

Image Enhancement


3D Image Processing

Torsten Möller

Overview

- What is an image?
 - quantization
 - color image
 - image pixels
- Quality
- Intensity transforms
- Histogram processing
- Filtering
- Smoothing
- Sharpening

What is a (digital) image?



- An image is made of **pixels** (=picture elements)
- the coordinate values are discretized

Image resolution

- the resolution of an image is the number of pixels per space unit



resolution too low: the image appears “pixelized”
enlargement needed (e.g. TV to HDTV conversion)



“pixelization” can be used for censorship

Quantization

- The pixel values are quantized: they belong to a discrete set of values, generally represented by integers between 0 and $N-1$



coded on 8 bits



coded on 4 bits



coded on 2 bits

Quantization

- The pixel values are quantized: they belong to a discrete set of values, generally represented by integers between 0 and $N-1$



coded on 8 bits
255 possible val.



coded on 4 bits




coded on 2 bits
4 possible val.

Quantization

- The pixel values are quantized: they belong to a discrete set of values, generally represented by integers between 0 and $N-1$



coded on 8 bits
255 possible val.



8 bits are enough (if no further processing, like histogram equalization) because humans do not see smaller intensity differences.

What is a color image?

image \mathbf{I}



red channel I^R



green channel I^G




blue channel I^B


Equivalently, each pixel value $\mathbf{I}[\mathbf{k}] \in \mathbb{R}^3$ has three components, each generally coded on 8 bits \rightarrow 24 bits/pixel

Representing an image

- 3 equivalent representations of the same image:



| | | |
|-----|-----|-----|
| 161 | 56 | 201 |
| 0 | 145 | 46 |
| 255 | 56 | 201 |



- A greyscale image is a two-dimensional array of real values. It can be viewed as a matrix, for instance.
- If the resolution is high, we get the impression that the image is defined continuously:
We don't see the pixels.



<http://jaredjared.com/2012/10/visual-acuity-dpi/>

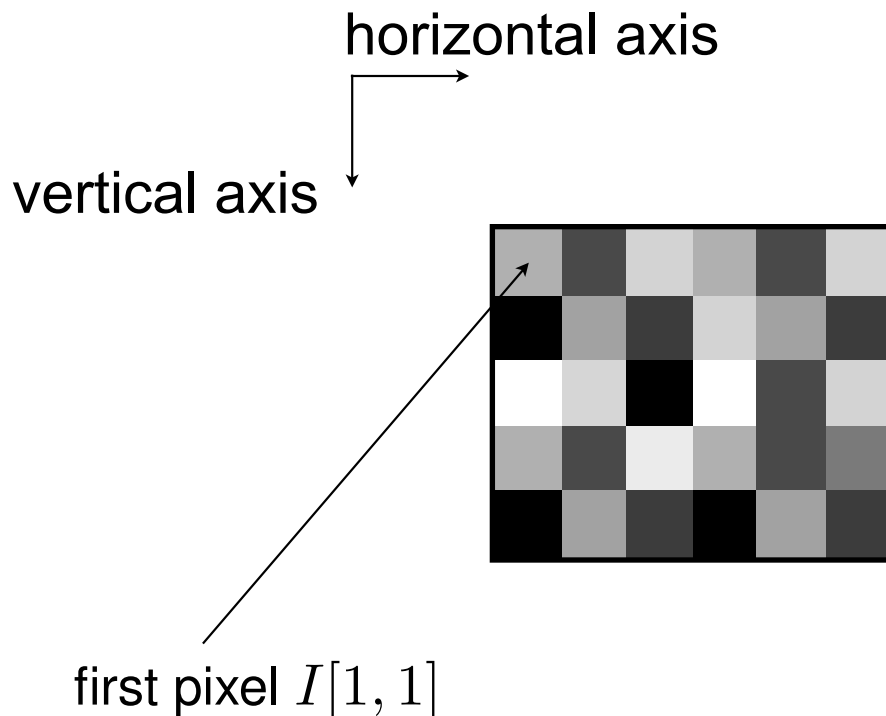
Indexing the pixels

An image is made of **rows** and **columns**

Usual way of indexing the pixel values:

$I[k_x, k_y]$ is the pixel value of the image I at location (k_x, k_y) , with $1 \leq k_x \leq W$ and $1 \leq k_y \leq H$.

