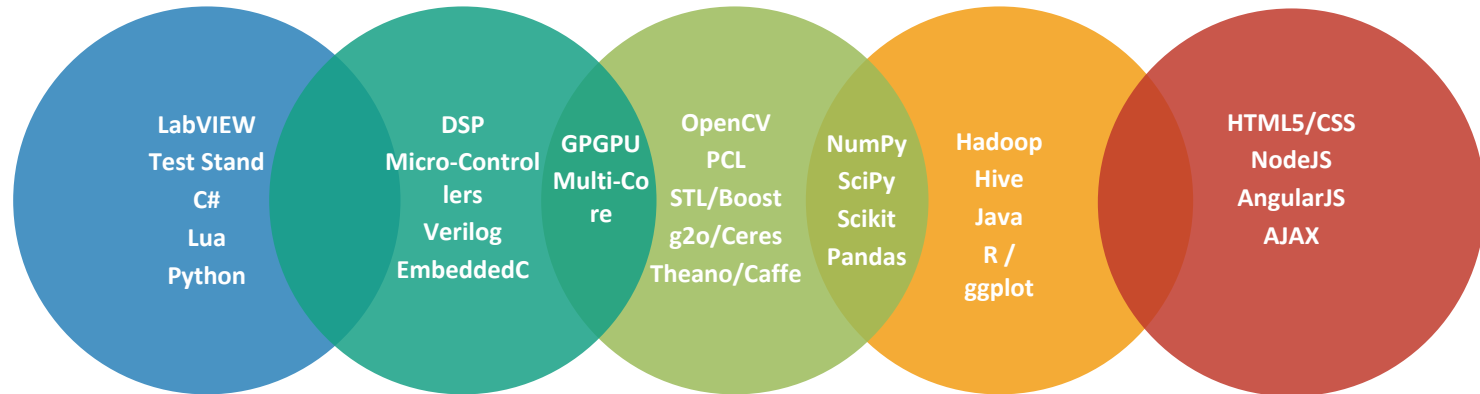
Data Analytics

Statistics, Machine Learning, Big Data, NLP, Text Analysis, Predictive Analytics



Web Technologies

Node.js, HTML, CSS, UI Design, C#, Java





Machine Vision
Cameras

06'



Spot It

**Smart Vision
Camera**

DSP, Image Processing
Library and Interactive
GUI for programming.

07'-08



Spot It V2

08-09



NEO
Smart Camera

Embedded DSP.
Fully Integrated
Functionality.

09-13'

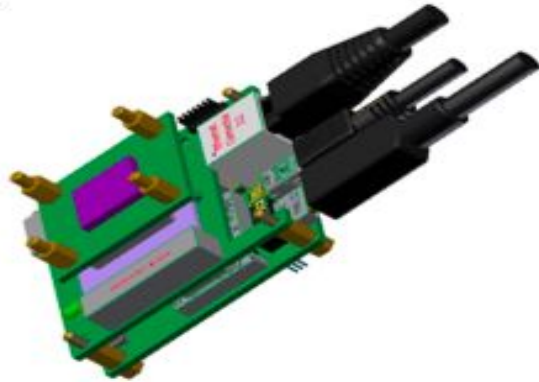


NUERA

2016 - 17



- ARM with 1.2 GHz processing power
- 8 GB Internal Flash memory and 1GB DDR3 RAM
- Powerful open source image processing library
- Option to interface 3 Image sensors (Parallel + MIPI)
- GIGE Ethernet, HDMI and Touch screen support
- GSM/4G support, WIFI/Bluetooth/Audio can be integrated
- RS232/CAN support to interface with PLCs
- Industrial standard with IP67 (Waterproof)
- Miniature design: 63.5 x 63.5 x 38 mm
- Interchangeable Lens
- **Deep Learning Support**



Optimization

MRF/DBN,
Graph Cuts, Linear Prog,
Convex Prog

Machine Learning

Regression,
SVM and Kernel
Methods,
Neural Networks, GMM,
Bayesian Techniques

Features/ Dimn Reduction

HOG, SIFT, SURF,
LBP, PCA, Sparse-PCA,
Isomap, Kernel-PCA,
MDS

3D Vision

Structured Light,
Time of Flight,
Reconstruction,
Structure from Motion

3D Reconstruction

Object Recognition

Object Tracking

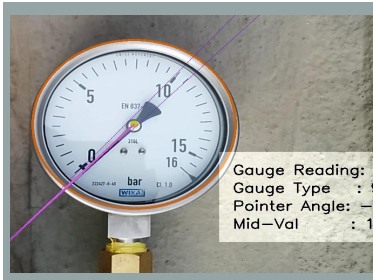
Depth Estimation

Structured Lighting
Point Cloud Generation
Post-processing and Metrology

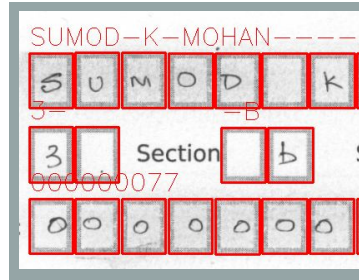
HOG, SIFT, SURF, LBP, FAST,
Boosting, SVM, ANN
PCA, Sparse-PCA, NCA, DBN

Region Based Tracking,
Kalman/Particle Filtering,
Markov Random Fields

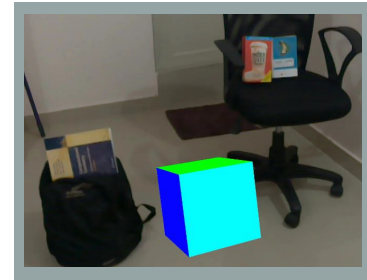
Dynamic Programming
Graph Cuts
Markov Random Fields



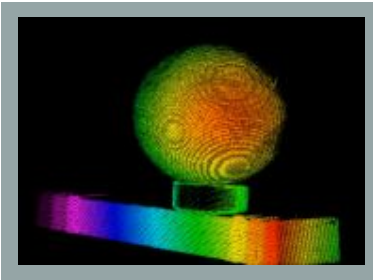
Project : Gauge Reader



Project : HCR



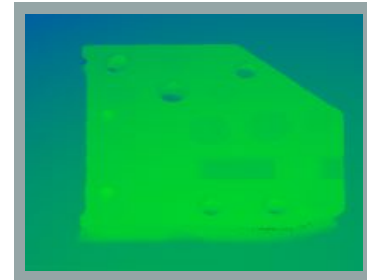
Project: Augmented Reality



Project : 3D Scanner



Project : Collision Detection



Product: 3D Pose Estimation

- We are in Machine Vision for the past 10+ years
- Deployed many multi-year projects in production
- First Smart camera from India
- Computer Vision
- 3D Vision
- Team of 12+ Engineers



Computer Vision & Machine Vision

Soliton Technologies

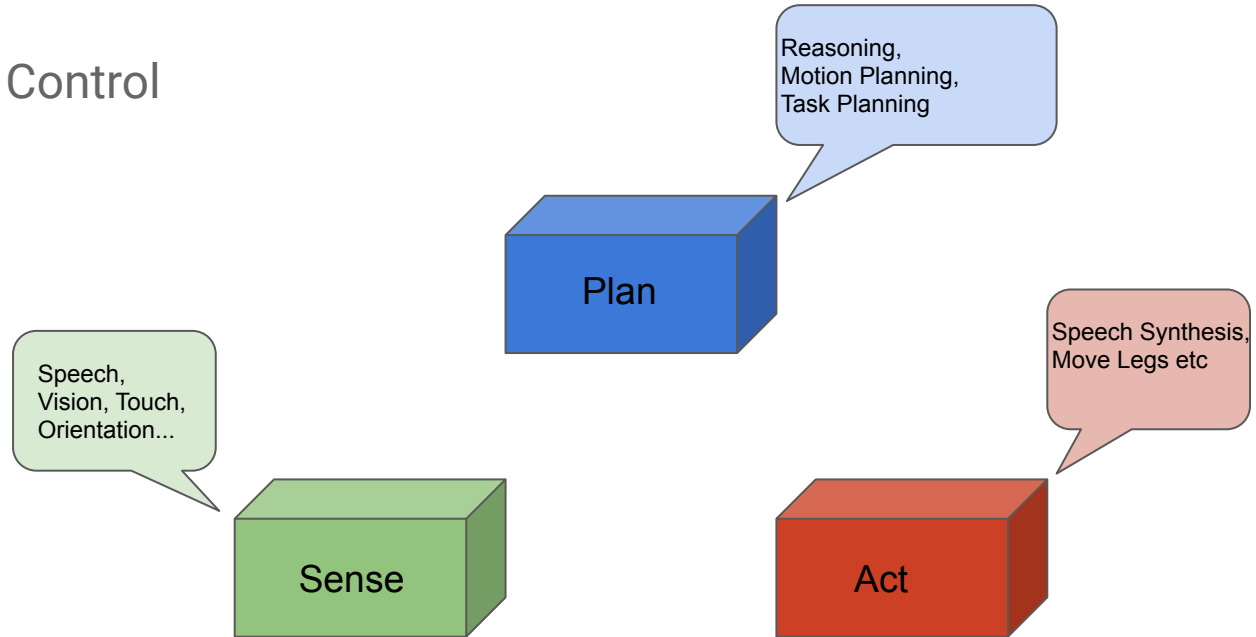


Let's start from us - Human

How we function..?

Humans in the World

- Perceive -> Reason -> Control

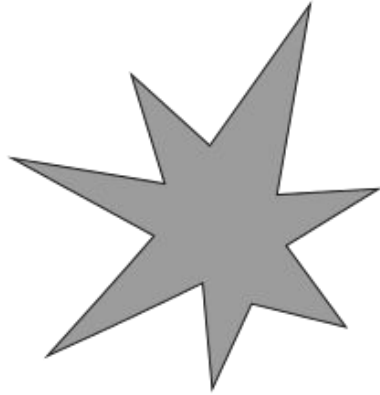


Human Vision

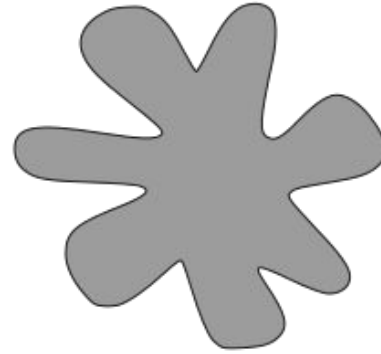
- Importance of vision for humans
 - 30% of cortex used in vision
 - McGurks Effect
- Decision makings
 - McGurk's Effect
 - Bouba / Kiki Experiment



Interconnections



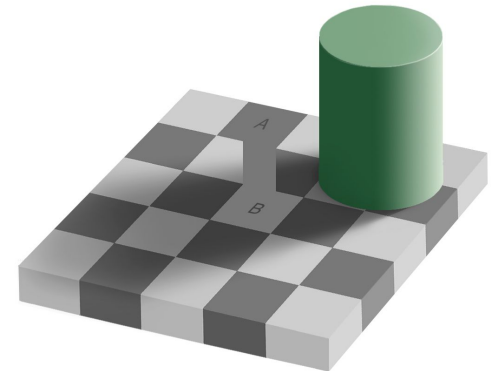
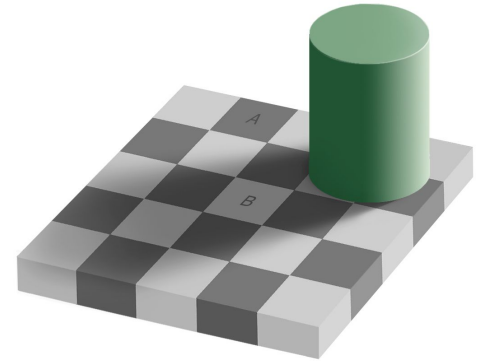
Kiki



Bouba

Human Vision

- MultiSensory, High Resolution, Immersive Inner Movie
- Arguable one of the most Complex Machines
 - Color Stereo Pair: 768 x 576 @ 30 fps : 80 MB/s
 - 10^{11} Neurons
 - 10^{14} - 10^{15} Synapses
 - Eons of evolution, learning right priors
 - Continuous flow of perception (with very few glitches)
- Perception
 - No eye inside an eye: just stream of electrical impulses :*Shades*
 - Brain combines them all to create our reality based on past experiences: *Sound Prediction & McGurks*
 - It is Brain's best guess of what is out there: *Ames Room*



Your Brain is Cheating you



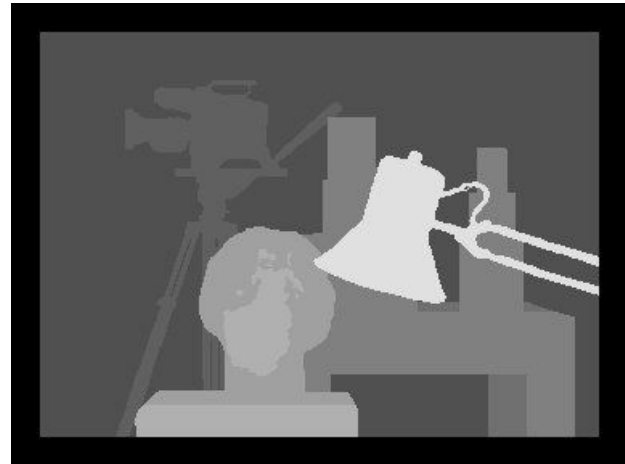
What is Computer Vision?

Ability of Computers to See and Understand from the Scene

- Saying if an object is present in the image
- Saying what are all the things present in the image
- Describe the scene

Computer Vision

- Image Processing, Computer Vision, Machine Vision
- Low Level: Edge Detection, Filtering..
- Mid Level: Segmentation..
- High Level: Object Detection, Stereo



Computer Vision in Games



Image Captioning



A person riding a motorcycle on a dirt road.



A group of young people playing a game of frisbee.



A herd of elephants walking across a dry grass field.

