Image Features (I)

Dr. Gerhard Roth

COMP 4102A
Winter 2015
Version 2

Image Features

Image features – may appear in two contexts:

- Global properties of the image (average gray level, etc) – global features
- Parts of the image with special properties (line, circle, textured region) – local features

Here, assume second context for image features:

- Local, meaningful, detectable parts of the image
- Should also be invariant to changes in the image

Detection of image features

- Detection algorithms – produce feature descriptors
  - Feature descriptors – often just high dimensional vectors
- Example – line segment descriptor: coordinates of mid-point, length, orientation

How invariant is the feature

- If feature can be extracted if conditions change then it is invariant to that condition
  - Extract the feature
  - Change that particular condition
  - Extract feature again, and compare the before and after
  - Little change makes feature invariant to that condition
- What are possible conditions?
  - Change scale, move closer or farther
  - Change viewpoint, a number of possibilities
    - Rotate the object in the camera plane
    - Rotate the object out of the camera plane
    - Translate the object
  - Change the lighting (darker, lighter, shadows, etc.)

How discriminatory is the feature?

- If we can use this feature to distinguish or classify objects then it must be discriminatory
- To be discriminatory it must be different for various types of objects otherwise it fails
- Shape and colour of street signs enables us to discriminate between types of signs
- But all speed limit signs have same shape and colour (square, black)
- Need to look at the text to discriminate between different speed limit signs

Edges in Images

Definition of edges

- Edges are significant local changes of intensity in an image.
- Edges typically occur on the boundary between two different regions in an image.

Origin of Edges

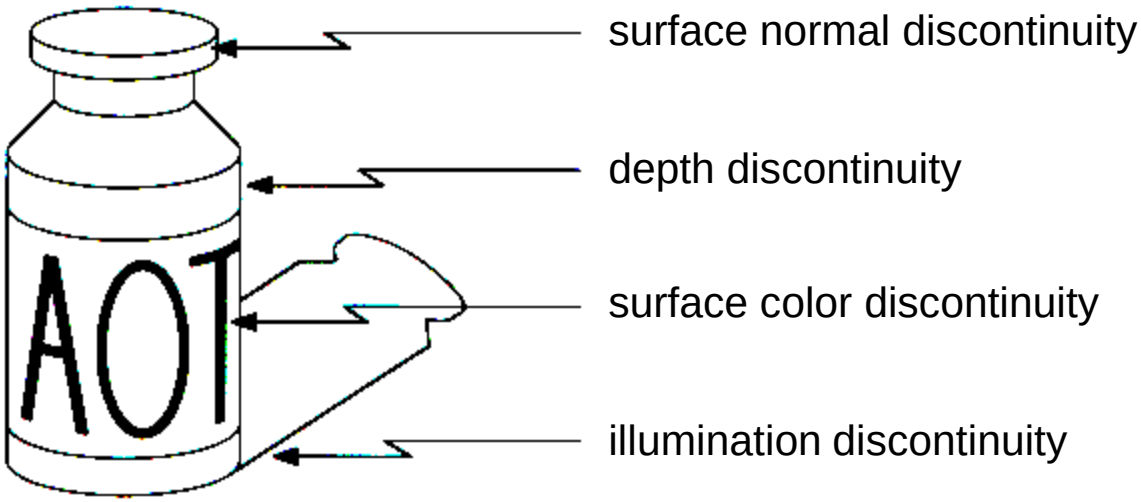

Edges are caused by a variety of factors

What causes intensity changes?

- Geometric events
  - object boundary (discontinuity in depth and/or surface color and texture)
  - surface boundary (discontinuity in surface orientation and/or surface color and texture)
- Non-geometric events
  - specularity
  - shadows (from other objects or from the same object)
  - inter-reflections

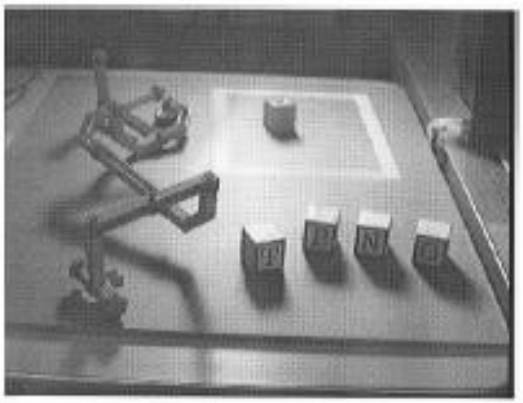

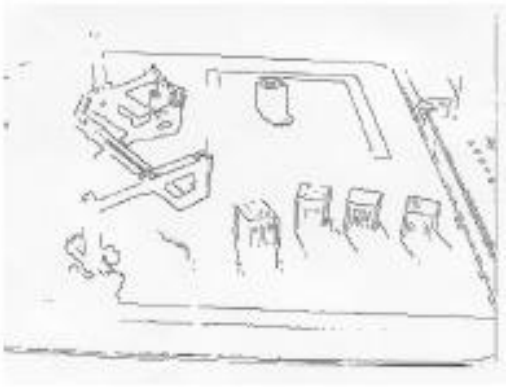

An edge is not a line...