W is the **width** and H is the **height**.



In matrix notations (e.g. in Matlab):
row-first indexing


$$I[k_y, k_x]$$

2D lattices


- An image is defined on a lattice.
- The most common is the Cartesian (a.k.a square) lattice.
- But other lattices exist and have interesting properties.





2D lattices



cartesian lattice



hexagonal lattice




Given two independent vectors \mathbf{r}_1 , \mathbf{r}_2 , the associated 2D lattice is the set of points in the plane, whose coordinates are linear combinations with integer weights of \mathbf{r}_1 , \mathbf{r}_2 .

The pixel shape is the **Voronoi cell** of the lattice.


The hexagonal lattice

- better isotropy:
12-symmetry,
6-connectivity



Grey level profile (cut)

- Viewing the pixel values along one row of an image
--> Matlab command `improfile`



Overview

- What is an image?
 - quantization
 - color image
 - image pixels
- Quality
- Intensity transforms
- Histogram processing
- Filtering
- Smoothing
- Sharpening

Image quality

- Common defaults in images:
 - blur (motion blur, out of focus blur ...)

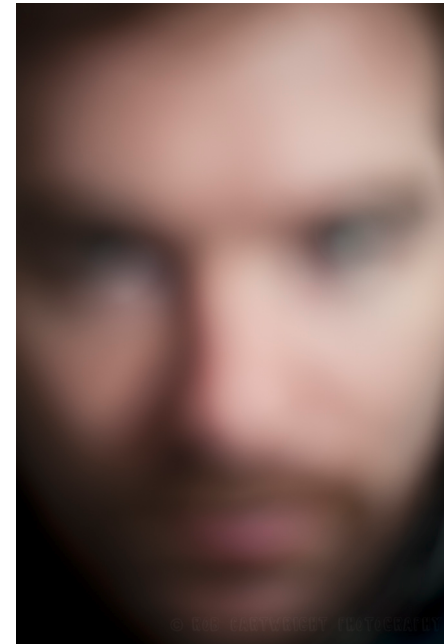


Image quality

- Common defaults in images:
 - blur (motion blur, out of focus blur ...)
 - ringing



Image quality

- Common defaults in images:
 - blur (motion blur, out of focus blur ...)
 - ringing
 - aliasing (staircasing, Moiré patterns)



Image quality

- Low-frequency Moiré artifacts appear when high-frequency content is sampled in an incorrect way.






Image quality

- We need to be able to measure the difference between two images, for instance an original image and a degraded one.
- Classical difference measures between two images I_1 and I_2 :

mean absolute error:
$$\text{MAE} = \frac{1}{WH} \sum_{k_x=1}^W \sum_{k_y=1}^H |I_1[k_x, k_y] - I_2[k_x, k_y]|$$

mean square error:
$$\text{MSE} = \frac{1}{WH} \sum_{k_x=1}^W \sum_{k_y=1}^H (I_1[k_x, k_y] - I_2[k_x, k_y])^2$$

peak signal to noise ratio (dB):
$$\text{PSNR} = 10 \log_{10} \left(\frac{255^2}{\text{MSE}} \right)$$

- There exist much more sophisticated quality measures and difference measures for images (SSIM...)

Overview

- What is an image?
 - quantization
 - color image
 - image pixels
- Quality
- Intensity transforms
- Histogram processing
- Filtering
- Smoothing
- Sharpening

Image enhancement

- Objective: process an image so that the result is more suitable than the original image for a specific application.
- In spatial domain: manipulation of the pixel values.
- In Fourier domain: manipulation of the frequency content.

Image enhancement

- Objective: process an image so that the result is more suitable than the original image for a specific application.
- In spatial domain: manipulation of the pixel values.
 - Example: taking the image negative $255-I$
 - we are more sensitive to low contrast in bright regions than in dark regions.