Deep Learning + Computer Vision in Prisma

Style transfer



Structure from Motion

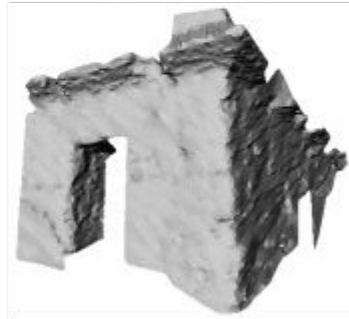
Input frames
(3 of 5)



Texture-mapped model



Shaded model



Output
3-D model

(Images courtesy of Luc Van Gool)



Seam Carving

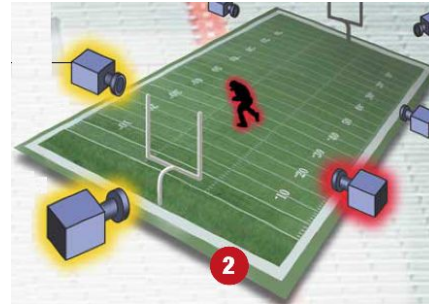
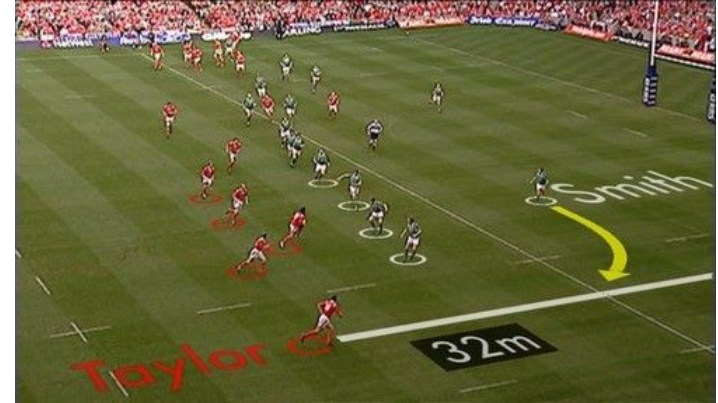
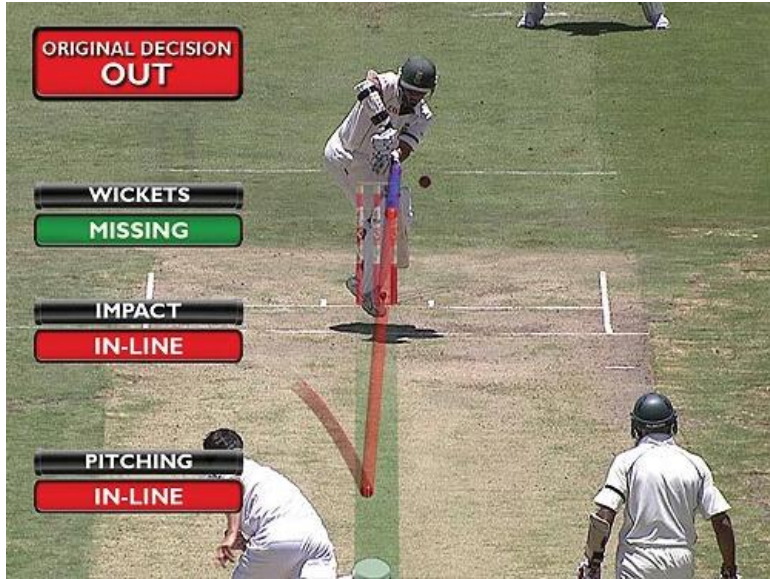


In Movies



- Bullet Time special-effect in The Matrix is an example of the application of SFM ideas.
- Linear array of cameras replaces moving camera.
- Green screen makes segmentation easy.

In International Games



Goals

Help building an understanding of How to solve real world problems using Computer Vision with ML and DL

Understand the Math and appreciate the beauty of it as opposed to simply Applying the technology

We can't teach you Vision in a 2 day Workshop. But we will try our best to Motivate you towards learning

Computer Vision

Digital image capture, storage and process.
Traditional Object detection and Limitations

Machine Learning

Evolution of ML. Understanding how Machines learn.
Limitations

Deep Learning

History and Biological motivation. Dive into Black Box (DL). Best practices

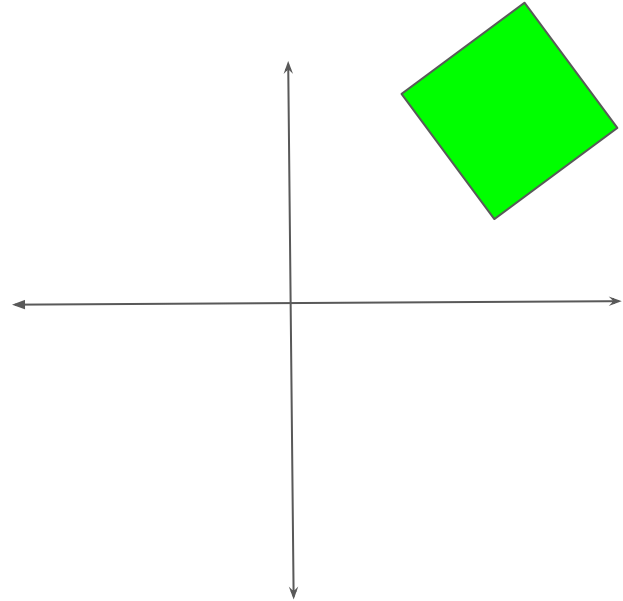
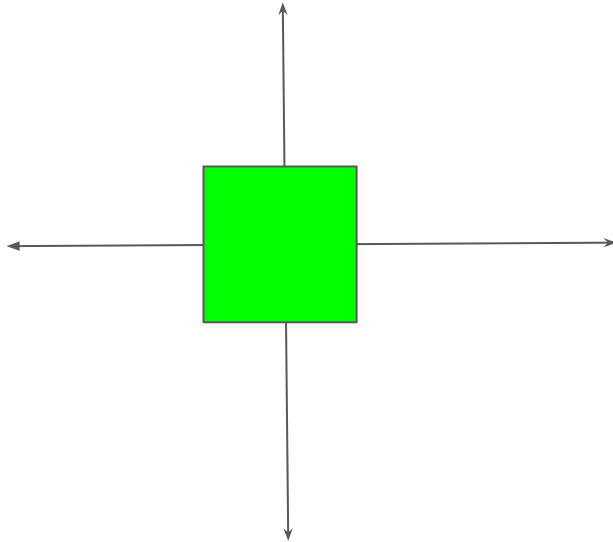
Euclidean Geometry?



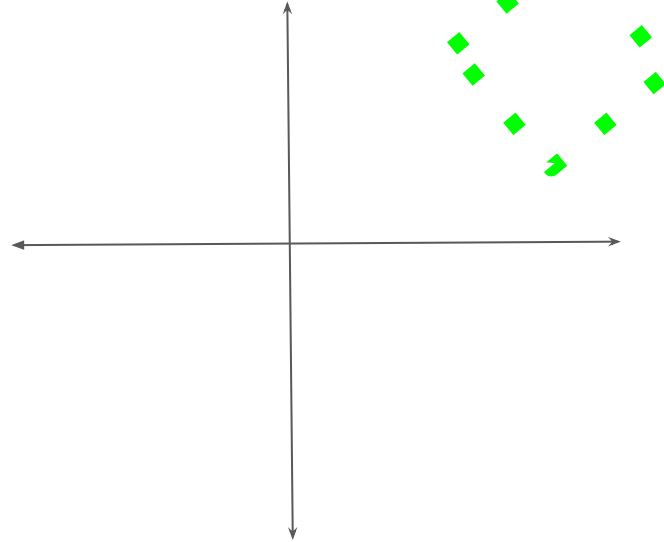
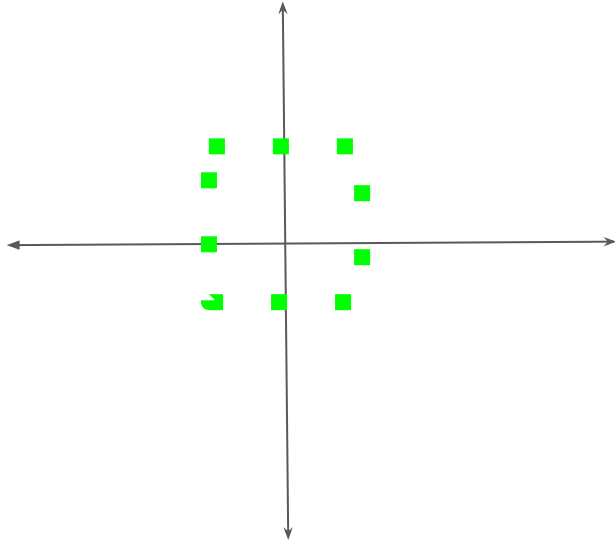
Projective Geometry



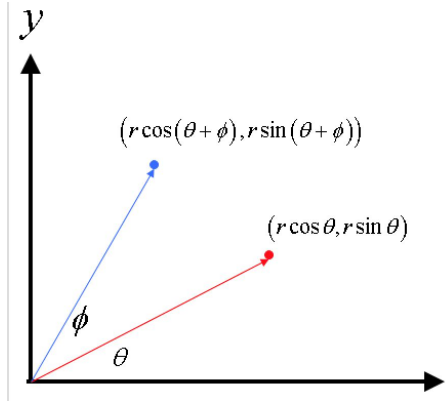
Motivation example - Euclidean Transform



Shapes are Complicated. Focus on Points



Transformation Matrices



Let's represent a point x, y in polar coordinates

$$x = r \cos \phi \quad y = r \sin \phi$$

Let θ be the rotation angle. Let x' and y' represent the rotated point. The rotated point can be expressed in polar coordinates as

$$\begin{aligned} x' &= r \cos(\theta + \phi) \\ &= r \cos \theta \cos \phi - r \sin \theta \sin \phi \\ &= x \cos \theta - y \sin \theta \end{aligned}$$

$$\begin{aligned} y' &= r \sin(\theta + \phi) \\ &= r \sin \theta \cos \phi + r \cos \theta \sin \phi \\ &= x \sin \theta + y \cos \theta \end{aligned}$$

Can we represent the same in Matrix?

$$\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} x\cos\theta - y\sin\theta \\ x\sin\theta + y\cos\theta \end{bmatrix}$$

Lets see how we can represent translation

But can we represent both translation and rotation in a single compact matrix, rather than splitting into a matrix multiplication and vector translation?

$$\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} tx \\ ty \end{bmatrix}$$

Solution - Homogeneous Coordinates

Lets represent each point by appending a one

$$\begin{bmatrix} x \\ y \end{bmatrix} \rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} tx \\ ty \end{bmatrix} = \begin{bmatrix} x\cos\theta - y\sin\theta + tx \\ x\sin\theta + y\cos\theta + ty \end{bmatrix}$$

Solution - Homogeneous Coordinates

$$\begin{bmatrix} \cos\theta & -\sin\theta & tx \\ \sin\theta & \cos\theta & ty \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} x\cos\theta - y\sin\theta + tx \\ x\sin\theta + y\cos\theta + ty \\ 1 \end{bmatrix}$$

Class I: Euclidean

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{bmatrix} \varepsilon \cos \theta & -\sin \theta & t_x \\ \varepsilon \sin \theta & \cos \theta & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}$$

(*iso*=same, *metric*=measure)

$$\varepsilon = \pm 1$$

orientation preserving: $\varepsilon = 1$

orientation reversing: $\varepsilon = -1$

$$\mathbf{x}' = \mathbf{H}_E \mathbf{x} = \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix} \mathbf{x}$$

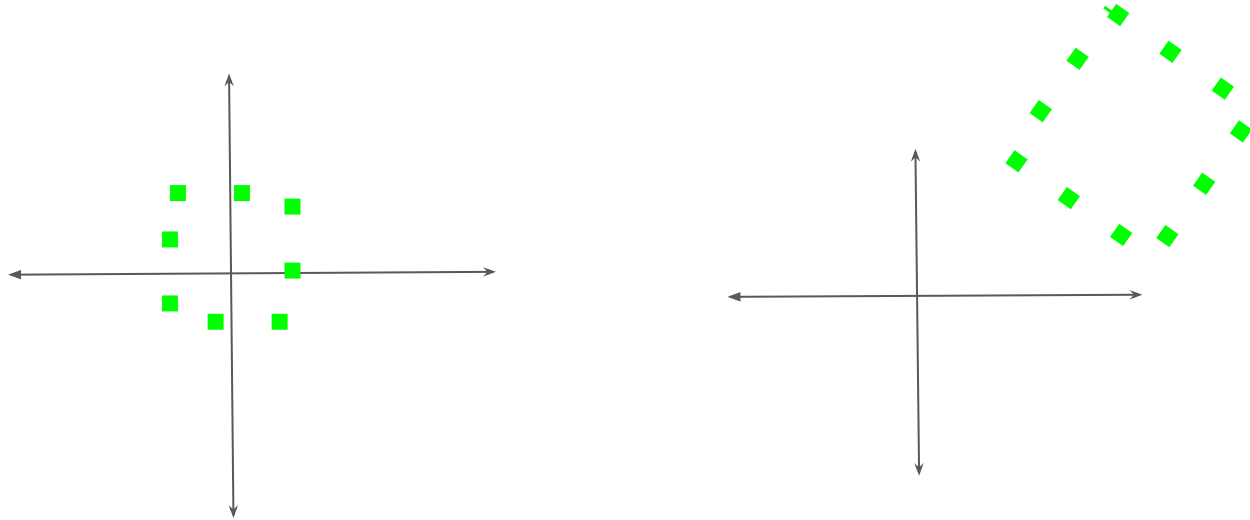
$$\mathbf{R}^\top \mathbf{R} = \mathbf{I}$$

3DOF (1 rotation, 2 translation)

special cases: pure rotation, pure translation

Invariants: length, angle, area

Ability to Scale up or Scale down



$$\begin{pmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = \begin{pmatrix} s_x x \\ s_y y \\ 1 \end{pmatrix} .$$

Ability to Scale up / Scale down

$$\begin{bmatrix} s & 0 & 0 \\ 0 & s & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} sx \\ sy \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} \cos\theta & -\sin\theta & tx \\ \sin\theta & \cos\theta & ty \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} s & 0 & 0 \\ 0 & s & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} s\cos\theta & -s\sin\theta & tx \\ s\sin\theta & s\cos\theta & ty \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Class II: Similarities

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{bmatrix} s \cos \theta & -s \sin \theta & t_x \\ s \sin \theta & s \cos \theta & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}$$

(isometry + scale)

also known as *equi-form* (shape preserving)

$$\mathbf{x}' = \mathbf{H}_S \mathbf{x} = \begin{bmatrix} s\mathbf{R} & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix} \mathbf{x}$$

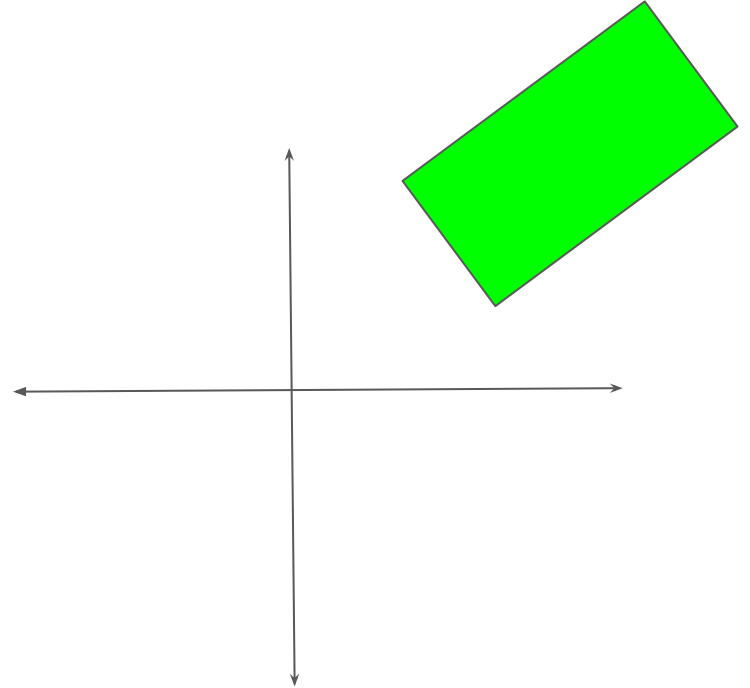
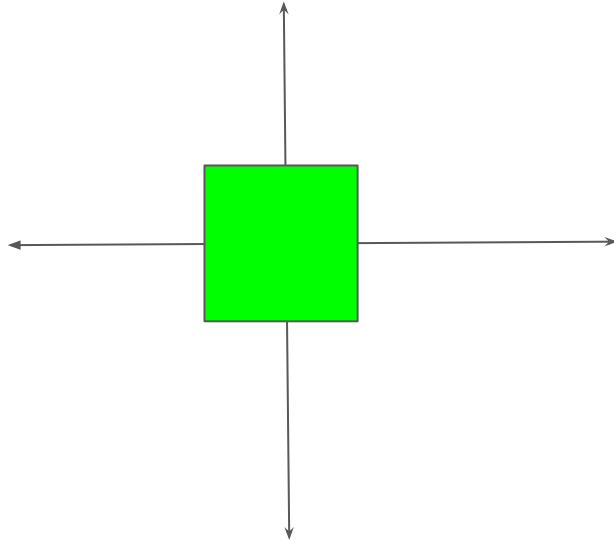
$$\mathbf{R}^\top \mathbf{R} = \mathbf{I}$$

4DOF (1 scale, 1 rotation, 2 translation)

metric structure = structure up to similarity (in literature)

Invariants: ratios of length, **angle**, ratios of areas, **parallel lines**

For Non-Uniform scaling?

