CSE 564: Scientific Visualization

Lecture 5: Image Processing

Klaus Mueller

Stony Brook University

Computer Science Department

Image Processing Definitions

- Purpose:
  - enhance certain features of the image
  - de-emphasize other features of the image

- Implemented as filters or transformations:
  - some operate on the entire set of pixels at once (global operations)

    examples: brightness and contrast enhancement
  - some operate only on a subset of pixels (local operations in a pixel neighborhood)

    examples: edge detection, contouring, image sharpening, blurring

Intensity Transformations

- Problem: we only have a fixed number of grey / color levels to work with (e.g., 8 bits / 24 bits)
- Need to use this 'real estate' wisely to bring out the image features that we want
- Intensity transformations T_p enhance certain intensity ranges at the cost of compressing others

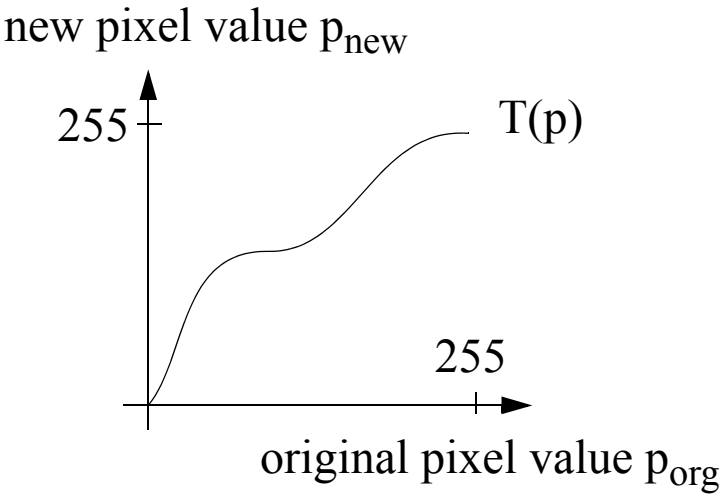

- Examples:

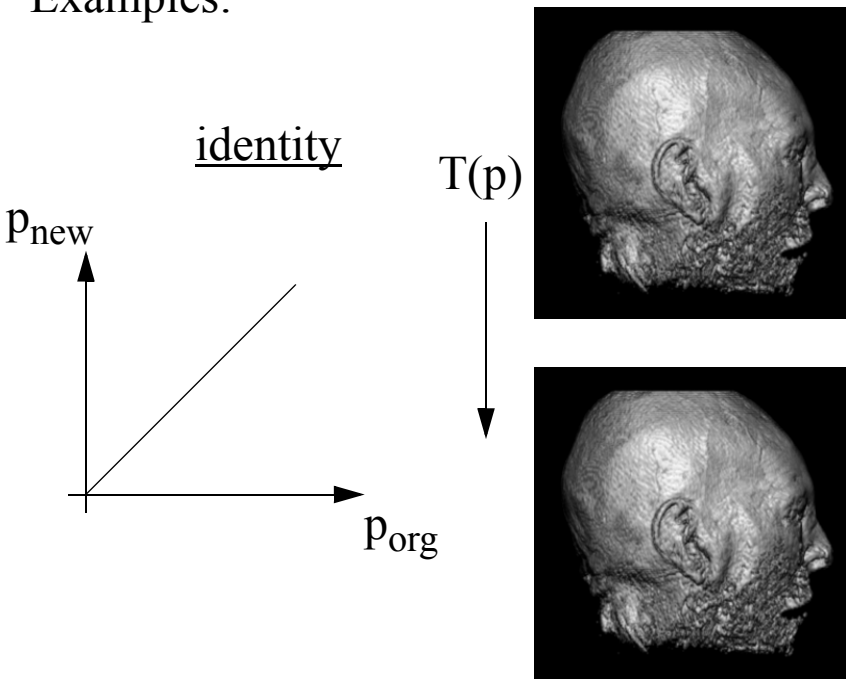

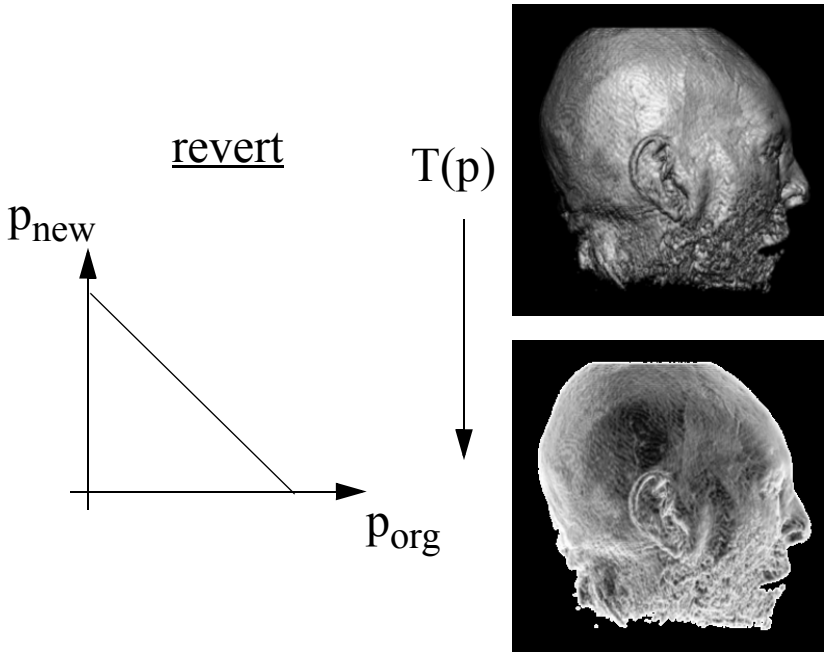

Intensity Transformations II

- More examples

contrast stretching

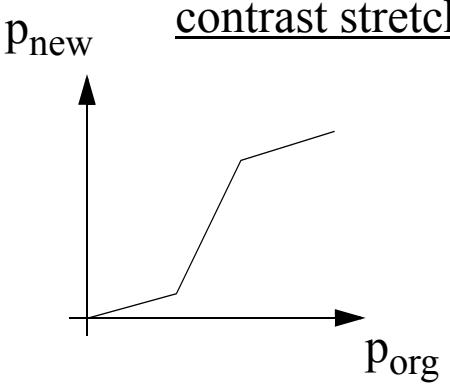

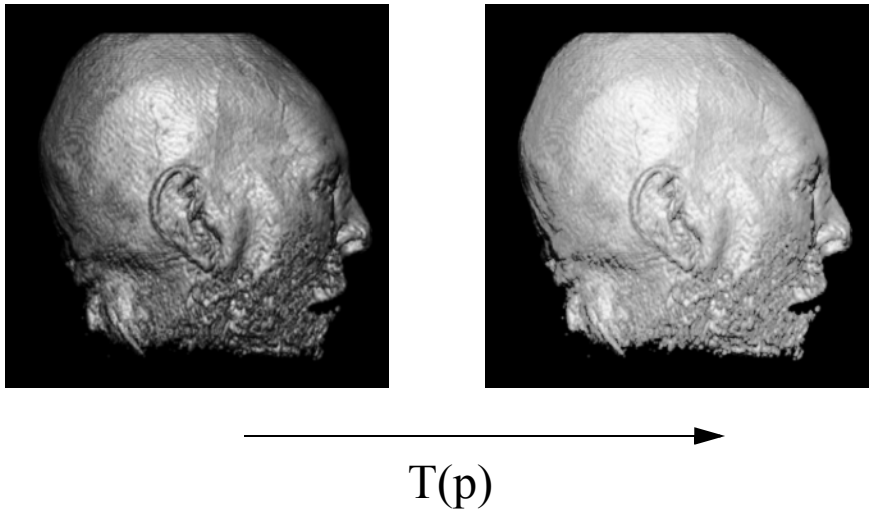

threshold

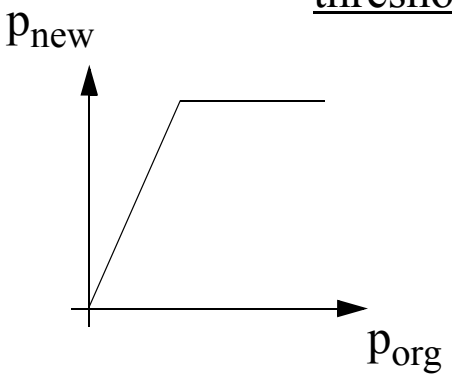

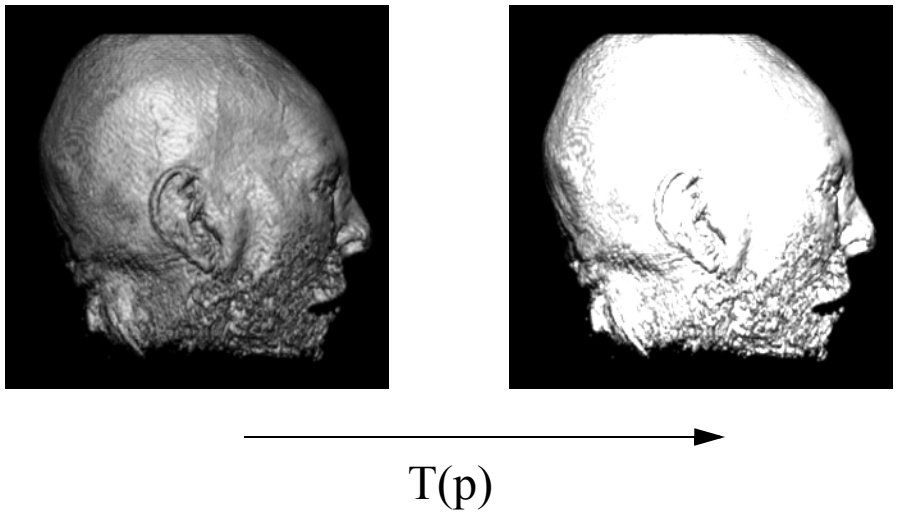

Color Transformations

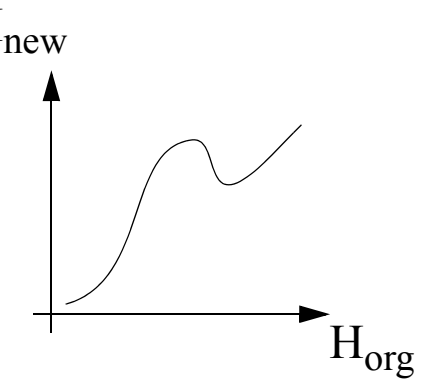

- Convert the image from RGB to HSV space