Image enhancement

- Objective: process an image so that the result is more suitable than the original image for a specific application.
- In spatial domain: manipulation of the pixel values.
 - Example: thresholding the pixel values







Image enhancement

- Intensity Transformations
 - image negatives
 - log transforms
 - gamma (power-law) transforms
 - contrast stretching
 - intensity-level slicing
 - bit-plane slicing

Log transform



a b

FIGURE 3.5
 (a) Fourier spectrum.
 (b) Result of applying the log transformation in Eq. (3.2-2) with $c = 1$.

Gamma (power-law) transform





FIGURE 3.8
 (a) Magnetic resonance image (MRI) of a fractured human spine.
 (b)–(d) Results of applying the transformation in Eq. (3.2-3) with $c = 1$ and $\gamma = 0.6, 0.4,$ and 0.3 , respectively. (Original image courtesy of Dr. David R. Pickens, Department of Radiology and Radiological Sciences, Vanderbilt University Medical Center.)

Gamma (power-law) transform

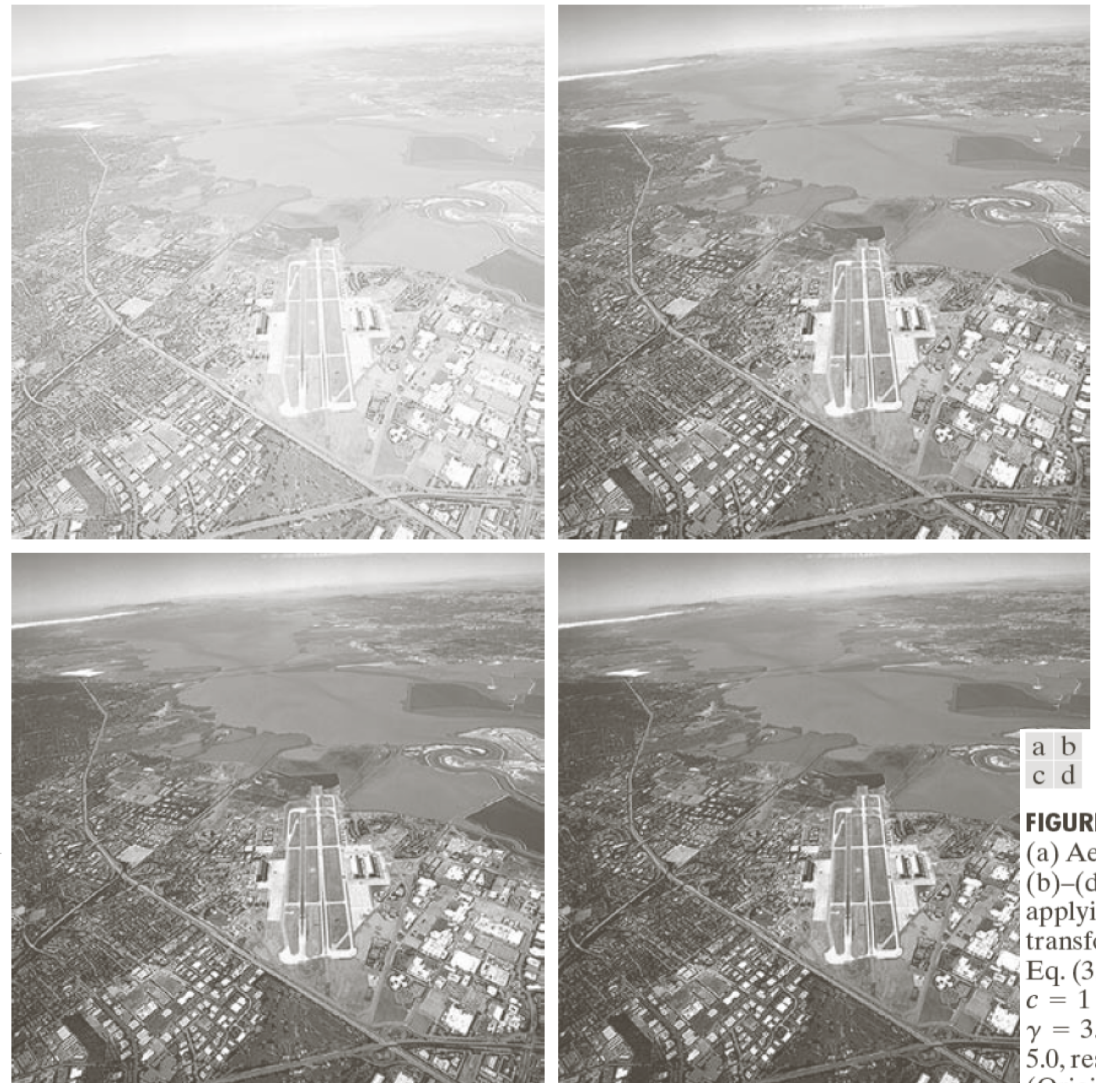




FIGURE 3.9 (a) Aerial image. (b)–(d) Results of applying the transformation in Eq. (3.2-3) with $c = 1$ and $\gamma = 3.0, 4.0,$ and 5.0 , respectively. (Original image for this example courtesy of NASA.)

Contrast stretching




a b
c d

FIGURE 3.10
Contrast stretching.
(a) Form of transformation function. (b) A low-contrast image. (c) Result of contrast stretching. (d) Result of thresholding. (Original image courtesy of Dr. Roger Heady, Research School of Biological Sciences, Australian National University, Canberra, Australia.)

Intensity-level slicing

a b


FIGURE 3.11 (a) This transformation highlights intensity range $[A, B]$ and reduces all other intensities to a lower level. (b) This transformation highlights range $[A, B]$ and preserves all other intensity levels.



a b c

FIGURE 3.12 (a) Aortic angiogram. (b) Result of using a slicing transformation of the type illustrated in Fig. 3.11(a), with the range of intensities of interest selected in the upper end of the gray scale. (c) Result of using the transformation in Fig. 3.11(b), with the selected area set to black, so that grays in the area of the blood vessels and kidneys were preserved. (Original image courtesy of Dr. Thomas R. Gest, University of Michigan Medical School.)

Bit-plane slicing