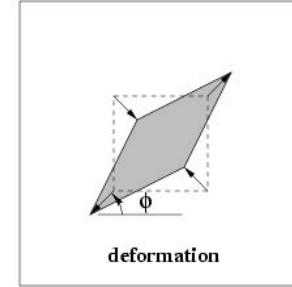
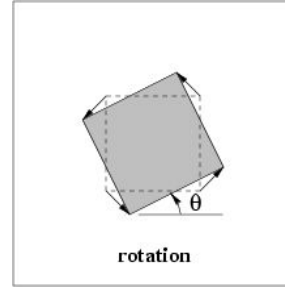


$$\begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} s_x x \\ s_y y \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} \cos\theta & -\sin\theta & tx \\ \sin\theta & \cos\theta & ty \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} s_x \cos\theta & -s_y \sin\theta & tx \\ s_x \sin\theta & s_y \cos\theta & ty \\ 0 & 0 & 1 \end{bmatrix}$$

Class III: Affine transformations

$$\begin{pmatrix} x' \\ y' \\ 1 \end{pmatrix} = \begin{bmatrix} a_{11} & a_{12} & t_x \\ a_{21} & a_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}$$



$$\mathbf{x}' = \mathbf{H}_A \mathbf{x} = \begin{bmatrix} \mathbf{A} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \mathbf{x} \quad \mathbf{A} = \mathbf{R}(\theta)\mathbf{R}(-\phi)\mathbf{D}\mathbf{R}(\phi) \quad \mathbf{D} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

6DOF (2 scale, 2 rotation, 2 translation)

non-isotropic scaling! (2DOF: scale ratio and orientation)

Invariants: parallel lines, ratios of parallel lengths, ratios of areas

General Affine Transform

$$\begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \text{ rotation} \quad \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \text{ translation}$$

$$\begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \text{ scaling} \quad \begin{bmatrix} 1 & s & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \text{ skew}$$

$$\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ 0 & 0 & 1 \end{bmatrix} \text{ general affine transform}$$

Projective Transform

All transforms learnt till now are planar transform

What happens when we try to map between two different planes

Why is this necessary? Why is it necessary to learn a transform that maps between planes?

Optical illusions - Inferring 3D world from 2D planes





Image formation

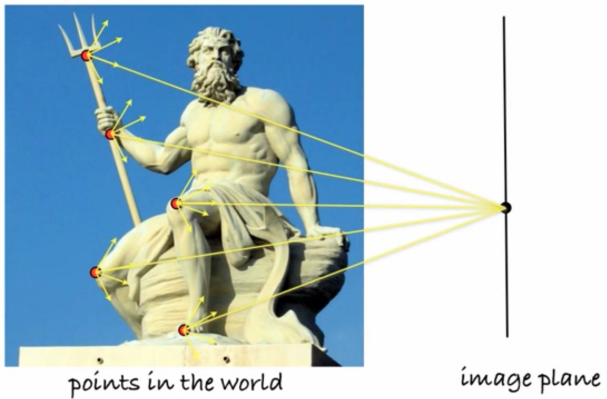
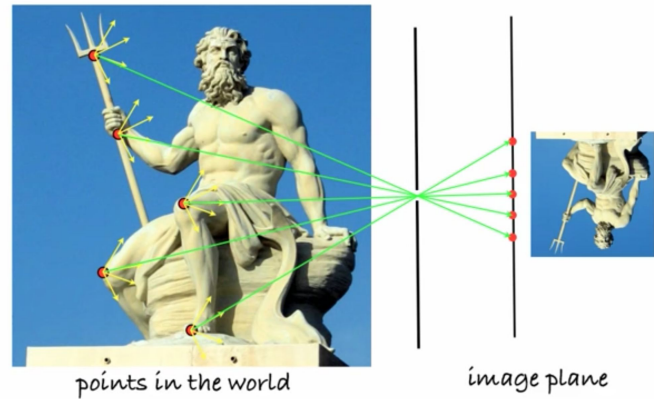


Image formation



Pin hole images



a university for the real world

CRICOS No. 00213J

The world's largest pin hole camera



<http://www.legacyphotoproject.com>

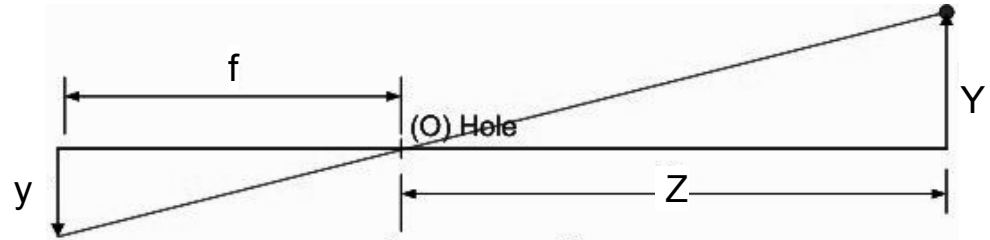
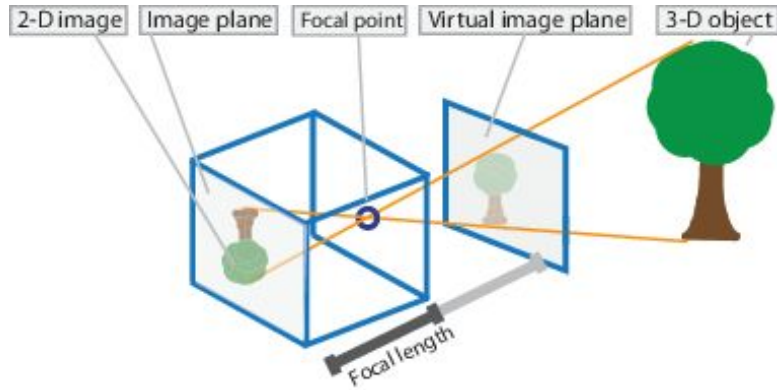


a university for the real world

CRICOS No. 00213J

Pinhole Camera Model

- Simplest model of imaging process

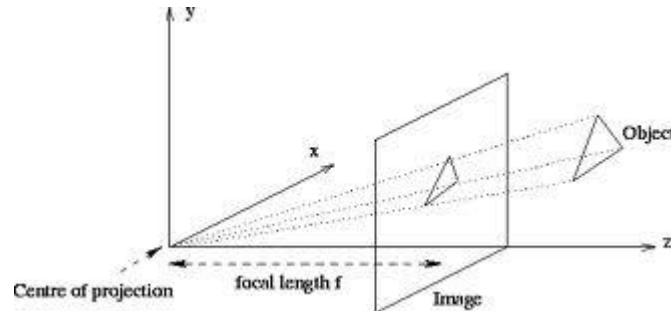
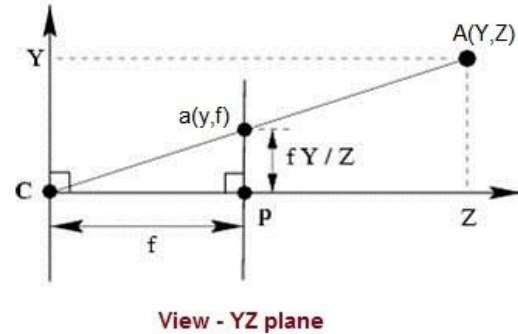
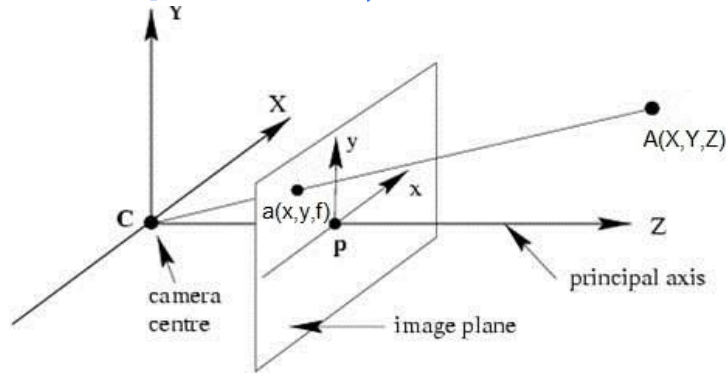


$$\frac{y}{f} = \frac{Y}{Z} \longrightarrow y = \frac{fY}{Z}$$

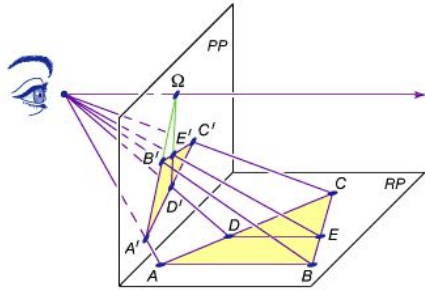
$$\frac{x}{f} = \frac{X}{Z} \longrightarrow x = \frac{fX}{Z}$$

Ref: “A Flexible New Technique for Camera Calibration”, Zhengyou Zhang, Technical Report, August 13, 2008

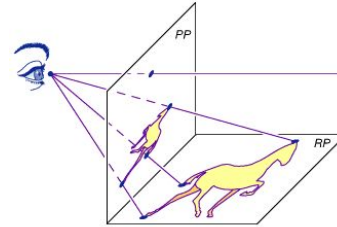
Pinhole Camera Model- Another Representation (Use of virtual planes)



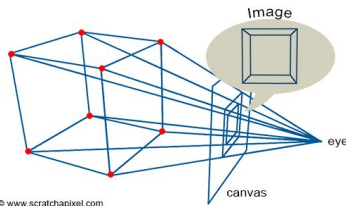
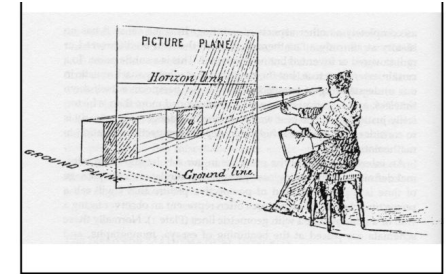
How human eye perceives the world?



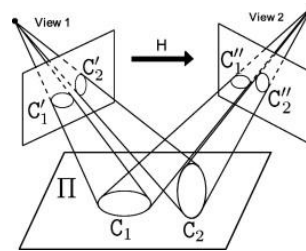
© 2002 Encyclopædia Britannica, Inc.



© 2002 Encyclopædia Britannica, Inc.

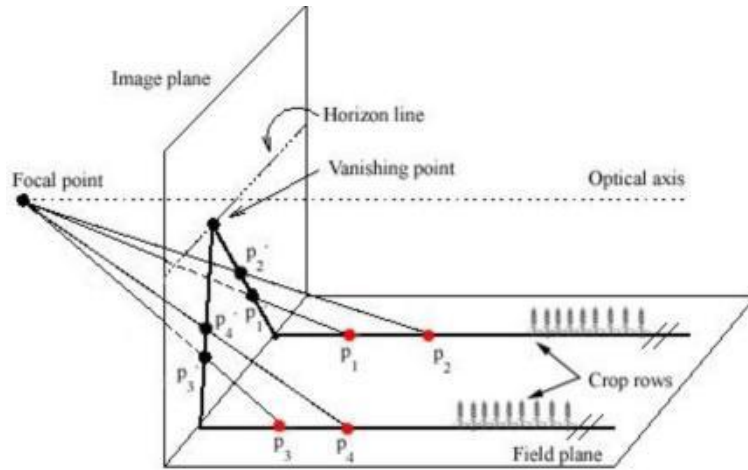


© www.scratchapixel.com

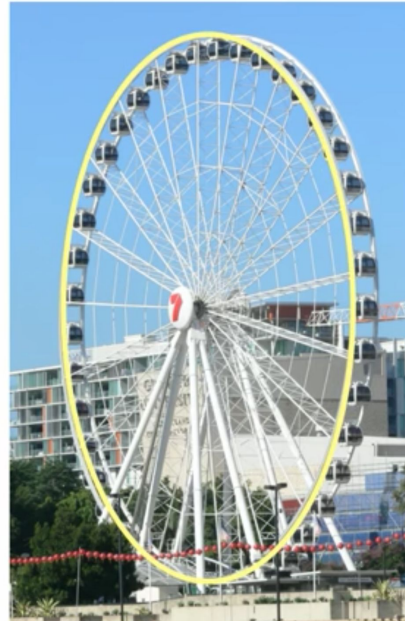


What is lost in Projective transforms?

Parallel lines meet



Distortion of Shapes



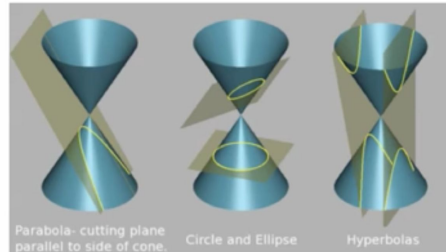
Construction of Perspective images

<https://www.geogebra.org/classic>

What is preserved in Projective Transform?

Perspective projection

- Maps
 - Lines \rightarrow lines
 - parallel lines not necessarily parallel
 - angles are not preserved
 - Conics \rightarrow conics



Significance of Previous statement



Every family of parallel lines should be changed into a family of intersecting lines and as a consequence a few of the intersecting lines may change into parallel lines

All previous transforms were dealing with points and hence were mathematically easy to derive.

This kind of projective transform actually requires one to analyse lines and hence mathematically involved?