Human visual system computes edges

- Regions of brain called V1 (in red) find edges

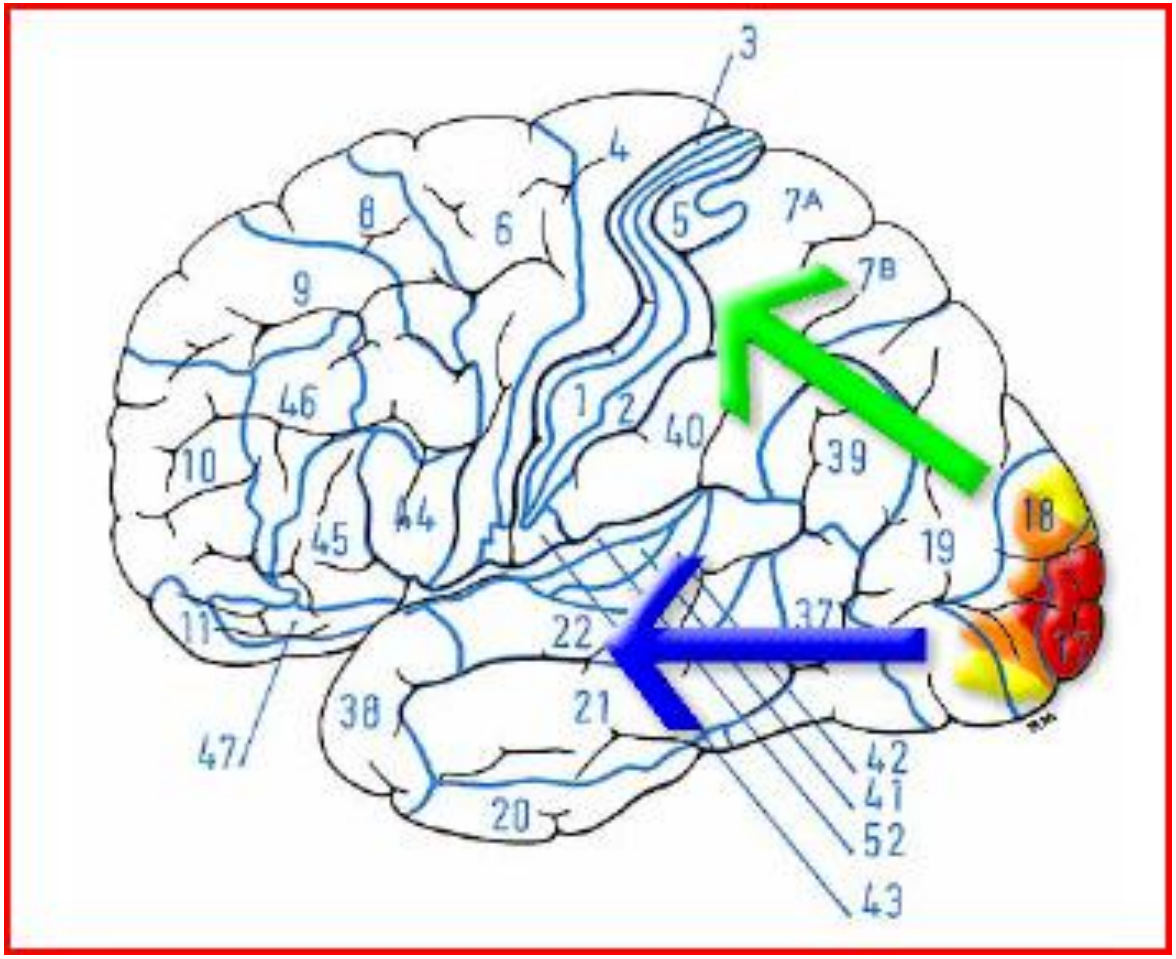

Simple and Complex cell

- These cells are local feature detectors

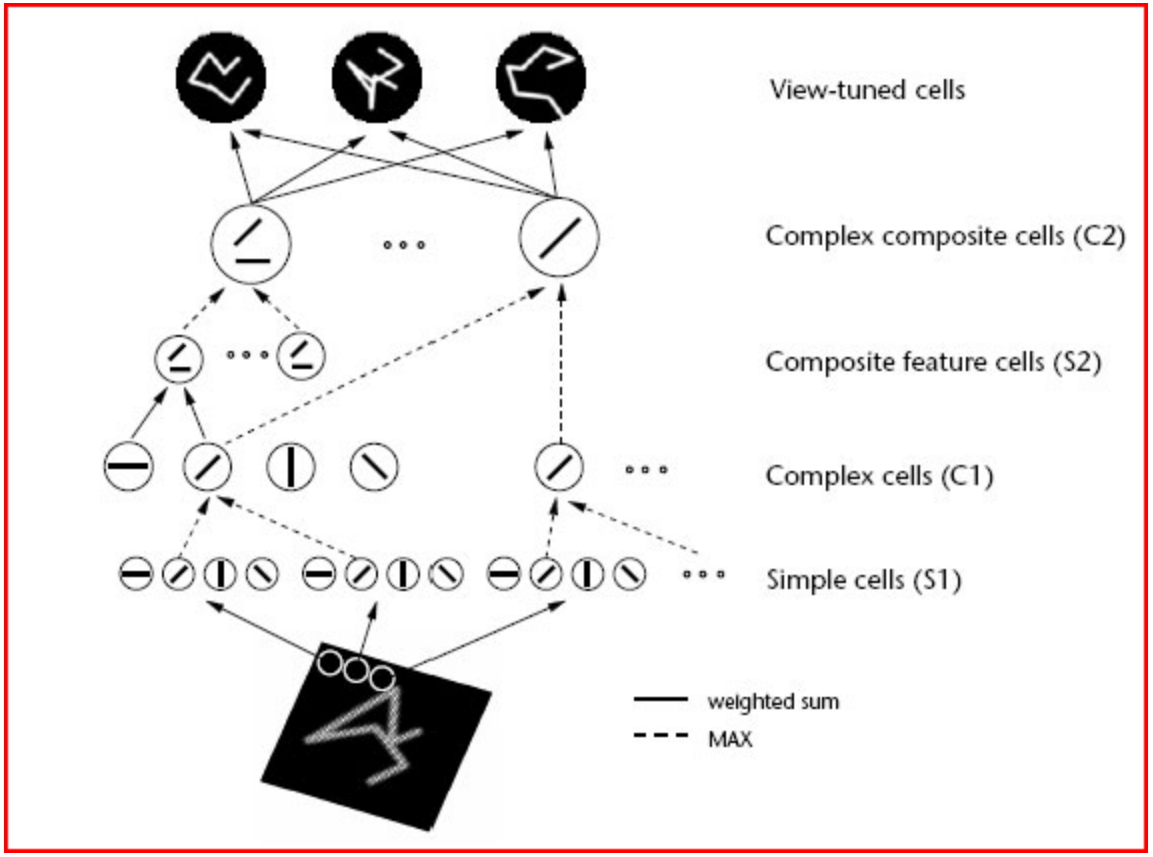