| | | |
|---|---|---|
| a | b | c |
| d | e | f |
| g | h | i |

FIGURE 3.14 (a) An 8-bit gray-scale image of size 500×1192 pixels. (b) through (i) Bit planes 1 through 8, with bit plane 1 corresponding to the least significant bit. Each bit plane is a binary image.

Bit-plane slicing



a b c


FIGURE 3.15 Images reconstructed using (a) bit planes 8 and 7; (b) bit planes 8, 7, and 6; and (c) bit planes 8, 7, 6, and 5. Compare (c) with Fig. 3.14(a).

Overview


- What is an image?
 - quantization
 - color image
 - image pixels
- Quality
- Intensity transforms
- Histogram processing
- Filtering
- Smoothing
- Sharpening

Histogram of an image

- The histogram of an image represents the number of pixels having a given gray value.
- We see three zones of repartition: dark, intermediate, and bright.



Matlab
command
`imhist`




Histogram of an image

- The histogram of an image represents the number of pixels having a given gray value.
- We see a uniform repartition, except the central peak.



Matlab
command
`imhist`




Histogram of an image

- The histogram of an image represents the number of pixels having a given gray value.
- Although the image seems to be black, we see that there are a few gray values



Matlab
command
`imhist`



Histogram of an image

- The histogram does not convey any spatial information.



original image
(Michelangelo's David)



after thresholding



after Floyd-
Steinberg dithering

same histogram




Enhancement by histogram modification



darker image




brighter image





Histogram modification

- Linear rescaling of the range. The distance between the peaks remains constant.




higher contrast





Histogram equalization

- Non-linear rescaling of the range. The distance between the peaks is proportional to their heights.



higher contrast


Matlab
command
`histeq`



Histogram equalization

- Idea -- stretch histogram non-uniformly such that final histogram is a uniform distribution

$$p_s(s) = p_r(r) \left| \frac{dr}{ds} \right|$$
$$s = T(r) = \int_0^r p_r(w) dw$$




a b

FIGURE 3.18 (a) An arbitrary PDF. (b) Result of applying the transformation in Eq. (3.3-4) to all intensity levels, r . The resulting intensities, s , have a uniform PDF, independently of the form of the PDF of the r 's.