Can we find an easier way out?

Ideal points and line at infinity

Let us assume that all parallel lines actually meet at an infinite point called the “ideal point”.

All the ideal points lie on a line called the “line at infinity”

Then our transformation should focus on transforming the “ideal points” into real points

This gives rise to the two axioms of projective geometry

Axiom 1: Any two distinct points are incident with exactly one line

Axiom 2: Any two distinct lines are incident with at least one point

Supreme power that Projective transform bestows

Map an ideal point into a real point (Thus transforms parallel lines into intersecting lines)

Map the line at infinity into a “horizon line”

How to realise ideal points in practice?

Solution : Homogeneous coordinates

$$(2, 4, 1) \text{ ----> } (2, 4)$$

$$(4, 8, 2) \text{ -----> } (2,4,1) \text{ -----> } (2,4)$$

Ideal points can be represented in the form $(x, y, 0)$. A set of parallel lines are hypothesised to meet in an ideal point

Now we should have the ability to transform $(x, y, 0)$ into a real point (a, b, c) where $c \neq 0$



Class IV: Projective transformations

$$\begin{bmatrix} \mathbf{A} & \mathbf{t} \\ \mathbf{v}^\top & v \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \\ 0 \end{pmatrix} = \begin{pmatrix} \mathbf{A} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \\ v_1 x_1 + v_2 x_2 \end{pmatrix}.$$

$$\mathbf{x}' = \mathbf{H}_P \mathbf{x} = \begin{bmatrix} \mathbf{A} & \mathbf{t} \\ \mathbf{v}^\top & v \end{bmatrix} \mathbf{x} \quad \mathbf{v} = (v_1, v_2)^\top$$

$$\begin{bmatrix} \mathbf{A} & \mathbf{t} \\ \mathbf{0}^\top & 1 \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \\ 0 \end{pmatrix} = \begin{pmatrix} \mathbf{A} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \\ 0 \end{pmatrix}.$$

8DOF (2 scale, 2 rotation, 2 translation, 2 line at infinity)

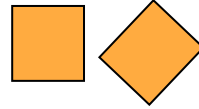
Action non-homogeneous over the plane

Invariants: cross-ratio of four points on a line (ratio of ratio)

Overview Transformations

Euclidean
3dof

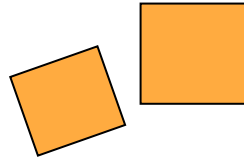
$$\begin{bmatrix} r_{11} & r_{12} & t_x \\ r_{21} & r_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix}$$



lengths, areas.

Similarity
4dof

$$\begin{bmatrix} sr_{11} & sr_{12} & t_x \\ sr_{21} & sr_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix}$$

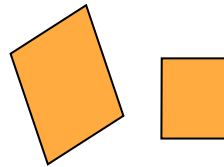


Ratios of lengths, angles.

The circular points I,J

Affine
6dof

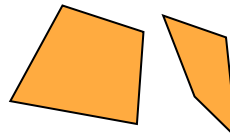
$$\begin{bmatrix} a_{11} & a_{12} & t_x \\ a_{21} & a_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix}$$



Parallelism, ratio of areas, ratio of lengths on parallel lines (e.g midpoints), linear combinations of vectors (centroids). **The line at infinity l_∞**

Projective
8dof

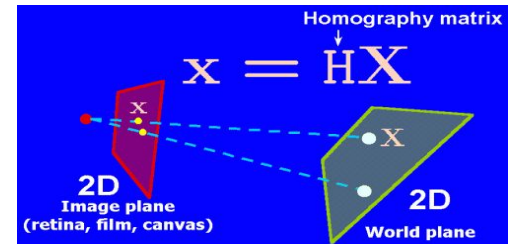
$$\begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix}$$



Concurrency, collinearity, order of contact (intersection, tangency, inflection, etc.), cross ratio



Projective Distortion



Hands - On

Demo

1. Geometric transforms ipython notebook
2. Projective transform image plane change
- 3.

Removing projective distortion



select four points in a plane with known coordinates

$$x' = \frac{x'_1}{x'_3} = \frac{h_{11}x + h_{12}y + h_{13}}{h_{31}x + h_{32}y + h_{33}}$$

$$y' = \frac{x'_2}{x'_3} = \frac{h_{21}x + h_{22}y + h_{23}}{h_{31}x + h_{32}y + h_{33}}$$

$$x'(h_{31}x + h_{32}y + h_{33}) = h_{11}x + h_{12}y + h_{13}$$

$$y'(h_{31}x + h_{32}y + h_{33}) = h_{21}x + h_{22}y + h_{23}$$

(linear in h_{ij})

(2 constraints/point, 8DOF \Rightarrow 4 points needed)

Remark: no calibration at all necessary

Application - License plate recognition





Image Stitching





Image Stitching

1. Repeat for all images
 - a. Detect 2D Features points in both the images
 - b. Find match between 2D features of both images
 - c. Find Homography between two images using matched points
(Do RANSAC to reduce effect of outliers)
 - d. Warp the image w.r.t any particular one image
 - e. Merge the images
1. Color Blend the final image

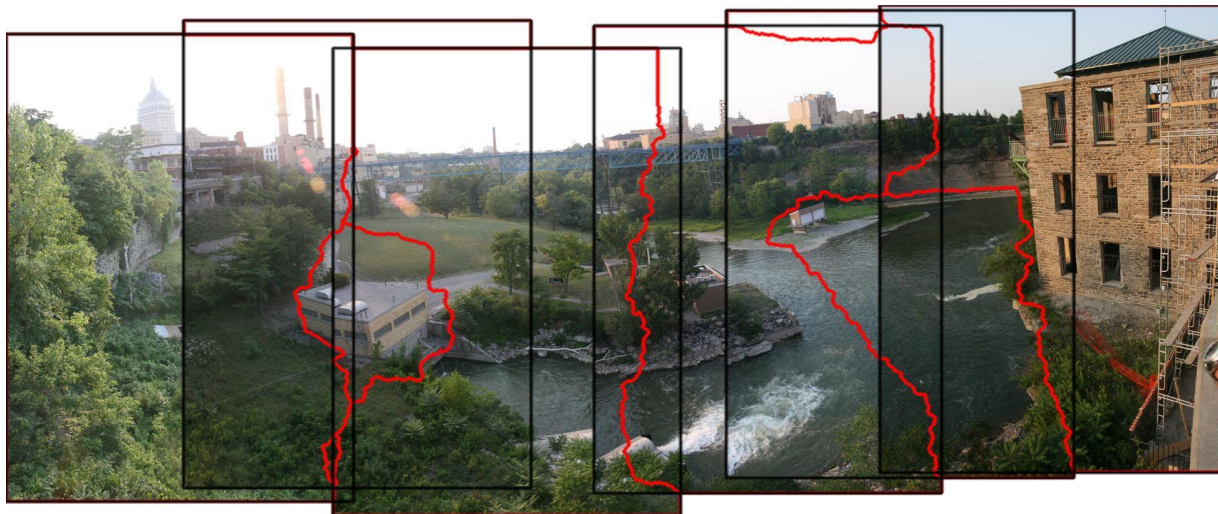


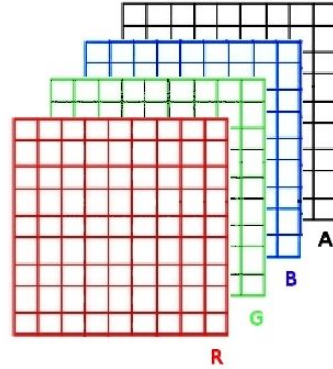


Image Representation

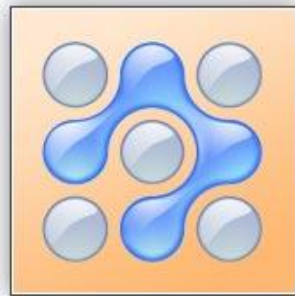
0.2235	0.1294	Blue	0.4196	0.2902	0.2902	0.4824	0.2588	0.2588	0.2588
0.5804	0.2902	0.0627	0.2902	0.2902	0.4824	0.2588	0.2588	0.2588	0.2588
0.5804	0.0627	0.0627	0.0627	0.2235	0.2588	0.2588	0.2588	0.2588	0.2588
0.5176	0.1922	0.0627	Green	0.1922	0.2588	0.2588	0.2588	0.2588	0.2588
0.5176	0.1294	0.1608	0.1294	0.1294	0.2588	0.2588	0.2588	0.2588	0.2588
0.5176	0.1608	0.0627	0.1608	0.1922	0.2588	0.2588	0.2588	0.2588	0.2588
0.5490	0.2235	0.5490	Red	0.7412	0.7765	0.7765	0.7765	0.7765	0.7765
0.490	0.3882	0.5176	0.5804	0.5804	0.7765	0.7765	0.7765	0.7765	0.7765
0.2588	0.2902	0.2588	0.2235	0.4824	0.2235	0.2235	0.2235	0.2235	0.2235
0.2235	0.1608	0.2588	0.2588	0.1608	0.2588	0.2588	0.2588	0.2588	0.2588
0.1608	0.2588	0.2588	0.2588	0.2588	0.2588	0.2588	0.2588	0.2588	0.2588



Transparent Images

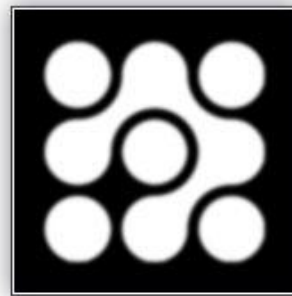


Use of Alpha Channel to create Transparent Image

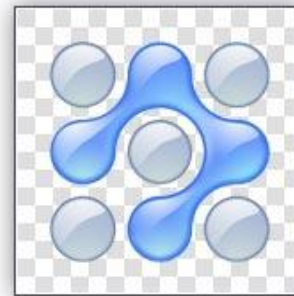


Original Image
RGB - 24 bpp

+

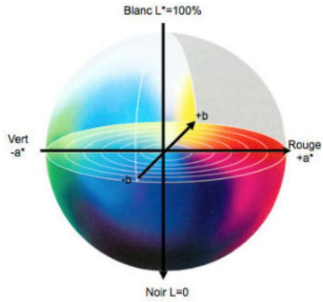
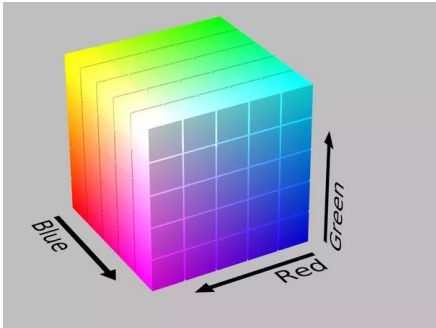
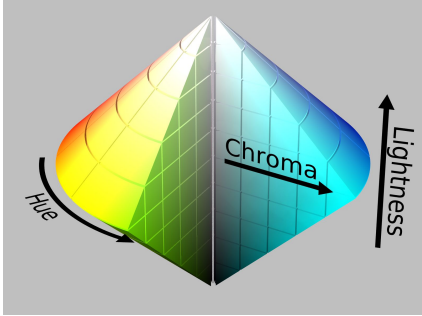
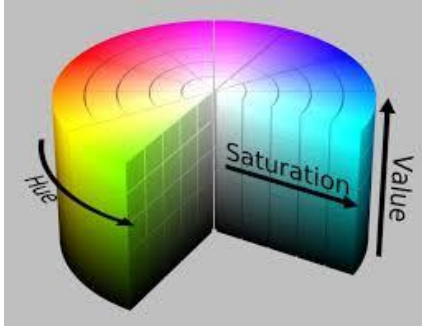


Alpha Channel
A - 8 bpp



Transparent Image
RGBA - 32 bpp

Different Color Spaces



Lab color space

Basic terminologies in images

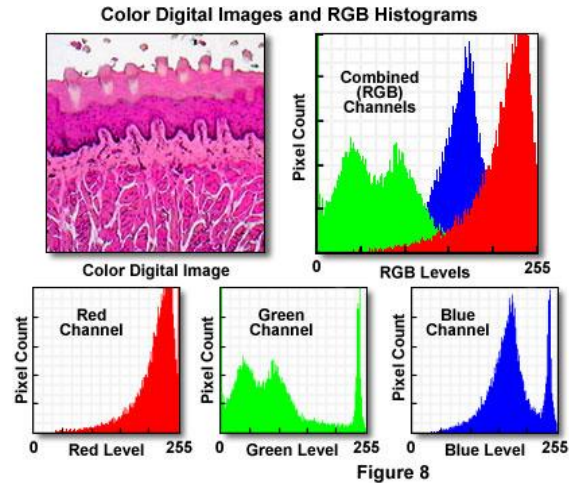
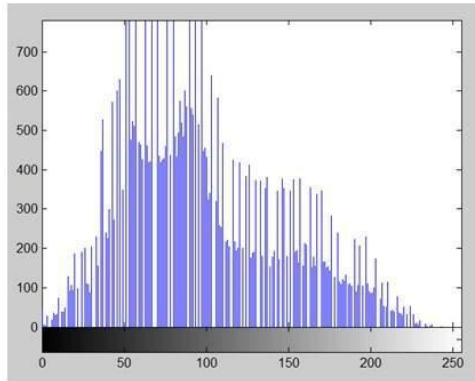
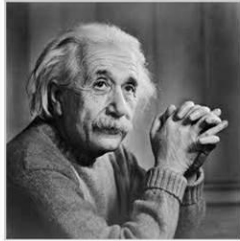
Brightness - Determines the intensity of the color presented by a pixel in a color image

Contrast - The difference between the lightest and darkest regions of an image.

Saturation - The component of the HSV color model that controls the amount of white mixed into the hue.