- Perform transformations of pixel H, S, V values via transfer functions
- Convert the transformed HSV image back into RGB space and display

Original

Hue transformed

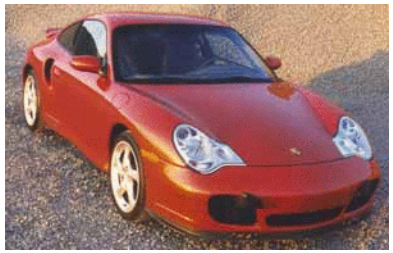

The Image Histogram

- A histogram lists the number of image pixels for each value

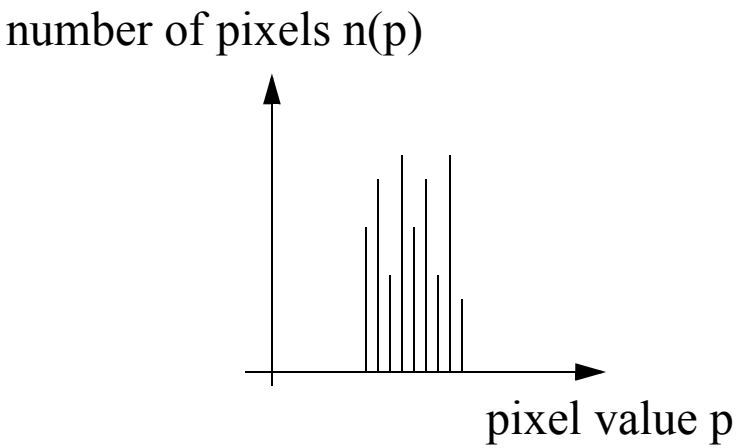

- The histogram reveals more insight about image contrast and brightness:

dark image

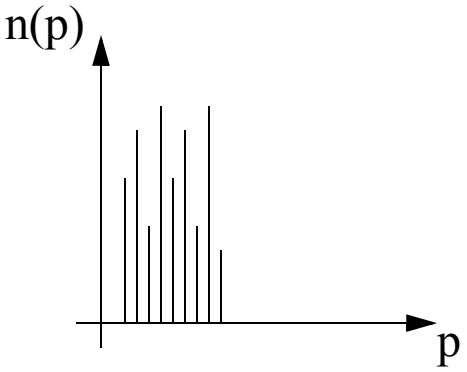

bright image

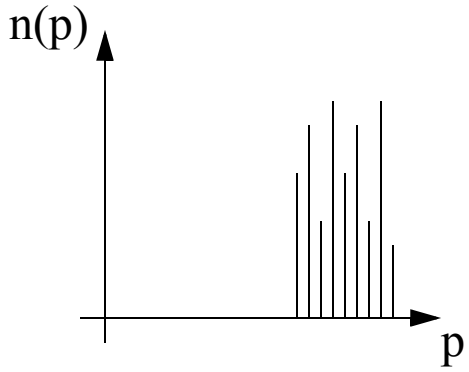

image with low contrast

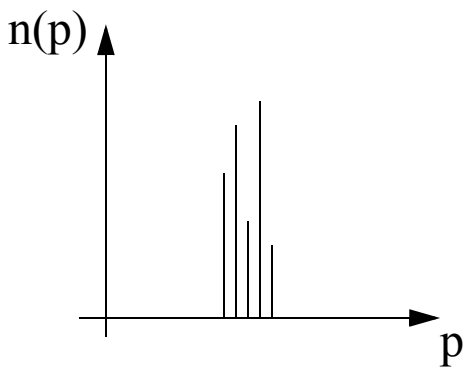

Histogram Processing

- Image contrast and brightness may be improved by modifying the histogram
- The 'contrast stretching' operation requires the user to manipulate the image's histogram
- *Histogram equalization* is an automatic procedure to spread out the value distribution

The discrete histogram equalization equation is: $p_{new}(k) = \sum_{j=0}^{k} \frac{n(p_{org}(j))}{n_{total}} \cdot p_{max}$

255

- For example, the equalization transformation for a dark image would be:

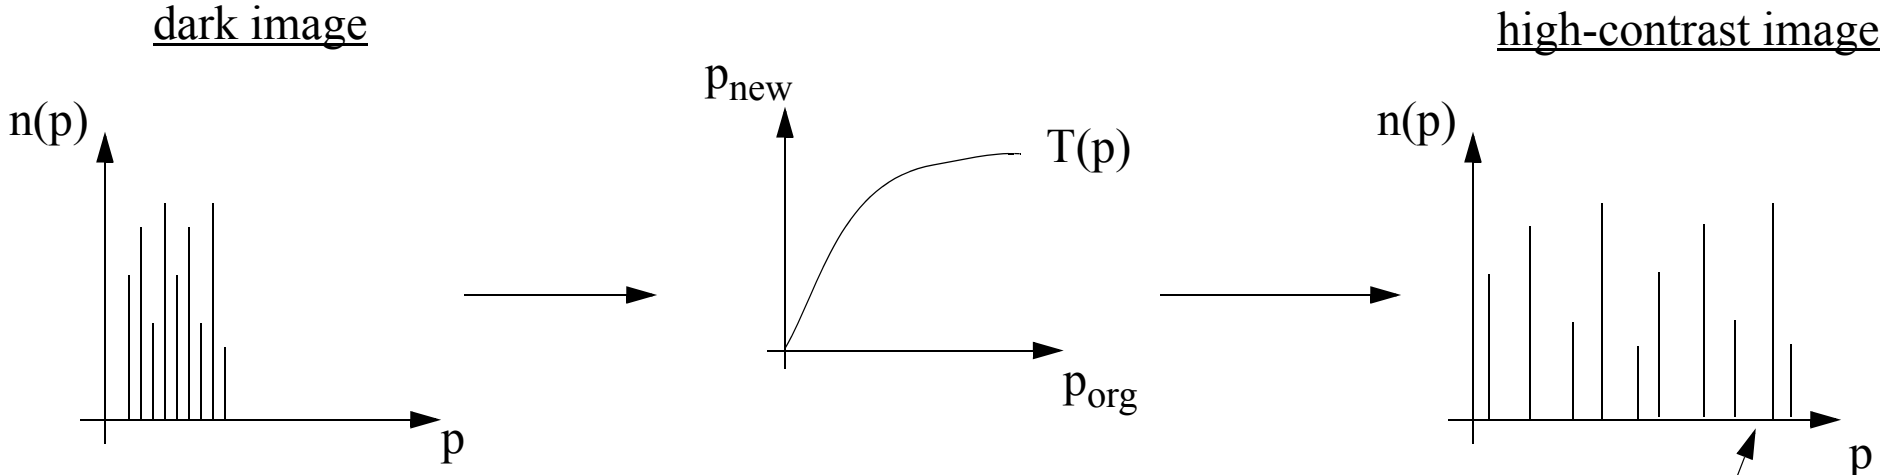

Note: the bars do not change in height, they are just shifted to different positions

Histogram Processing - Examples

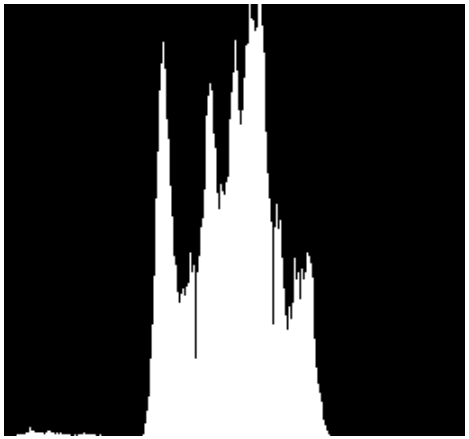

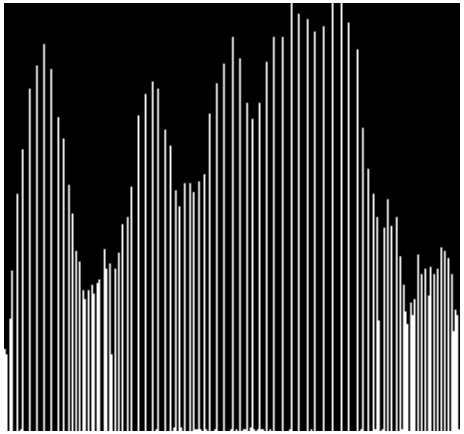

before equalization

after equalization

before and after - guess which

Discrete Convolution

- Instead of global transformations, we can also modify the image based on local information:
  - smoothing (examples: noise removal, preparation for image reduction)
  - edge enhancement (examples: de-blurring, sharpening, preparation for feature extraction)
- We use discrete convolution for these operations:
  - place a weight matrix or *mask* at each pixel location
  - this mask weighs the pixel's neighborhood and determines the output pixel's value