Histogram equalization

- theory developed for continuous histogram, reality deal with discrete approx.



a b c

FIGURE 3.19 Illustration of histogram equalization of a 3-bit (8 intensity levels) image. (a) Original histogram. (b) Transformation function. (c) Equalized histogram.

Histogram equalization




FIGURE 3.20 Left column: images from Fig. 3.16. Center column: corresponding histogram-equalized images. Right column: histograms of the images in the center column.

Histogram matching




FIGURE 3.21 Transformation functions for histogram equalization. Transformations (1) through (4) were obtained from the histograms of the images (from top to bottom) in the left column of Fig. 3.20 using Eq. (3.3-8).

Histogram matching

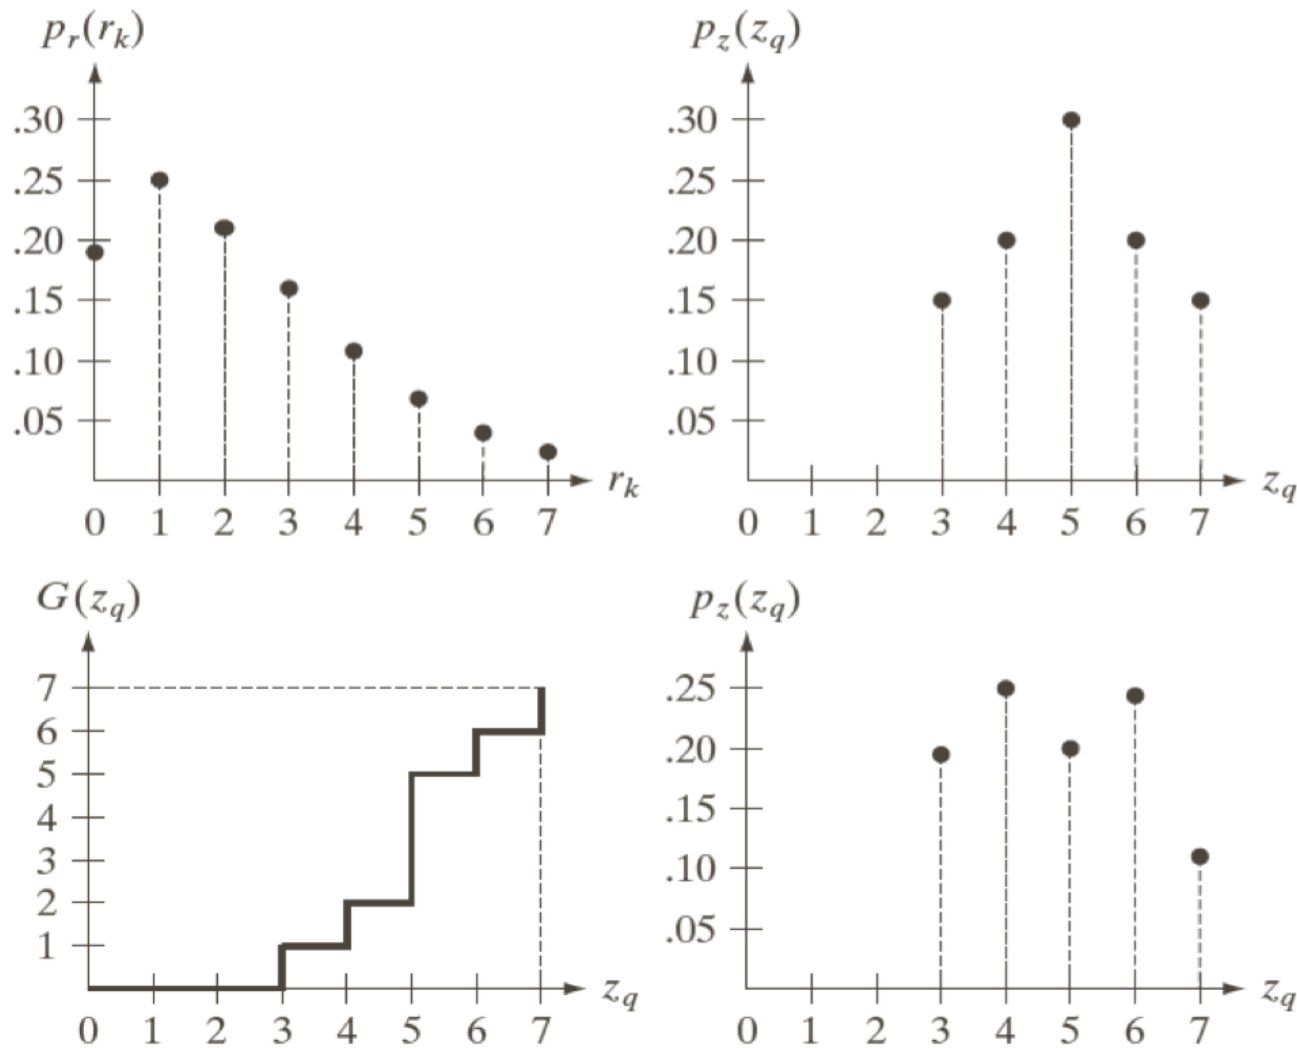
- Idea -- match the histogram to arbitrary function, instead of the constant function