Histograms

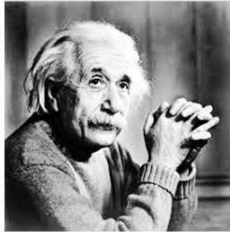
Histogram is a measure of the frequency of occurrence of each pixel value in the image



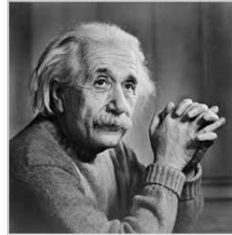
Histogram equalisation

Histogram equalization is a technique for adjusting image intensities to enhance contrast. A very good preprocessing step to reduce the effect input image lightness variations

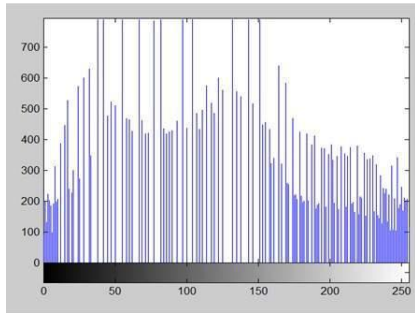
New Image



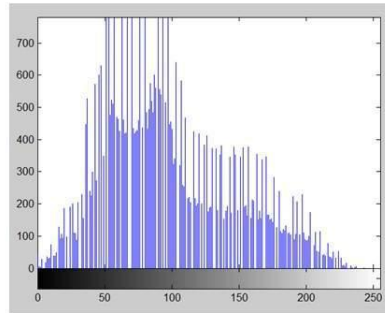
Old image



New Histogram

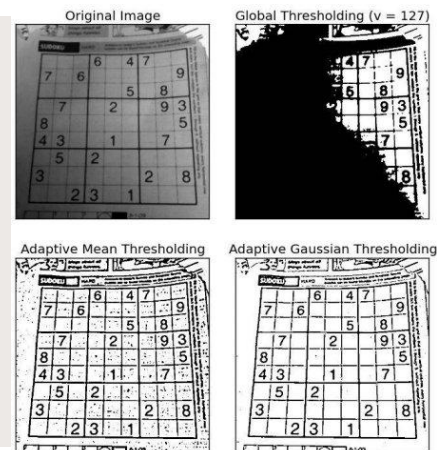
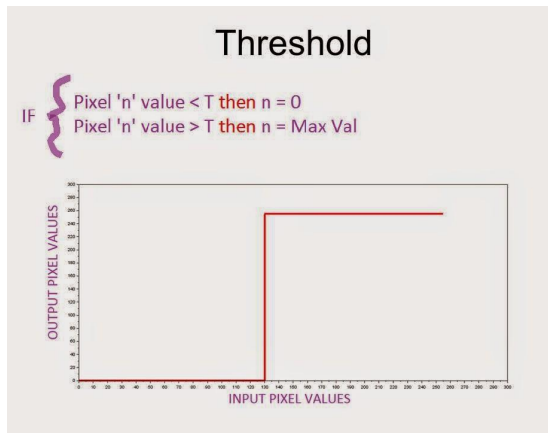


Old Histogram



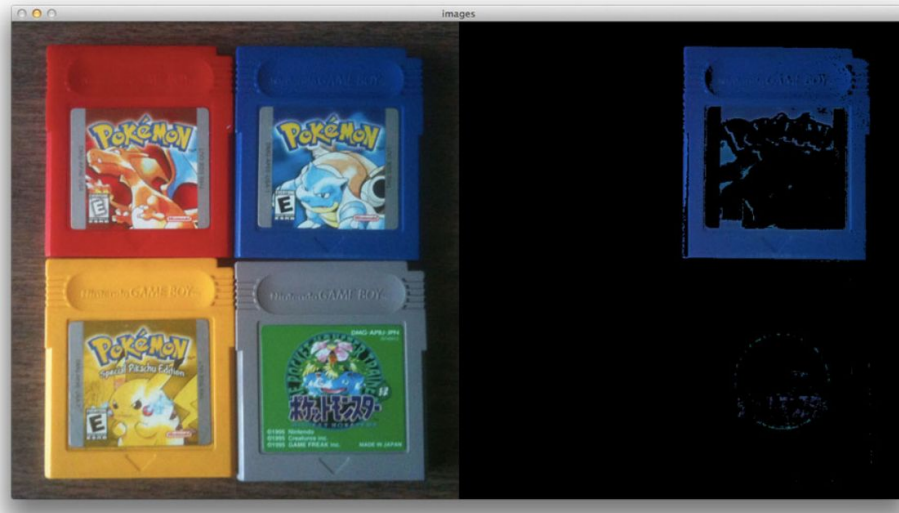
Thresholding

1. Converts a grayscale image to a binary image by thresholding all values below the threshold (a max limit) to zero and all values above the threshold (a max limit) to the highest value
 - a. Binary
 - b. Binary inverted
 - c. Adaptive



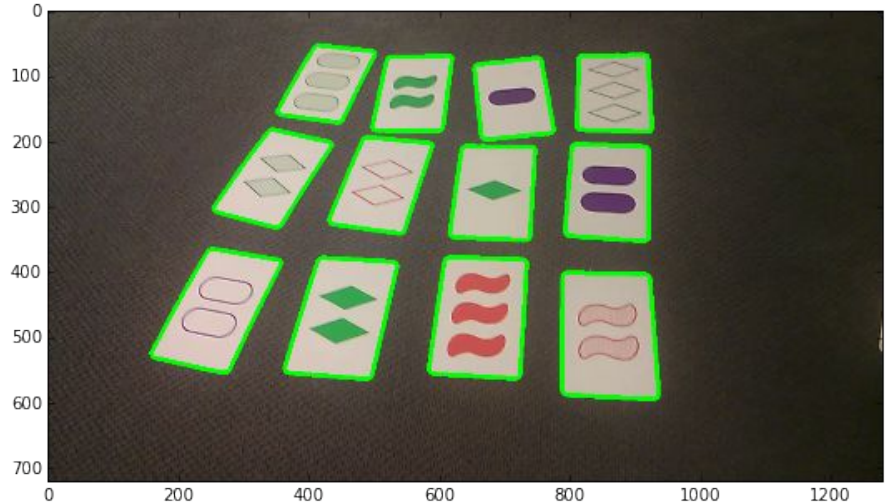
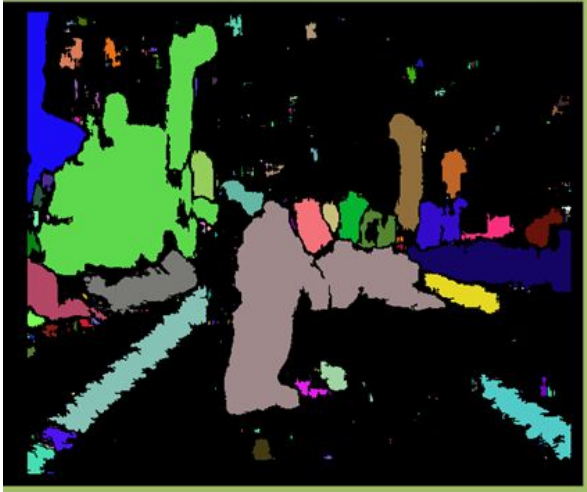
Masking a color in an image

1. Find the desired range for the color in HSV
2. Get / Convert image to HSV format.
3. Keep all pixels within the range as it is and replace other pixels by zero



Contour detection

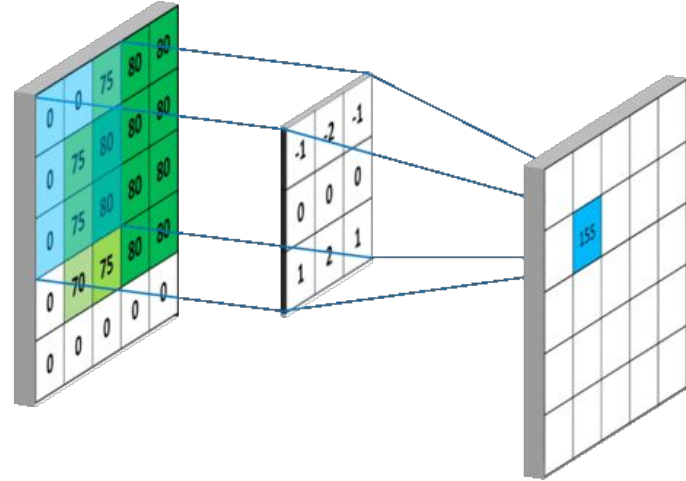
1. Contour is a curve joining all points in a boundary.
2. Contour are useful for shape analysis, object detection and analysis



Convolution

0 ₂	0 ₀	0 ₁	0	0	0	0
0 ₁	2 ₀	2 ₀	3	3	3	0
0 ₀	0 ₁	1 ₁	3	0	3	0
0	2	3	0	1	3	0
0	3	3	2	1	2	0
0	3	3	0	2	3	0
0	0	0	0	0	0	0

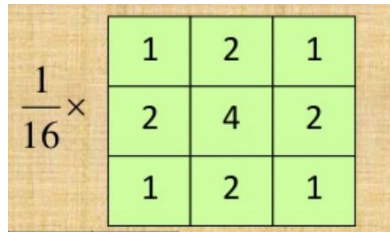
1	6	5
7	10	9
7	10	8



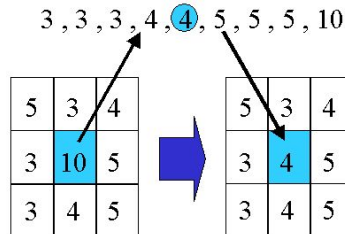
Smoothing

1. Smoothing removes noise
2. Different types of smoothing are available
 - a. Gaussian
 - b. Median
 - c. Average
 - d. Weighted average

Weighted average filter



Median filter



Box filter

$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

Gaussian filter

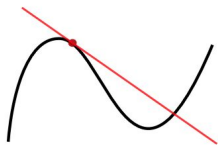
$\frac{1}{273}$

1	4	7	4	1
4	16	26	16	4
7	26	41	26	7
4	16	26	16	4
1	4	7	4	1

How to find gradients?

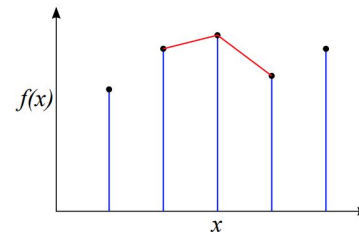
Given a function $f : \mathbb{R} \rightarrow \mathbb{R}$, its derivative is defined as

$$\frac{df}{dx}(x) = \lim_{\epsilon \rightarrow 0} \frac{f(x + \epsilon) - f(x)}{\epsilon}$$



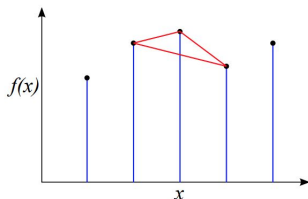
Derivative of f is the slope of the tangent to the graph of f

Derivatives of Discrete Functions



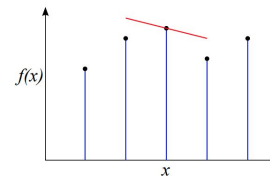
Slopes (derivatives) don't match on left and right

Derivatives of Discrete Functions



Instead take the average of the two (or secant)

Derivatives of Discrete Functions



Instead take the average of the two (or secant)

What is an edge?

What is an Edge?

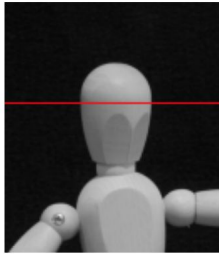
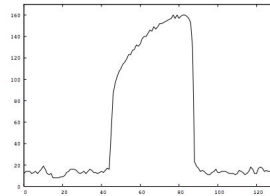


Image Value vs X-Position



An abrupt transition in intensity between two regions

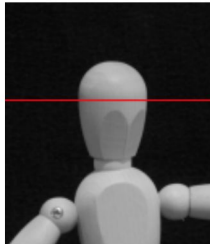


Image X-Derivative vs X-Position

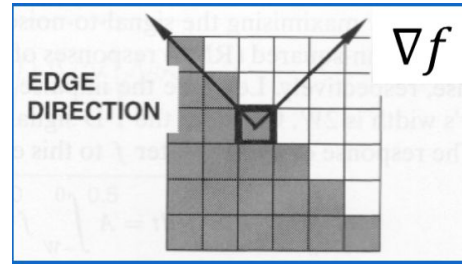
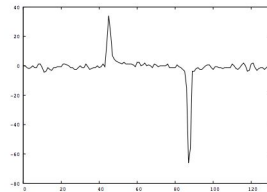


Image derivatives are high (or low) at edges

How to find gradients?

Finite Differences

Forward Difference

$$\Delta_+ f(x) = f(x+1) - f(x) \quad \text{right slope}$$

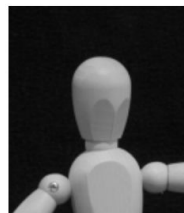
Backward Difference

$$\Delta_- f(x) = f(x) - f(x-1) \quad \text{left slope}$$

Central Difference

$$\Delta f(x) = \frac{1}{2}(f(x+1) - f(x-1)) \quad \text{average slope}$$

x-derivative using central difference:



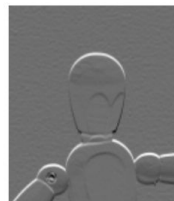
$$* \begin{bmatrix} \frac{1}{2} & 0 & -\frac{1}{2} \end{bmatrix} =$$



y-derivative using central difference:



$$* \begin{bmatrix} 0.5 \\ 0 \\ -0.5 \end{bmatrix} =$$



How to find gradients?

Finite Differences

Forward Difference

$$\Delta_+ f(x) = f(x+1) - f(x) \quad \text{right slope}$$

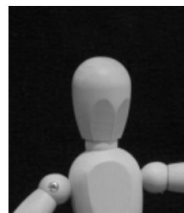
Backward Difference

$$\Delta_- f(x) = f(x) - f(x-1) \quad \text{left slope}$$

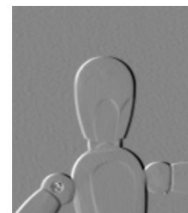
Central Difference

$$\Delta f(x) = \frac{1}{2}(f(x+1) - f(x-1)) \quad \text{average slope}$$

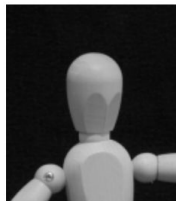
x-derivative using central difference:



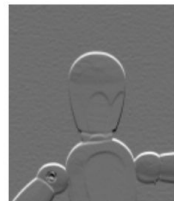
$$* \begin{bmatrix} \frac{1}{2} & 0 & -\frac{1}{2} \end{bmatrix} =$$



y-derivative using central difference:



$$* \begin{bmatrix} 0.5 \\ 0 \\ -0.5 \end{bmatrix} =$$



But let's smooth our gradients a little.

Prewitt

$$H_x^P = \frac{1}{3} \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} * [0.5 \ 0 \ -0.5] = \frac{1}{6} \begin{bmatrix} 1 & 0 & -1 \\ 1 & 0 & -1 \\ 1 & 0 & -1 \end{bmatrix}$$

What is H_y^P in Prewitt?

Sobel

$$H_x^S = \frac{1}{4} \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} * [0.5 \ 0 \ -0.5] = \frac{1}{8} \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix}$$

$$H_y^S = \frac{1}{4} [1 \ 2 \ 1] * \begin{bmatrix} 0.5 \\ 0 \\ -0.5 \end{bmatrix} = \frac{1}{8} \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$$

Now we are set to calculate our gradients?