Result is an “edge like” representation

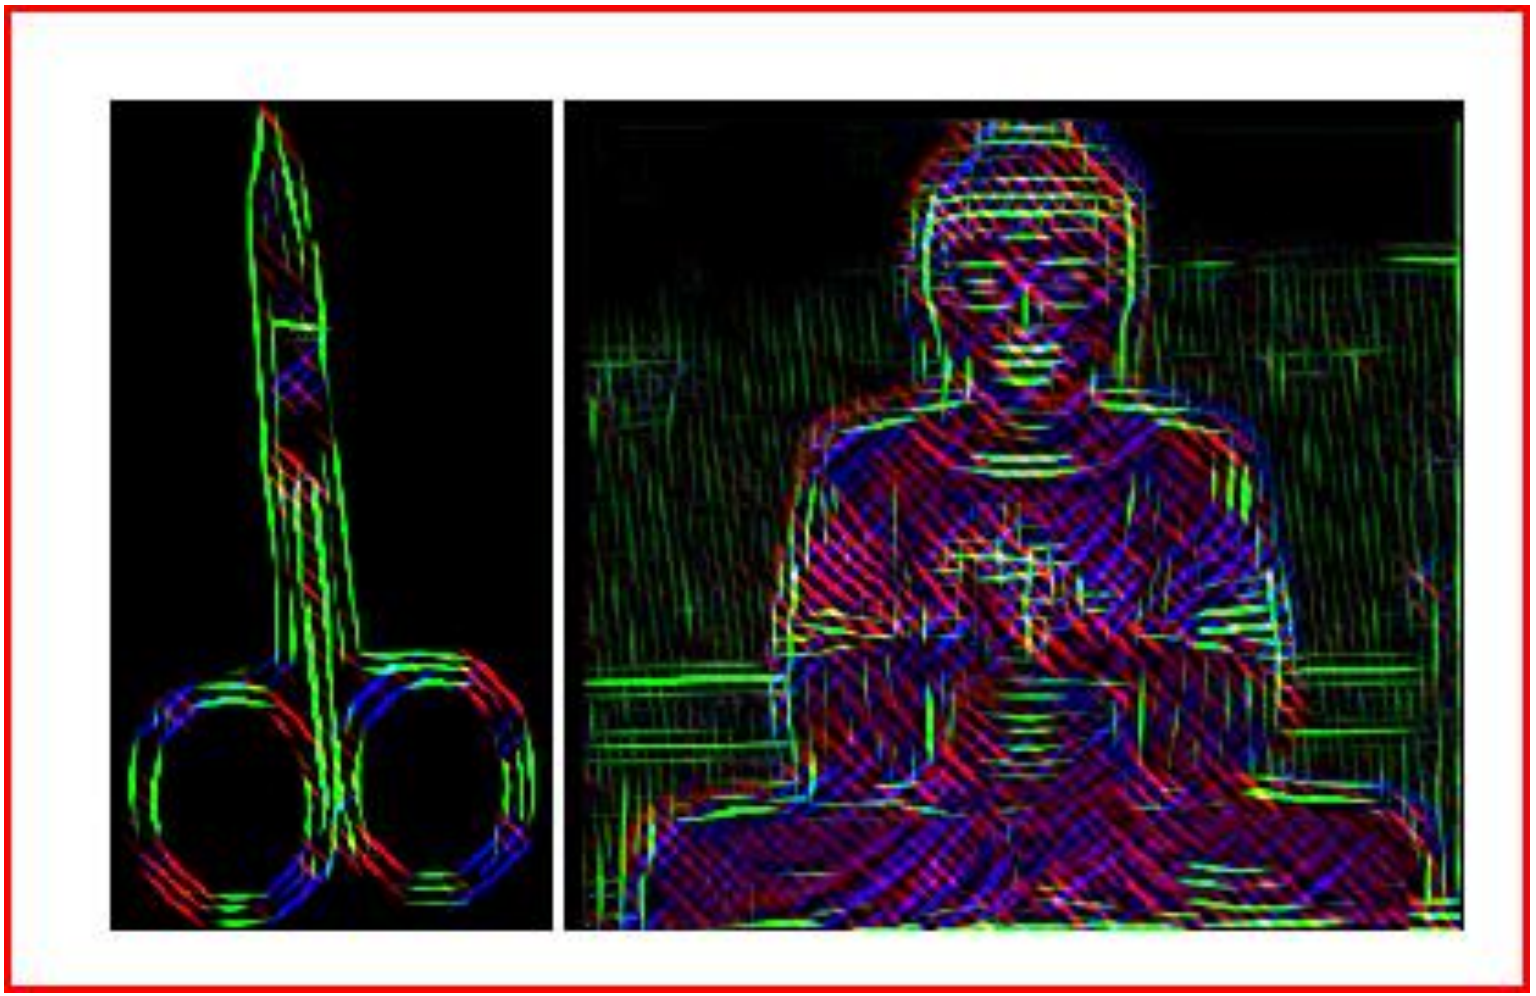

Edge Pixel Descriptors

- Edges are a connected set of edge pixels, each edge pixel has:
- Edge normal: unit vector in the direction of maximum intensity change.
  - Often called edge gradient (orthogonal to the edge direction)
- Edge direction: unit vector to perpendicular to the edge normal.
- Edge position or center: the pixel position at which the edge is located.
- Edge strength: related to the local image contrast along the normal.