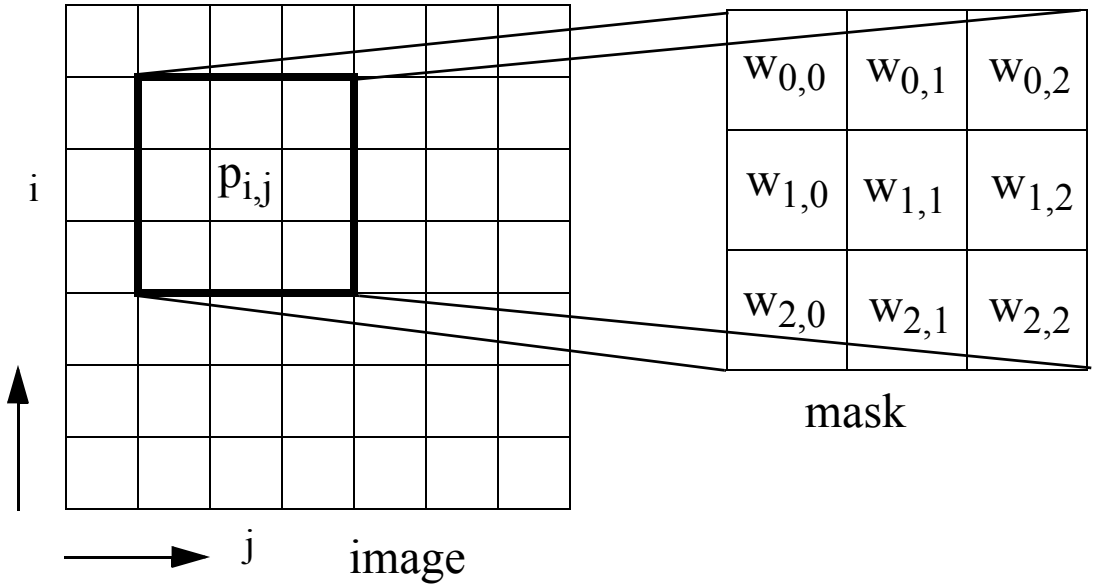


```
for each i, j
    temp = 0
    for each k, l
        temp += p^org_{i-1+k, j-1+l} · w_{k,l}
    p^new_{i,j} = temp
```

$$p^{new}_{i,j} = \sum_{k=0}^{2} \sum_{l=0}^{2} p^{org}_{i-1+k,\, j-1+l} \cdot w_{k,l}$$

- Important: Do not replace the computed values into original image, but write to an output image

Image Smoothing

- Smoothing mask:

  - averages the pixel neighborhood
  - each pixel's value is replaced by its local average
  - can be used to remove high frequency noise
  - larger masks smooth more and cut more noise
  - always make sure that sum of all mask elements equals 1.0, else image changes brightness
  - a blurry image results
  - jagged edges are replaced by blur
    (for example, see circle)

$$1/9 \times \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

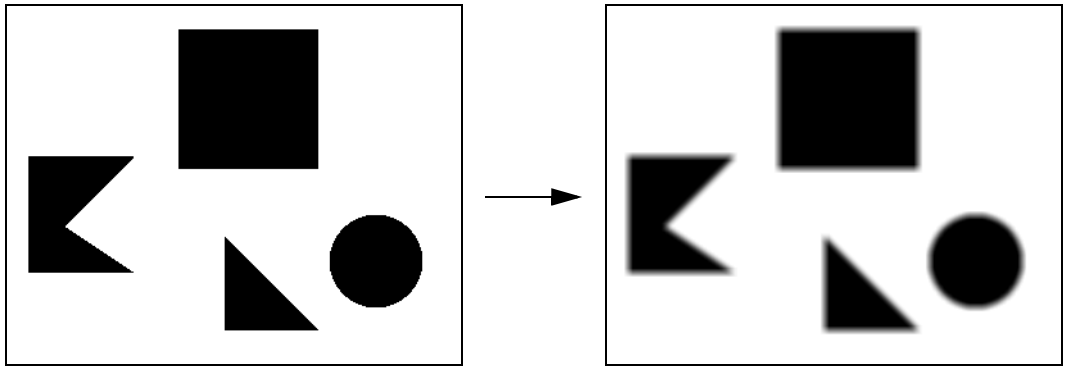

- Median Filter:

  - special filter that deals better with spike and spot noise
  - in contrast to the smoothing masks, it preserves edges and image sharpness
  - instead of adding all neighborhood pixels, it sorts them and picks the median as output pixel
  - images are generally less blurry than with convolution-based smoothing