Derivatives in 2 Dimensions

Given function

$$f(x, y)$$

Gradient vector

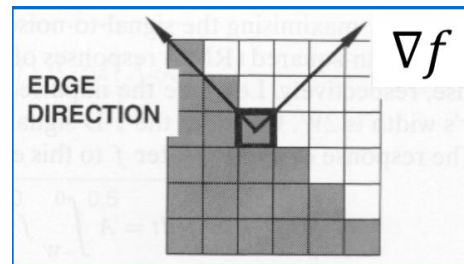
$$\nabla f(x, y) = \begin{bmatrix} \frac{\partial f(x, y)}{\partial x} \\ \frac{\partial f(x, y)}{\partial y} \end{bmatrix} = \begin{bmatrix} f_x \\ f_y \end{bmatrix}$$

Gradient magnitude

$$|\nabla f(x, y)| = \sqrt{f_x^2 + f_y^2}$$

Gradient direction

$$\theta = \tan^{-1} \frac{f_x}{f_y}$$



Morphological operations

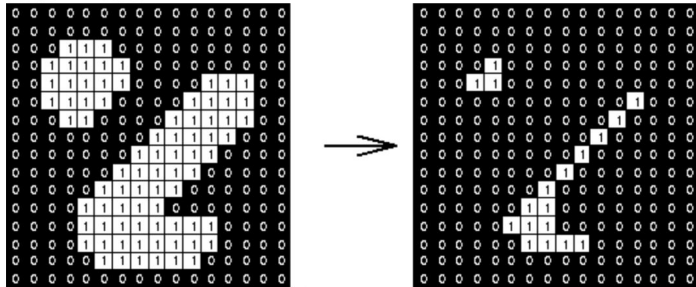
A collection of non-linear operations that is related to shape or morphology of features in an image.

There are four kinds of morphological operations in an image

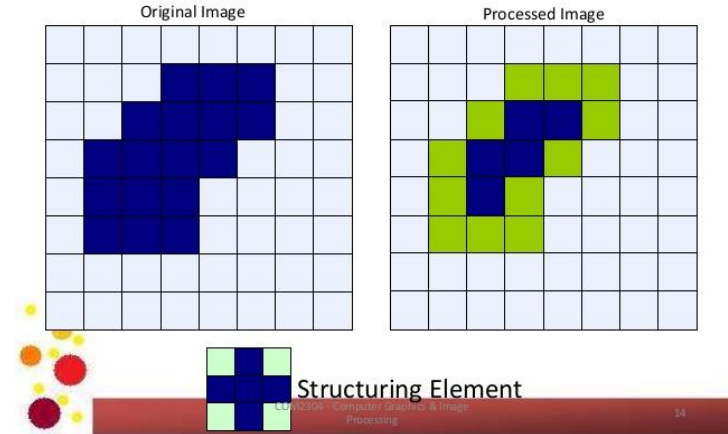
1. Erosion
2. Dilation
3. Opening
4. Closing

Erosion

The structuring element is superimposed on every pixel in the image. If all pixels within the structuring element are foreground, then the input pixel retains its value. If any pixel within the structuring element is background, then the input pixel is replaced by background.



Erosion: a 3x3 square structuring element
(www.cs.princeton.edu/~pshilane/class/mosaic/).



Greyscale image

Binary image by thresholding

Erosion: a 2x2 square structuring element

Dilation

The structuring element is superimposed on every pixel in the image. If all pixels within the structuring element are background, then the input pixel retains its value. If any pixel within the structuring element is foreground, then the input pixel is replaced by foreground.

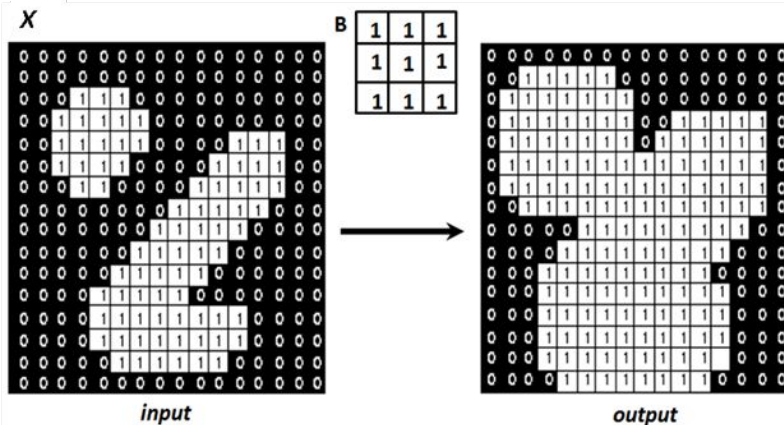
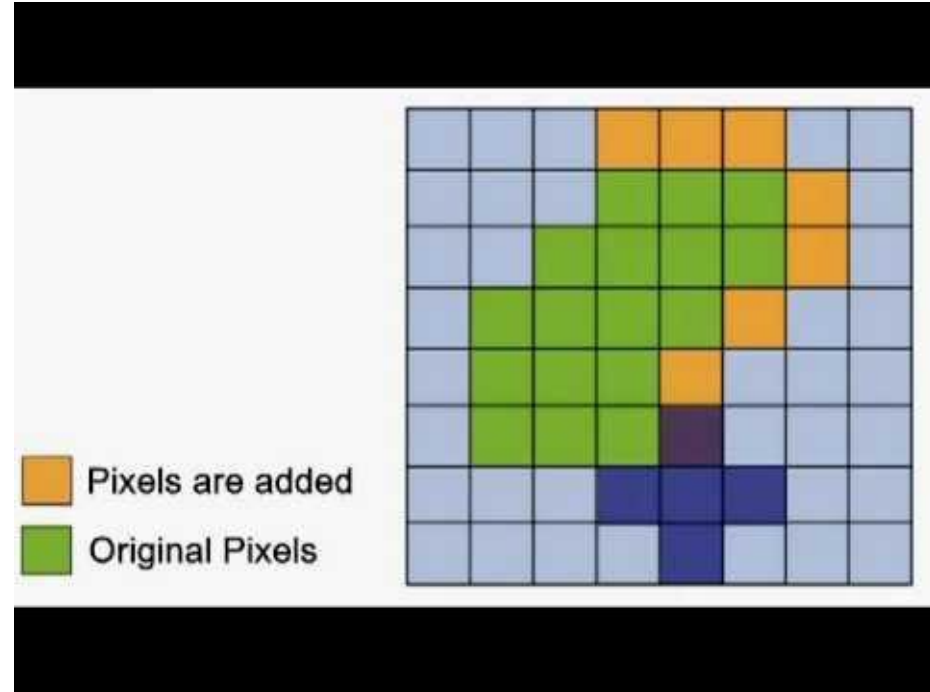
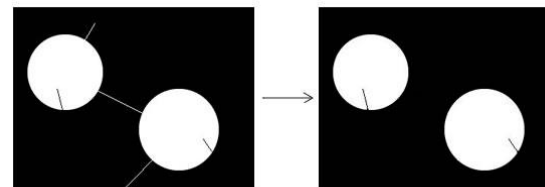


Figure 1. Effect of dilation using a 3X3 square structural element B.

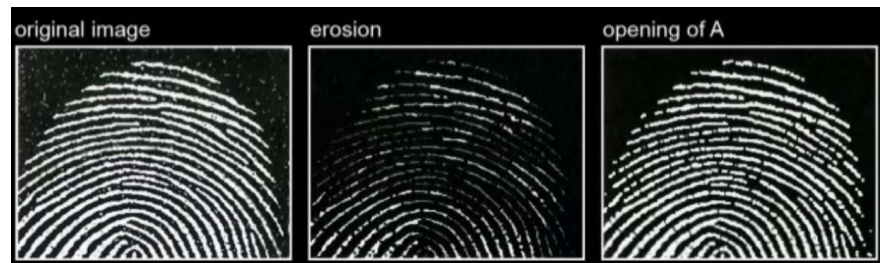
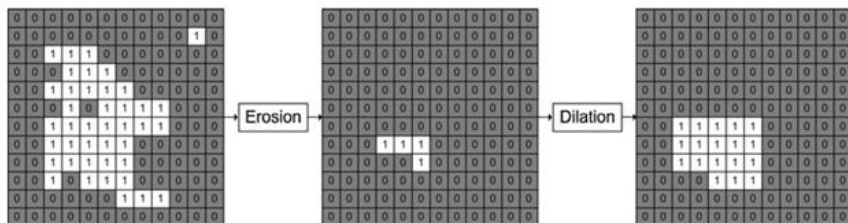


Opening

1. Erosion followed by dilations
2. Smooths contours and eliminates protrusions



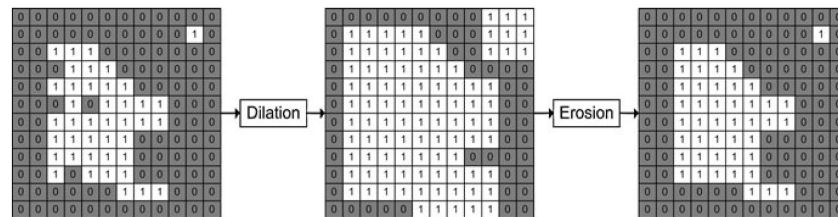
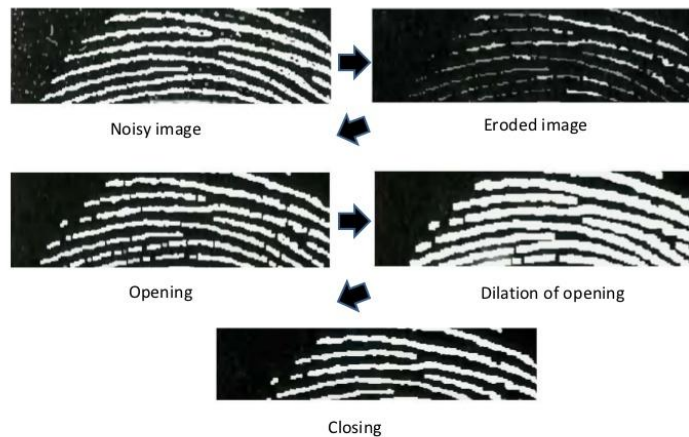
Morphological Opening



Closing

1. First dilate and then erode
2. This smooths sections of contours, fuses narrow breaks.

A morphological filter



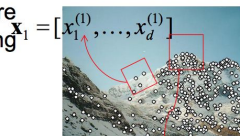
Establishing correspondence

Local features: main components

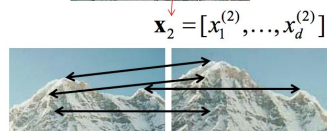
- 1) Detection: Identify the interest points



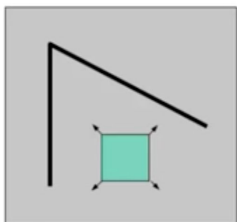
- 2) Description :Extract feature vector descriptor surrounding each interest point.



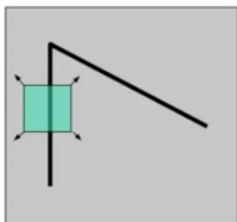
- 3) Matching: Determine correspondence between descriptors in two views



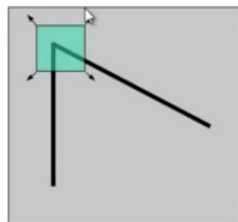
Properties of an interest point



“flat” region:
no change in
all directions



“edge”:
no change along
the edge direction



“corner”:
significant change
in all directions

Properties of Interest Point Detectors

- Detect all (or most) true interest points
- No false interest points
- Well localized.
- Robust with respect to noise.
- Efficient detection

Some of the commonly used interest point detector

1. Harris point detector
2. Shi-Tomasi detector
3. Features from accelerated Segment Test (FAST)

Feature descriptor

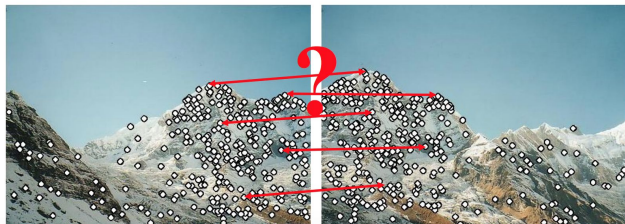
Feature descriptors

We know how to detect good interest points

Next question:

How to match image regions around interest points?

Answer: Need feature descriptors



Few of the desired properties of key point descriptors

1. Rotation Invariance
2. Scale Invariance

Some of the commonly used interest point descriptors

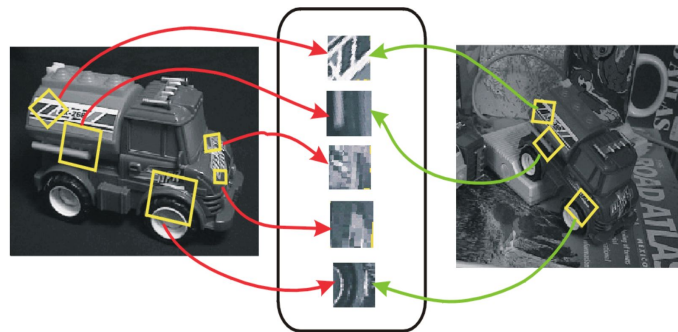
1. Scale Invariant Feature Transform (SIFT)
2. Speeded Up Robust Features (SURF)
3. Binary Robust Independent Elementary Features (BRIEF)
4. Oriented FAST and rotated BRIEF (ORB)

Feature matchers

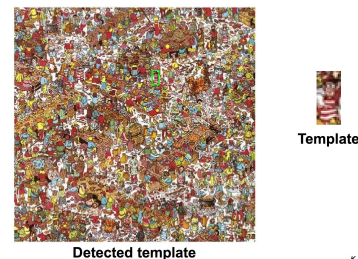
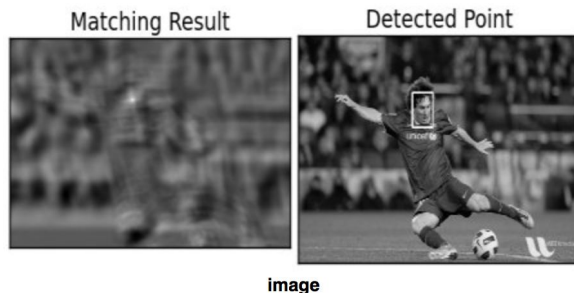


Process of feature matching is as follows

1. Find the distance of each descriptors in the given image to every descriptor in the template image
2. Find the best match for each descriptor and its corresponding distance.
3. Remove the pairs which are below a threshold
4. Still a few bad pairs will be left. To work around this, use louse's ratio. Find the ratio of the distance of the first best match to the second best match. If the ratio > 0.6 , eliminate the pair



Template matching



1. Given a template image and an input image to search for, search where exactly is the template image located in the input image.
2. Only translation and intensity variations are allowed. No other geometric operation is allowed
3. But how do we measure similarity between two patches of images.
4. Searching through the whole image might be time consuming?

How do we measure similarity between two patches

The similarity between two patches g_1 and g_2 can be found by the following ways

1. Sum of squared differences (SSD)

$$\sum (g_2(m) - g_1(m))^2$$

But these provide no invariance
against brightness and contrast !!!!