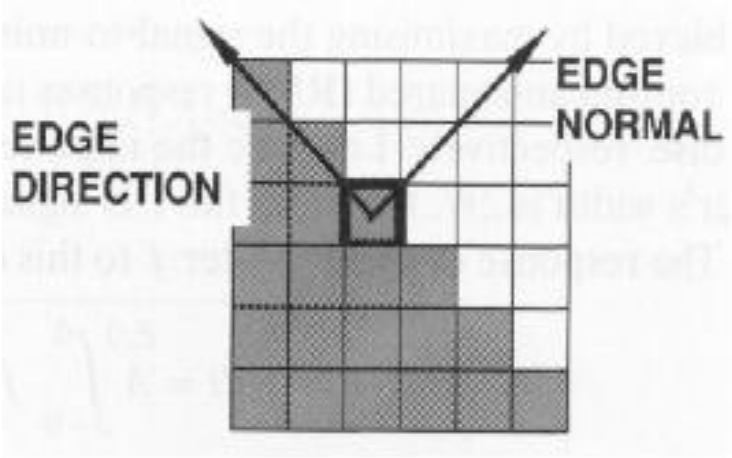

Applications of Edge Detection

- Produce a line drawing of a scene from an image of that scene.
- Important features can be extracted from the edges of an image (e.g. corners, lines, curves).
- These features are used by higher-level computer vision algorithms (e.g., segmentation, recognition, retrieval).

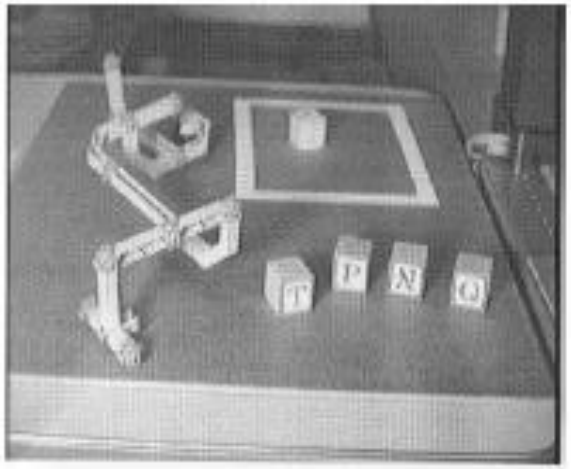

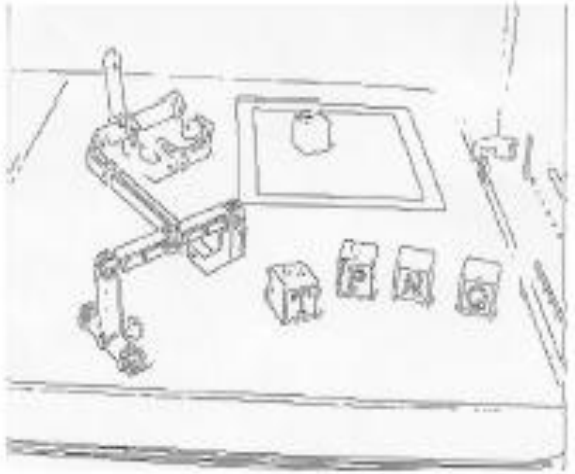

Three Steps of Edge Detection

- Noise smoothing
  - Suppress the noise without affecting the true edges
    - Often blur the image with Gaussian kernel of a given sigma
- Edge enhancement
  - Design a filter responding to edges, so that the output of the filter is large at edge pixels, so edges are localized as maxima in the filters response
- Edge localization
  - Decide which local maxima in the filters output are edges, and which are caused by noise. This usually involves:
    - Thinning the edges to 1 pixel width (non-maxima suppression)
    - Establish the minimum value to declare a local maxima as a true edge (thresholding)