Image Sharpening

- This operation enhances the edges, it is the opposite of smoothing
- It has little effect in smooth areas with no edges
- Note that an edge means that there is a high local derivative or gradient
- Sharpening masks implement some sort of image differentiation: $\nabla \boldsymbol{f} = \begin{bmatrix} \frac{\partial f}{\partial x} \\ \frac{\partial f}{\partial y} \end{bmatrix}$
- In most cases we are only interested in the magnitude: $\nabla f = mag(\nabla \boldsymbol{f}) = \left[\left(\frac{\partial f}{\partial x} \right)^2 + \left(\frac{\partial f}{\partial y} \right)^2 \right]^{\frac{1}{2}}$

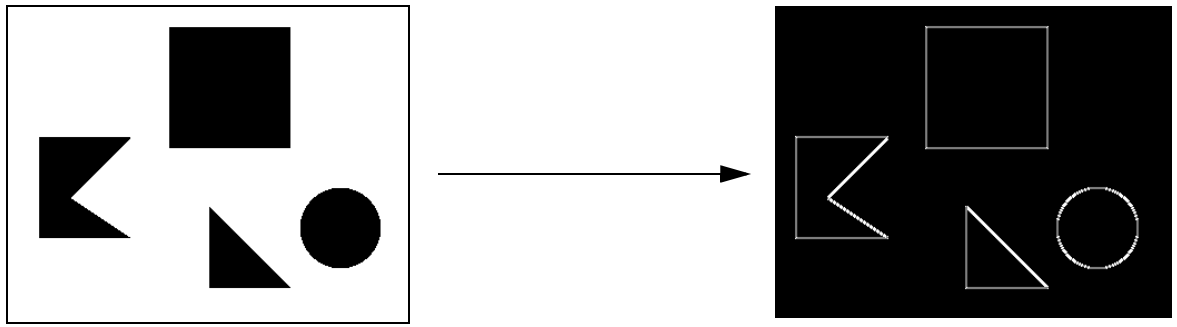

Image Sharpening - The Sobel Mask

1	2	1
0	0	0
-1	-2	-1

dy

Sobel

1	0	-1
2	0	-2
1	0	-1

dx

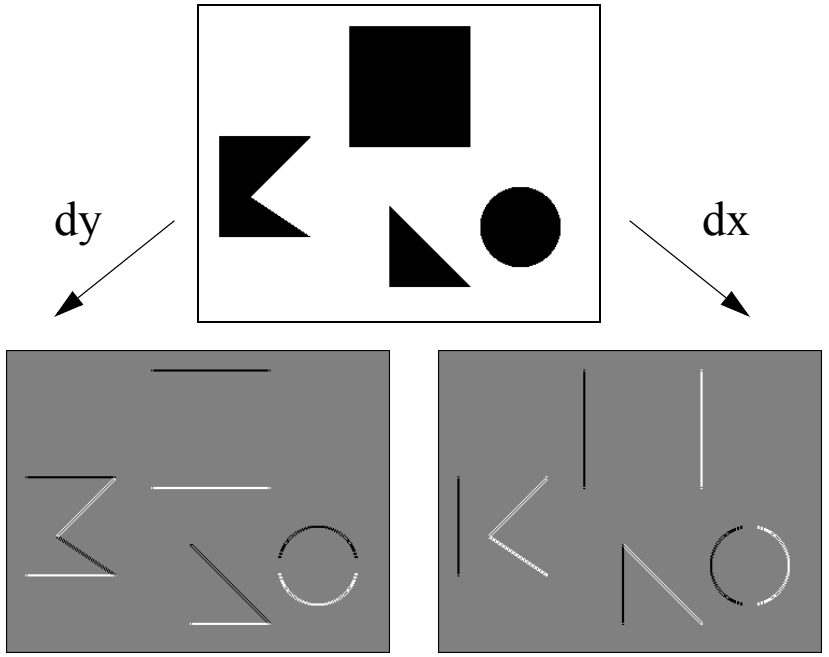

- The Sobel filter comes in a pair of two masks:

  - one mask computes an image for the x- derivative (dx), the other for the y-derivative (dy)

- Note that the dy-masks do some smoothing in the x-direction (dx-mask smoothes in y)

  - this decreases the sensitivity to noise (sharpening tends to magnify high frequency noise)