1. Sum of absolute difference (SAD)

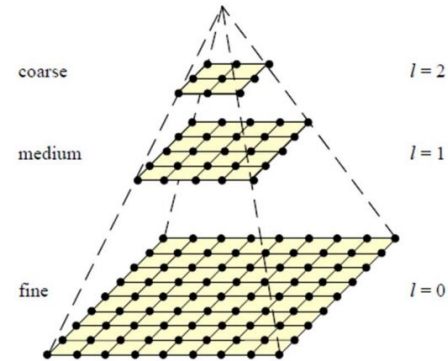
$$\sum |g_2(m) - g_1(m)|$$

1. Maximum of difference

$$\text{Max} = \max_m |g_2(m) - g_1(m)|$$

Solution - Normalised Correlation (Convolution) And Image pyramids

$$\frac{\sum_{i=-N}^N (F(i)I(x+i))}{\sqrt{\sum_{i=-N}^N (I(x+i))^2} \sqrt{\sum_{i=-N}^N (F(i))^2}} .$$



Day 1 - FN 2 (Basic Image Processing)

A photograph of a lunch break. In the foreground, a sandwich is cut into two halves and placed in a clear plastic container. The sandwich is filled with lettuce, tomato, and meat. To the right, a bowl contains shredded white cheese. In the background, there is a whole apple and a glass of orange juice. The image is overlaid with a semi-transparent dark grey filter.

Lunch Break

Feed yourselves well to feed the Machines more..!

Day 1 - AN 1 (Basic Image Processing)



15

Coffee

Est. 2016

Day 1 - AN 2 (Machine Learning Intro)

Session II (ML)

- Object Classification
- Classification: Simple Version
- Features
- Classification: CV based
- Linear classifiers
- SVM



What is Machine Learning?

- Arthur Samuel (1959): Field of study that gives computers the ability to learn without being explicitly programmed
- A well-posed learning problem (1998): A computer program is said to learn from experience E with respect to some task T and some performance measure P , if its performance measure P , as measured by P , improves with experience E .

What is Learning?

- Set of training samples : $[x_1, \dots, x_n]$ & $[y_1, \dots, y_n]$
- Hypothesis function ($w_1x_1 + w_2x_2 = y$)
- Optimize the loss function
 - Update / change the coefficients such that they minimize the error
 - Perform till convergence
- Getting optimal values for 'w' is called learning
- Now for a new x, predict it's y

Classification

Handwritten character Classification

Q: As a Human, how do I learn to classify?

No.0 / Answer:7, Predict:[7] No.1 / Answer:2, Predict:[2] No.2 / Answer:1, Predict:[1] No.3 / Answer:0, Predict:[0] No.4 / Answer:4, Predict:[4]



No.5 / Answer:1, Predict:[1] No.6 / Answer:4, Predict:[4] No.7 / Answer:9, Predict:[9] No.8 / Answer:5, Predict:[5] No.9 / Answer:9, Predict:[9]



No.10 / Answer:0, Predict:[0] No.11 / Answer:6, Predict:[6] No.12 / Answer:9, Predict:[9] No.13 / Answer:0, Predict:[0] No.14 / Answer:1, Predict:[1]

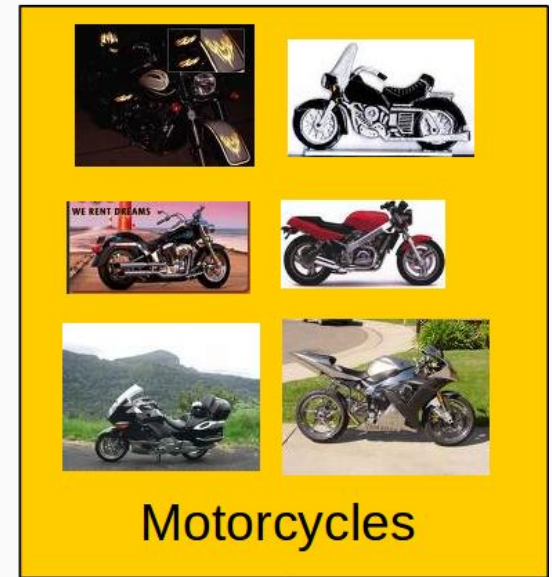
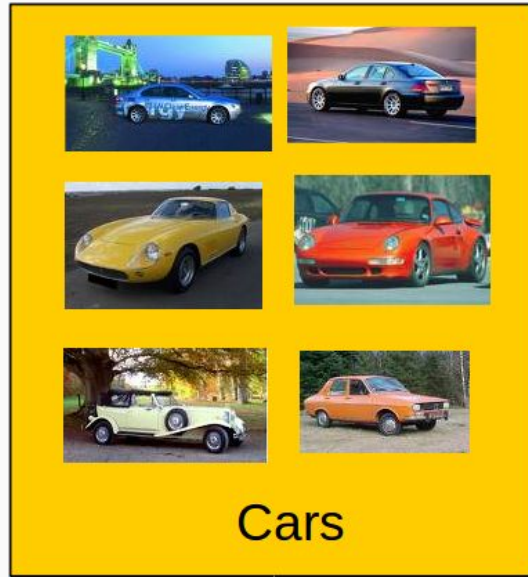


No.15 / Answer:5, Predict:[5] No.16 / Answer:9, Predict:[9] No.17 / Answer:7, Predict:[7] No.18 / Answer:3, Predict:[3] No.19 / Answer:4, Predict:[4]



Classification

1. Learn Patterns
2. Learn Rules



What is this?

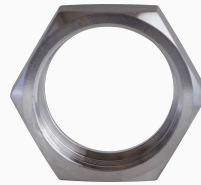


Simple Problem

How to classify a new image into any of the two categories?

Simplified Problem

- Conveyor Belt: 2D Image of Nut & Bolt
- Supervised Learning Problem
 - Train on available data
 - Test on new data



Nut



Bolt

Rules

We can classify a new image into any of these two categories by forming a set of rules.

Q: List down few rules by which we can classify

Bolts

1. Longer
2. Thinner
3. Cylindrical in shape
4. **More compact**

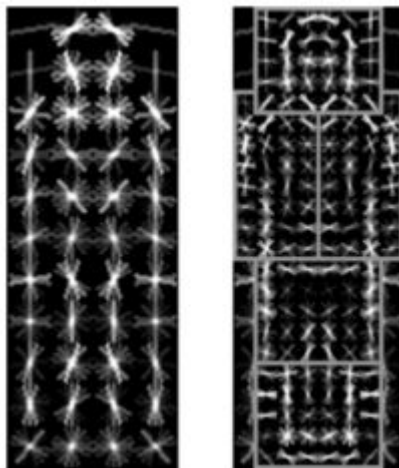
Nuts

1. **Circular**
2. Has cavity in the centre
3. Less area

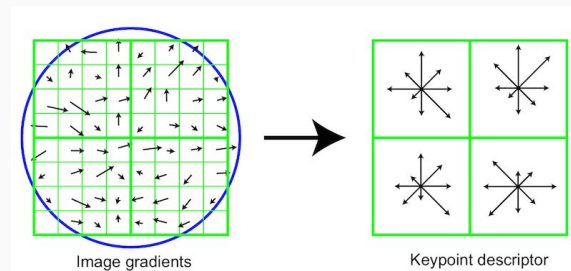
Features

In Machine Learning terms, we call them as features

Low-level Features



HoG



SIFT

Data

The graph shows the distribution of data with respect to the below features

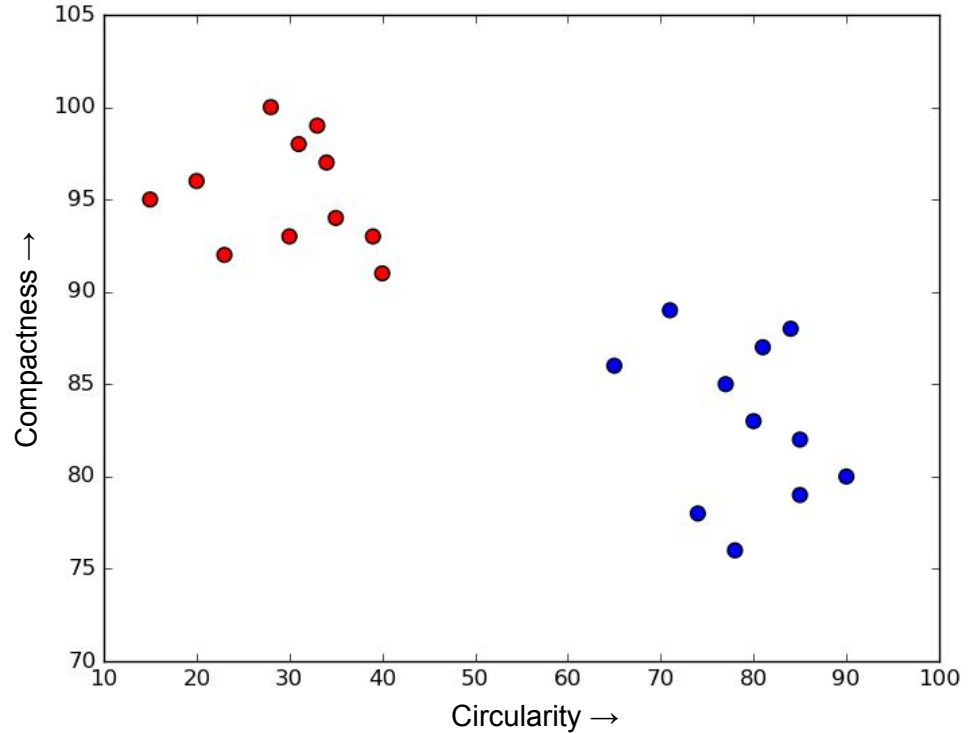
Features

1. Circularity
2. Compactness

Now, How can you classify the data?

● Bolt

● Nut



Classification

It is a simple if with two conditions

```
if Circularity > 55 and  
Compactness < 90:
```

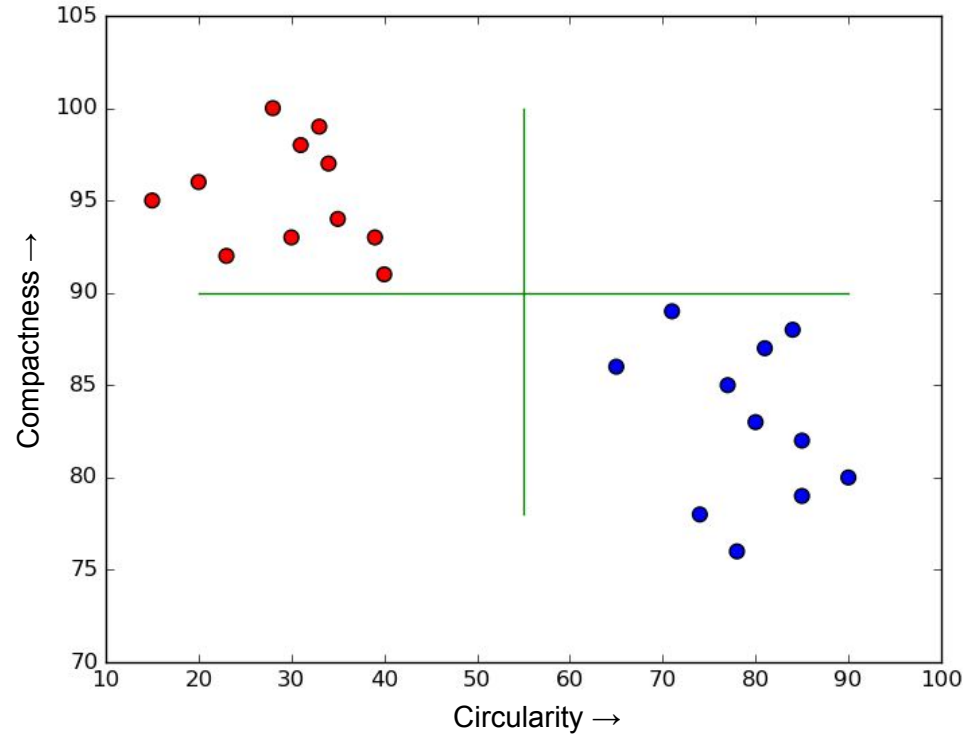
Nut

```
elif Circularity < 55 and  
Compactness > 90:
```

Bolt

● Bolt

● Nut



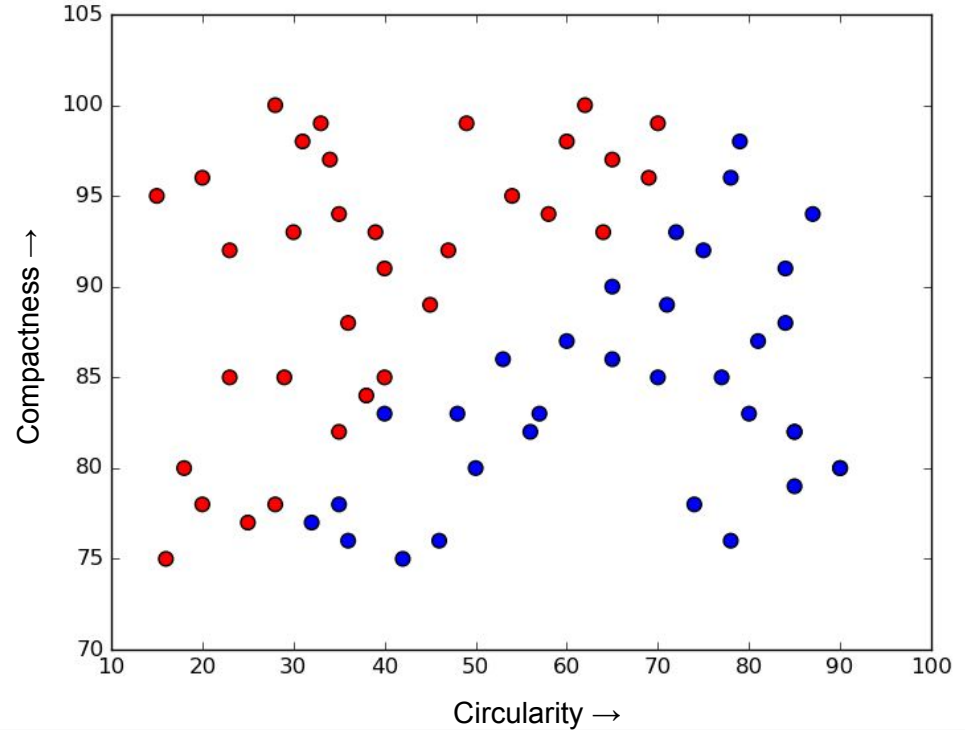
Complex Data

The Data becomes more complex to classify using a simple If Classifier

Now, How can you classify the data?