Images as Functions

1-D

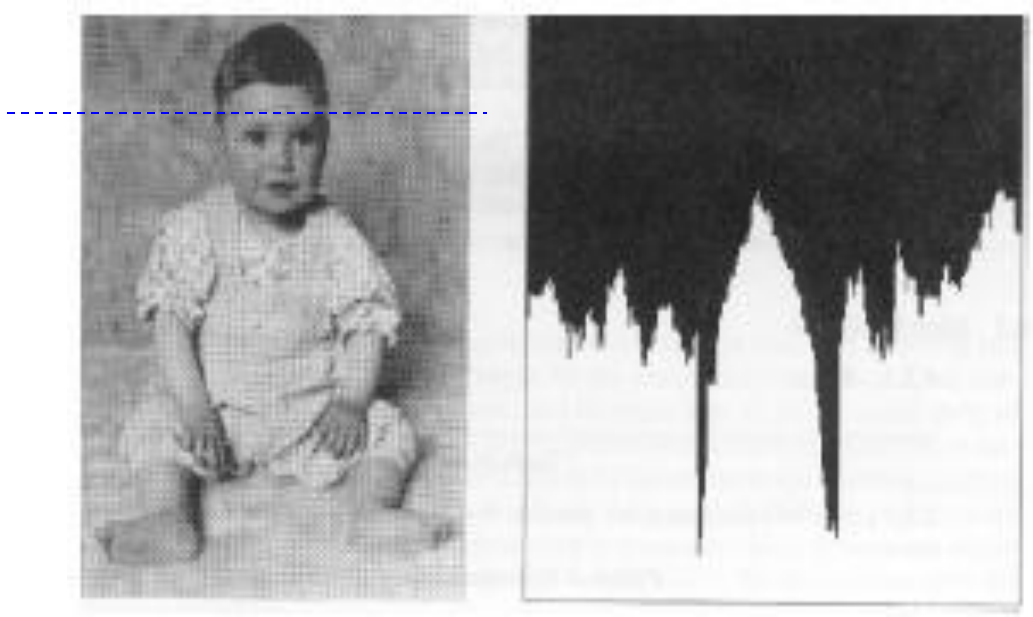

$$I = f(x)$$

Images as Functions

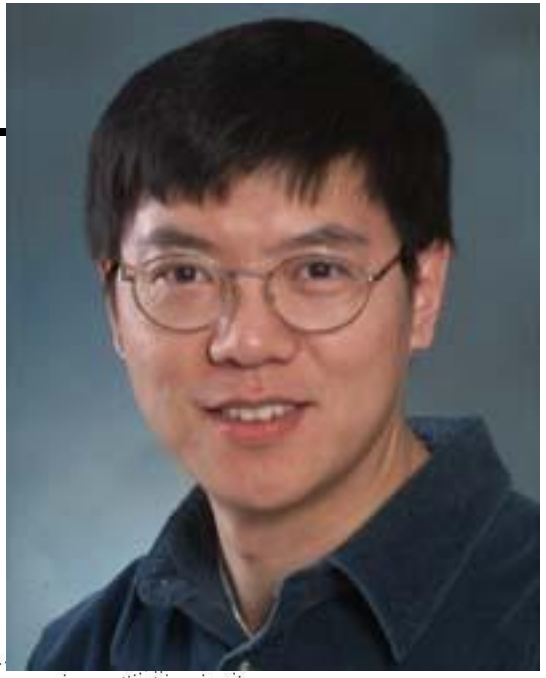

2-D

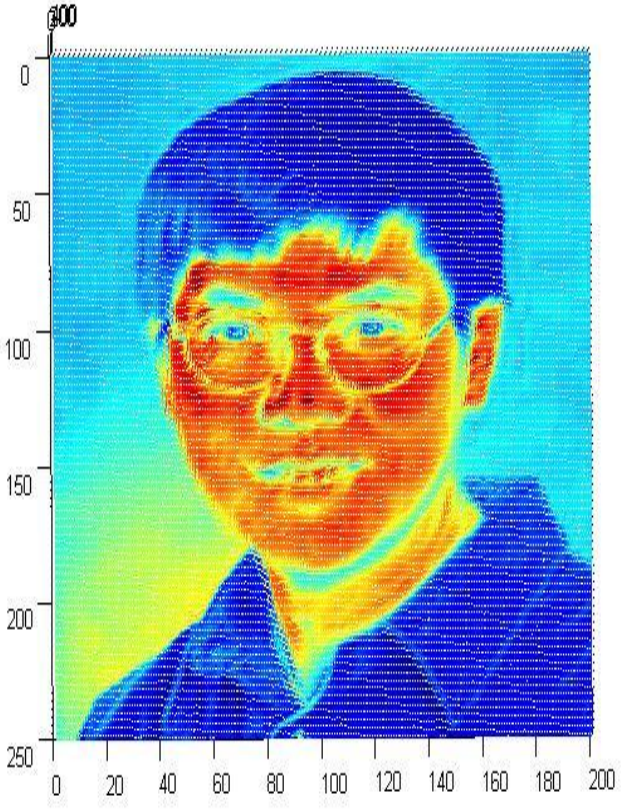

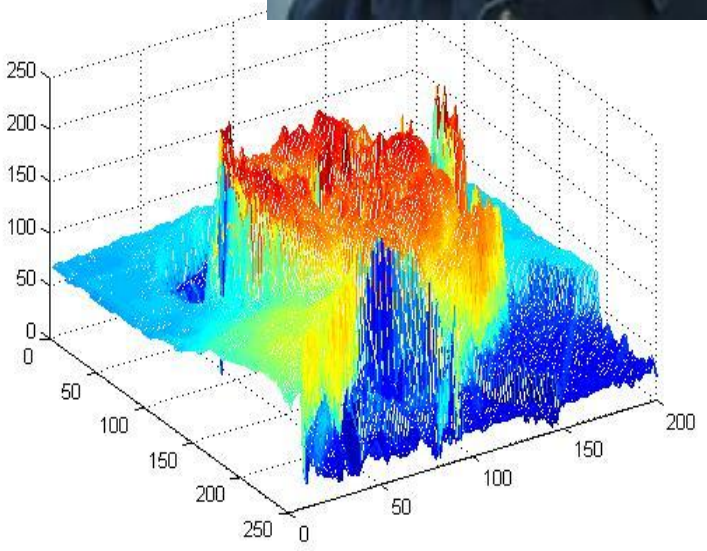

Red channel intensity

$$I = f(x, y)$$

Edge Detection using Derivatives

- Calculus describes changes of continuous functions using ***derivatives***.
- An image is a 2D function, so operators finding edges are based on ***partial derivatives***.
- Points which lie on an edge can be detected by either:
  - **detecting local maxima or minima of the first derivative**
  - **detecting the zero-crossing of the second derivative**
- Here we assume that there is no smoothing in the edge detection process
  - We are only looking at enhancement and localization

Edge Detection Using Derivatives

image

profile of a horizontal line

first derivative

second derivative

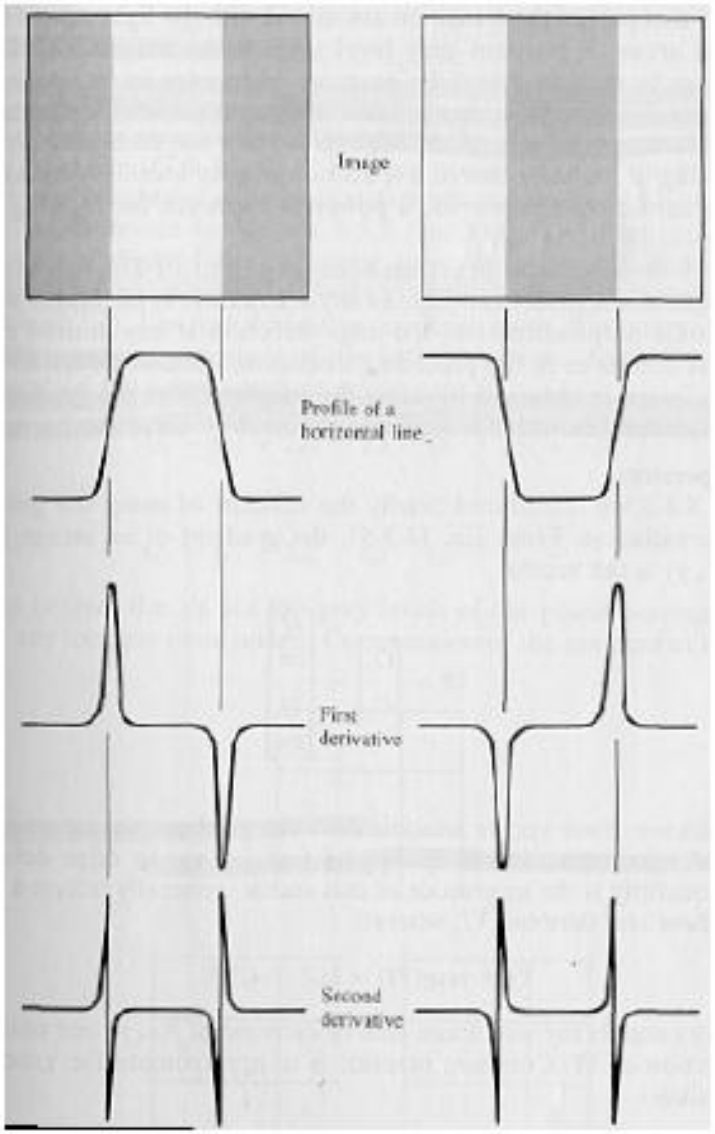

Finite Difference Method

We approximate derivatives with differences.

Derivative for 1-D signals:

Continuous function

$$f'(x) = \lim_{h \to 0} \frac{f(x+h) - f(x)}{h}$$

Discrete approximation

$$f'(x) \approx \frac{f_{i+1} - f_i}{i+1-i} = f_{i+1} - f_i$$

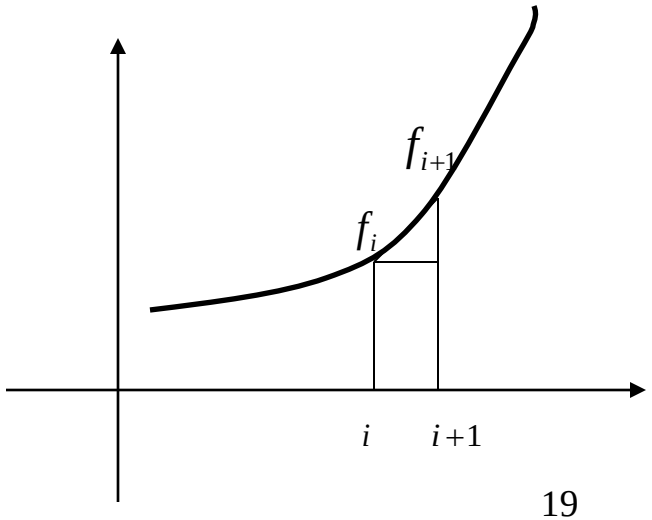

Finite Difference and Convolution

Finite difference on a 1-D image

$$f'(x) \approx f(x_{i+1}) - f(x_i)$$

is equivalent to convolving with kernel: $\begin{bmatrix} -1 & 1 \end{bmatrix}$

Finite Difference – 2D

Continuous function:

$$\frac{\partial f(x,y)}{\partial x} = \lim_{h \to 0} \frac{f(x+h,y) - f(x,y)}{h}$$

$$\frac{\partial f(x,y)}{\partial y} = \lim_{h \to 0} \frac{f(x,y+h) - f(x,y)}{h}$$

Discrete approximation:

$$I_x = \frac{\partial f(x,y)}{\partial x} \approx f_{i+1,j} - f_{i,j}$$

$$I_y = \frac{\partial f(x,y)}{\partial y} \approx f_{i,j+1} - f_{i,j}$$

Convolution kernels:

$$\begin{bmatrix} -1 & 1 \end{bmatrix}$$

$$\begin{bmatrix} -1 \\ 1 \end{bmatrix}$$

Image Derivatives

Image I

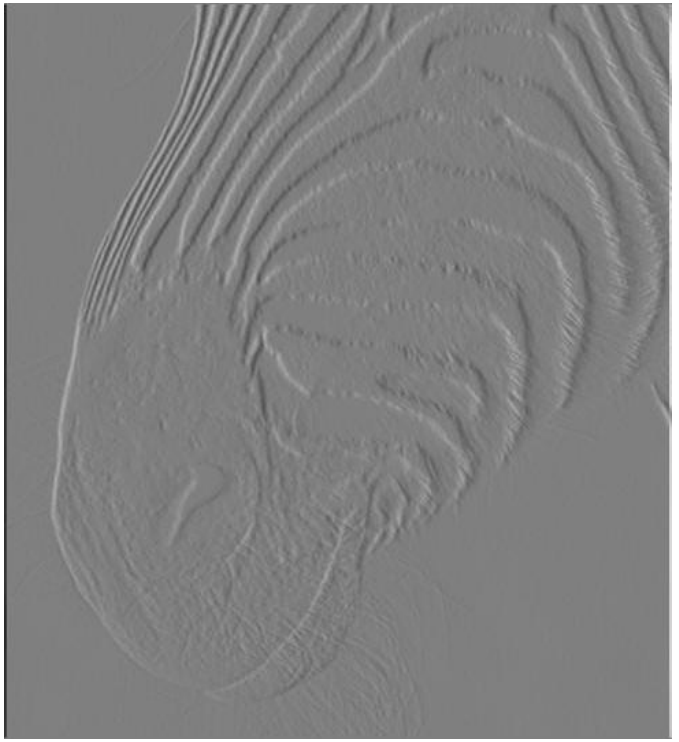

$$I_x = I * \begin{bmatrix} -1 & 1 \end{bmatrix}$$

Image Derivatives

Image I

$$I_x = I * \begin{bmatrix} -1 & 1 \end{bmatrix}$$

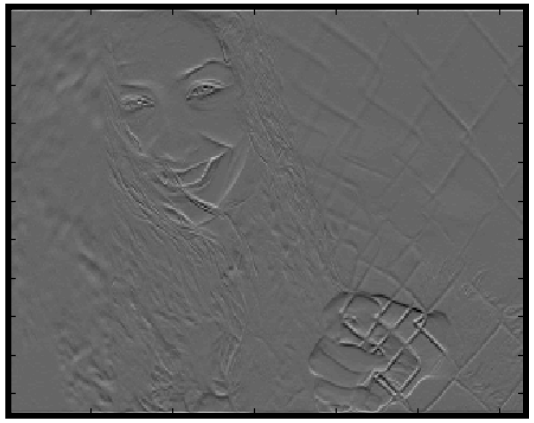

$$I_y = I * \begin{bmatrix} -1 \\ 1 \end{bmatrix}$$