- Note that pixel values below zero will occur at edges with negative gradients

- We get two images, img_{dx} and img_{dy}, their pixels are combined by:

$$img_{new} = \left(img_{dx}^{\ 2} + img_{dy}^{\ 2} \right)^{\frac{1}{2}} \qquad or \qquad img_{new} = \left| img_{dx} \right| + \left| img_{dy} \right|$$

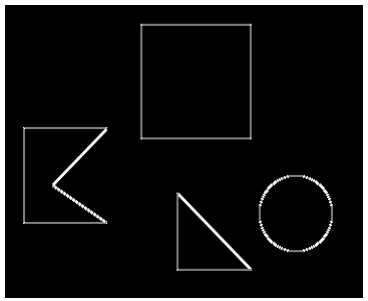

More Image Processing Techniques

- Image subtraction: compute the difference between two images

  - example: X-ray angiography to enhance perfused vessels

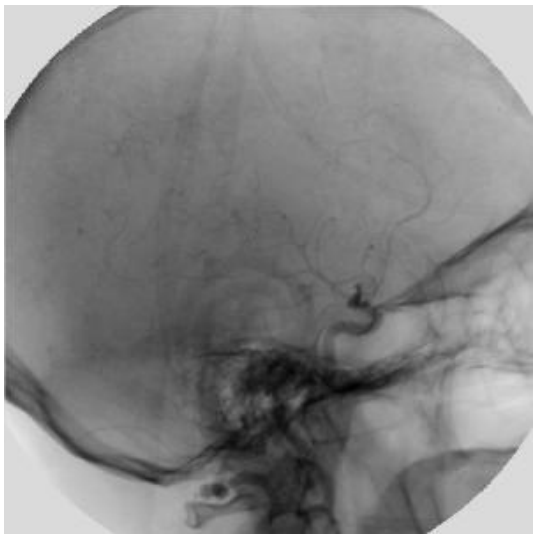

perfused

-

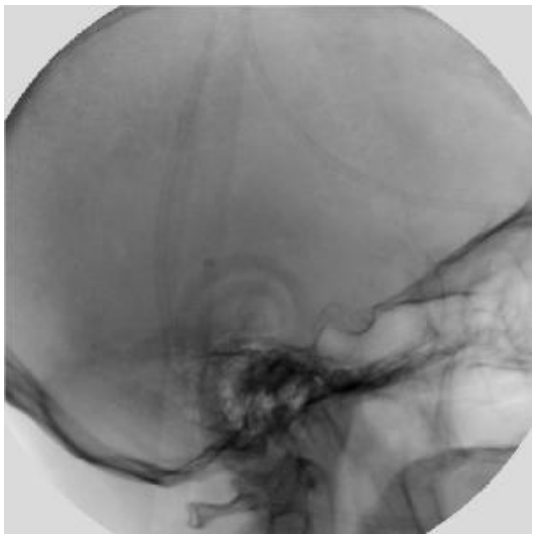

non-perfused (mask)

=

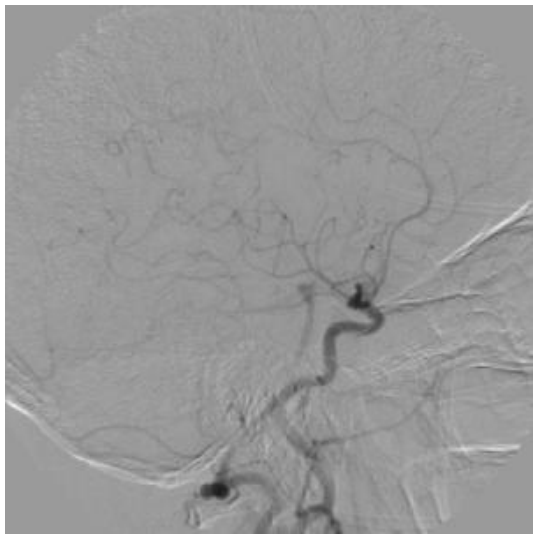

contrast-enhanced

- Image averaging:

  - Averaging of a series of noisy images of the same object removes noise and leaves object