● Bolt

● Nut

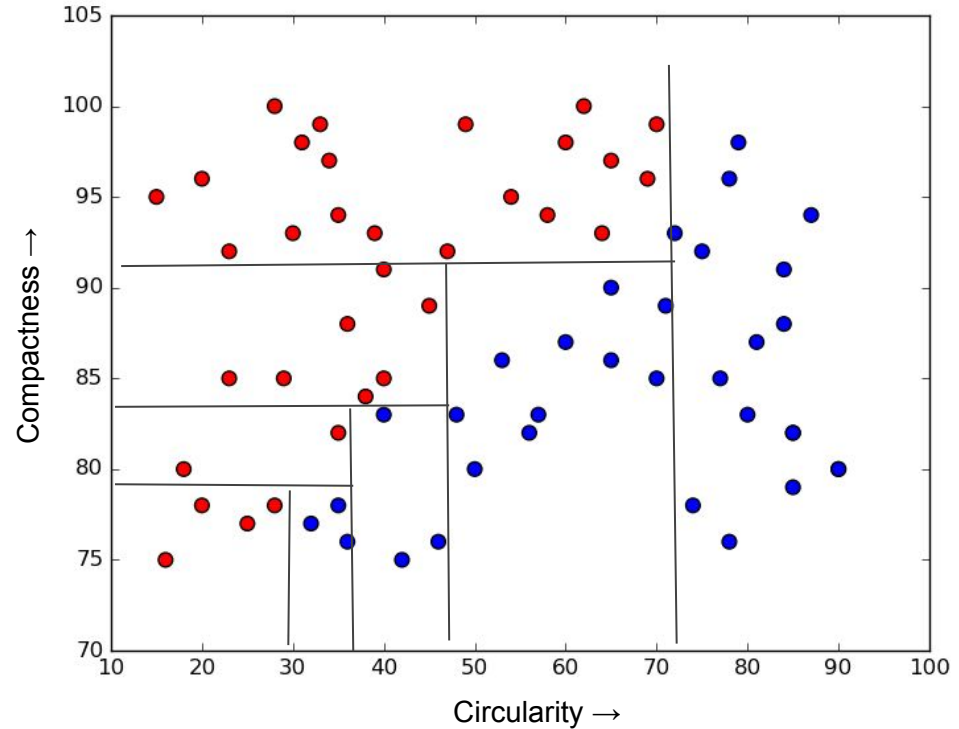


Decision Tree

- Non-parametric Supervised Learning
- Approximating data by a set of if-then-else decision rules
- Greedy algorithm
- Whitebox Model :)

• Bolt

• Nut

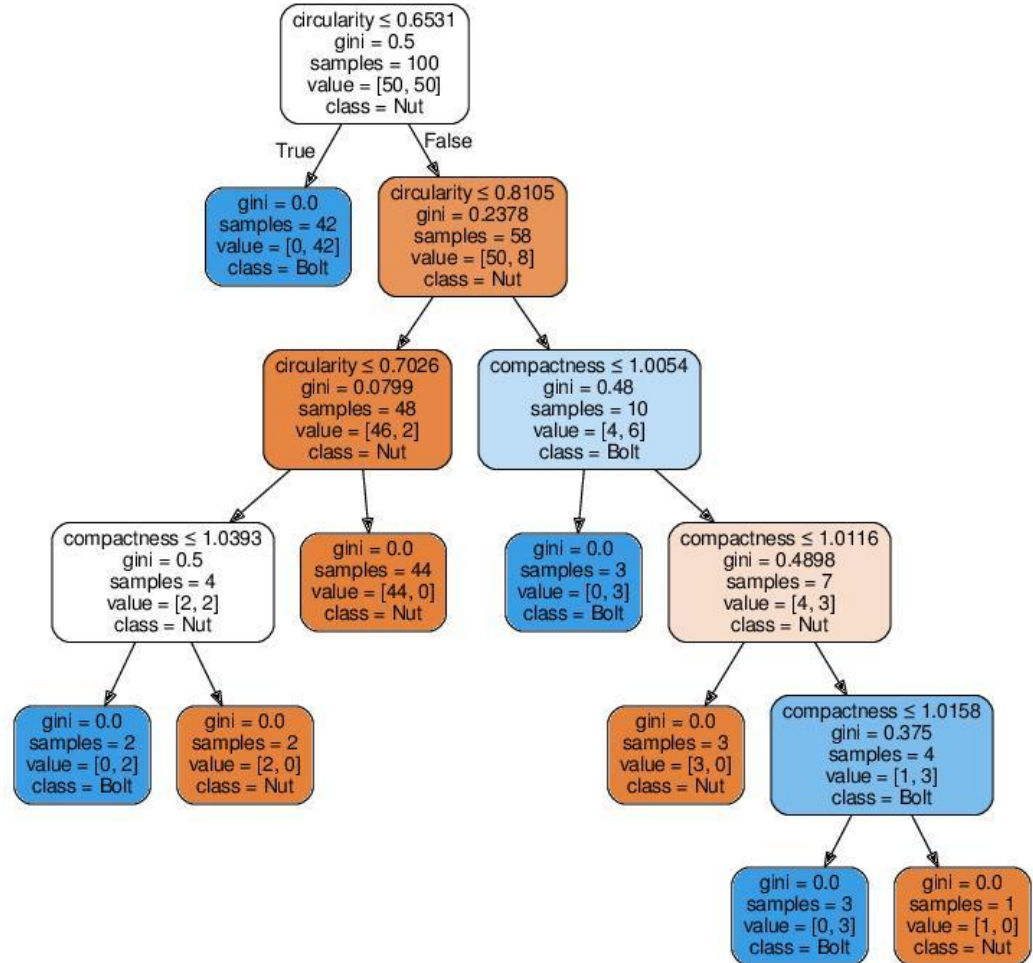


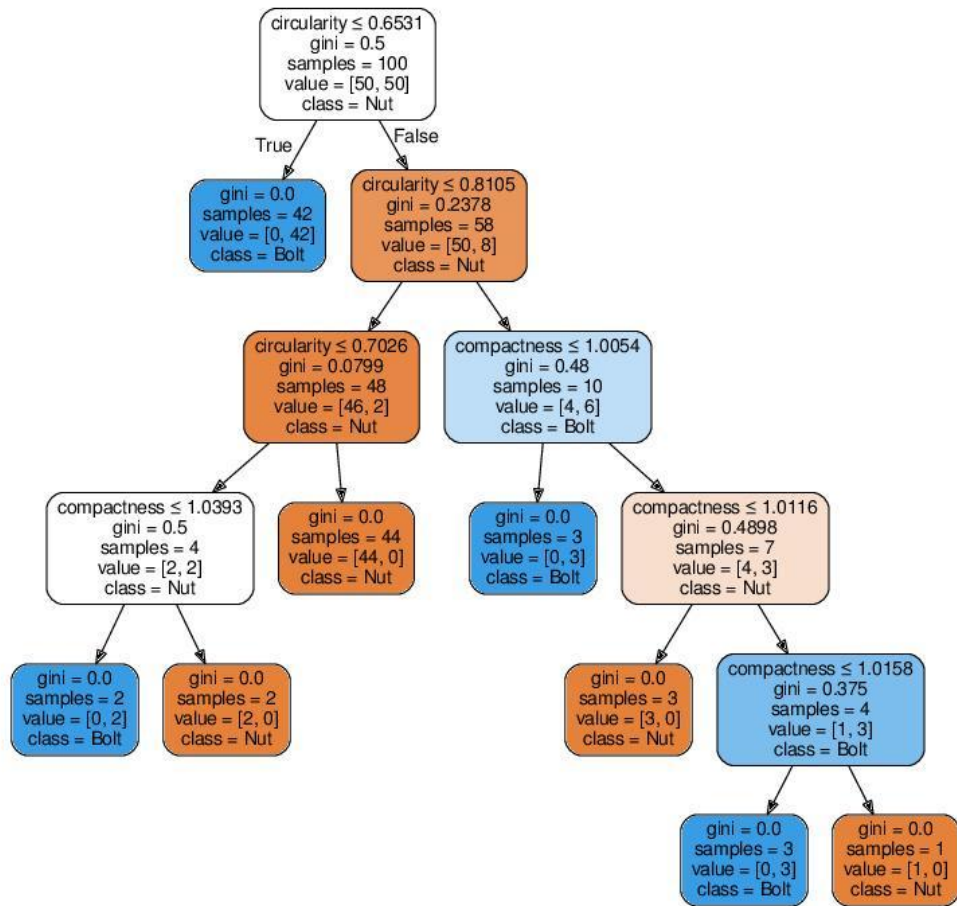
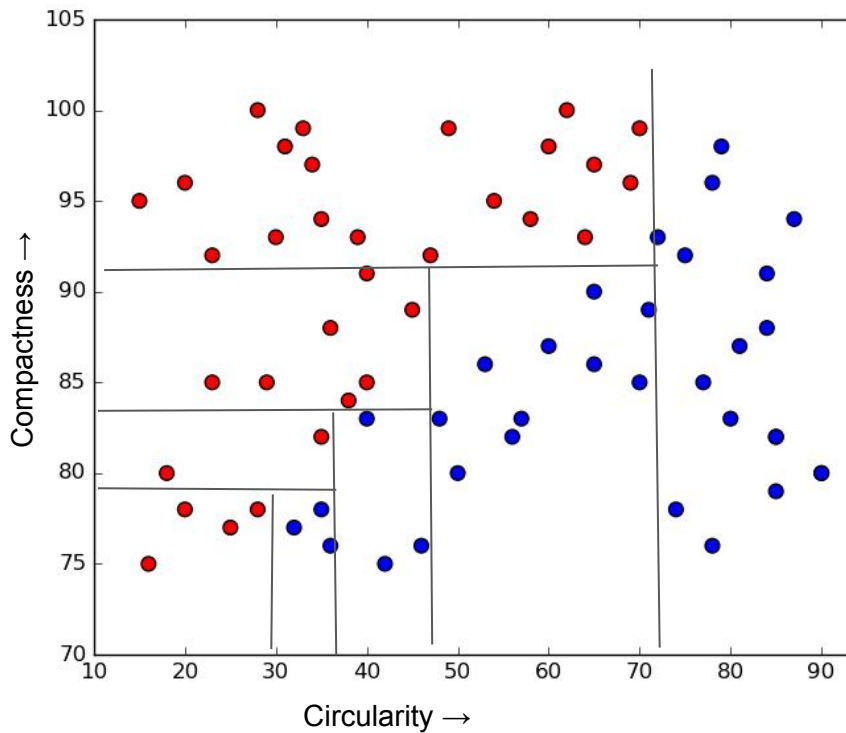
Decision Tree

Decision Tree generated for the Nuts and Bolts classification problem

Cons

- May grow huge and hard to manage / visualize
- Problem of overfitting







Let's Get our Hands Dirty

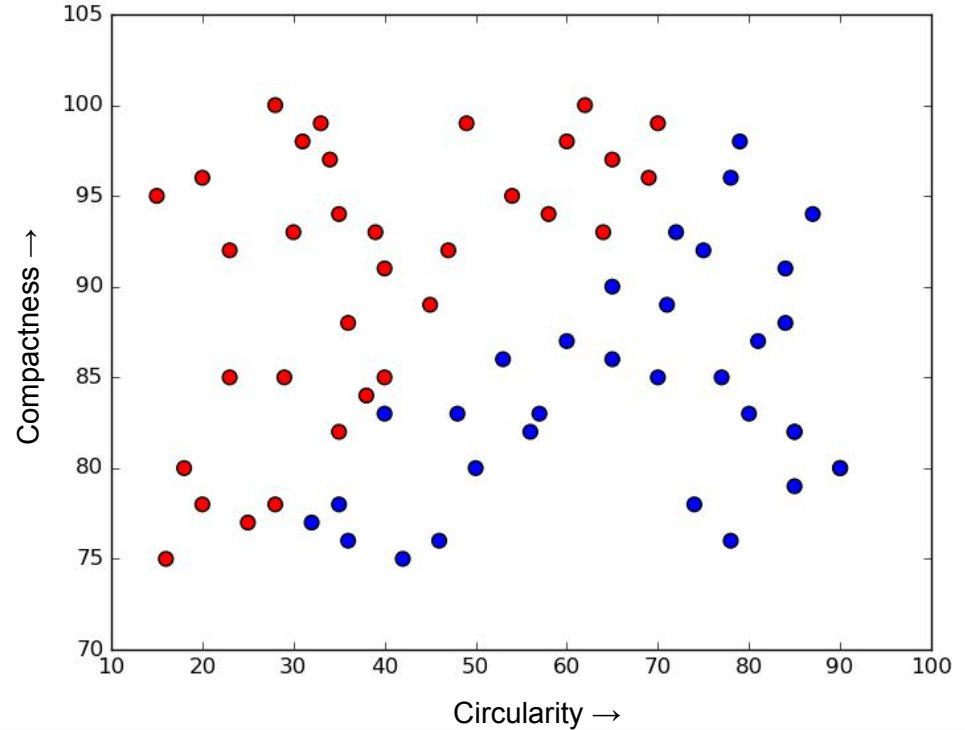
[IPython Notebook](#)

Closing Thoughts

What is the Cleaner way of Solving this problem?

● Bolt

● Nut



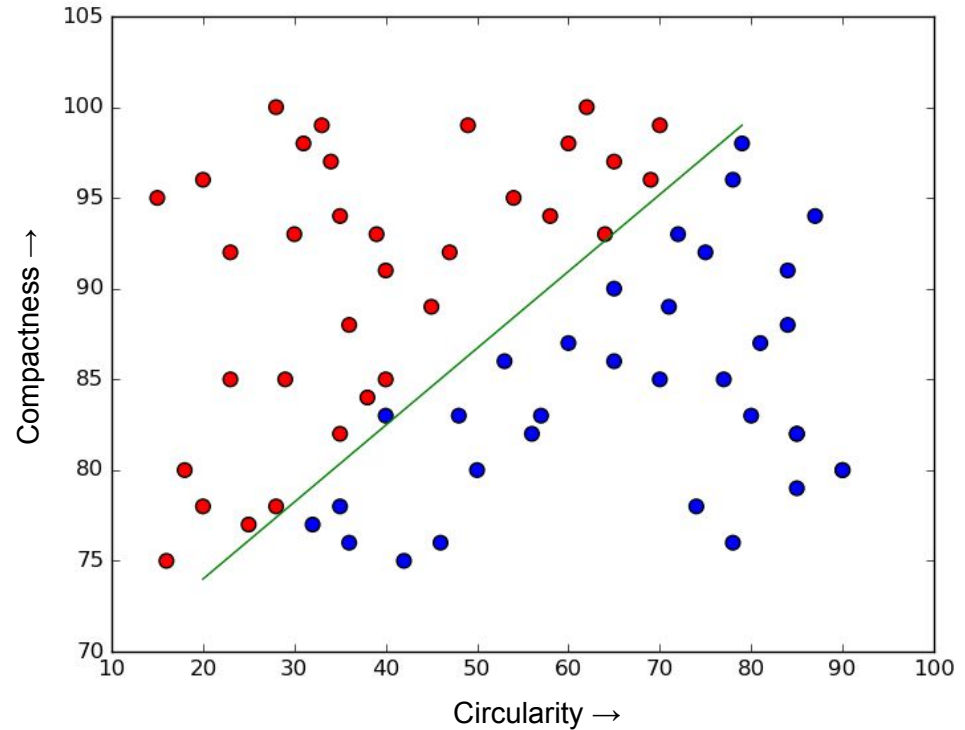
Linear Classifier

A line should classify the data

But, How do I come up with a Line which separates both the classes of data optimally?

● Bolt

● Nut





Now, the Session has
Ended..!