$$s = G(z) = \int_0^r p_z(w) dw$$

$$s = T(r) = \int_0^r p_r(w) dw$$

$$z = G^{-1}[T(r)] = G^{-1}(s)$$








Histogram matching




| | |
|---|---|
| a | b |
| c | d |

FIGURE 3.22
 (a) Histogram of a 3-bit image. (b) Specified histogram. (c) Transformation function obtained from the specified histogram. (d) Result of performing histogram specification. Compare (b) and (d).

Histogram matching




Local histogram equalization



a b c

FIGURE 3.26 (a) Original image. (b) Result of global histogram equalization. (c) Result of local histogram equalization applied to (a), using a neighborhood of size 3×3 .

Local histogram equalization



a b c

FIGURE 3.27 (a) SEM image of a tungsten filament magnified approximately $130\times$. (b) Result of global histogram equalization. (c) Image enhanced using local histogram statistics. (Original image courtesy of Mr. Michael Shaffer, Department of Geological Sciences, University of Oregon, Eugene.)

Overview

- What is an image?
 - quantization
 - color image
 - image pixels
- Quality
- Intensity transforms
- Histogram processing
- Filtering
- Smoothing
- Sharpening

Mechanics of filtering




FIGURE 3.28 The mechanics of linear spatial filtering using a 3×3 filter mask. The form chosen to denote the coordinates of the filter mask coefficients simplifies writing expressions for linear filtering.

Mechanics of filtering

Correlation:

$$w(x, y) \star f(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot f(x + s, y + t)$$

Convolution:

$$w(x, y) * f(x, y) = \sum_{s=-a}^a \sum_{t=-b}^b w(s, t) \cdot f(x - s, y - t)$$

| | | |
|-------------|------------|------------|
| $w(-1, -1)$ | $w(-1, 0)$ | $w(-1, 1)$ |
| $w(0, -1)$ | $w(0, 0)$ | $w(0, 1)$ |
| $w(1, -1)$ | $w(1, 0)$ | $w(1, 1)$ |

Correlation vs. Convolution





FIGURE 3.29 Illustration of 1-D correlation and convolution of a filter with a discrete unit impulse. Note that correlation and convolution are functions of *displacement*.

Overview

- What is an image?
 - quantization
 - color image
 - image pixels
- Quality
- Intensity transforms
- Histogram processing
- Filtering
- Smoothing
- Sharpening

Smoothing linear filters



a b


FIGURE 3.32 Two 3×3 smoothing (averaging) filter masks. The constant multiplier in front of each mask is equal to 1 divided by the sum of the values of its coefficients, as is required to compute an average.

Smoothing linear filters

FIGURE 3.33 (a) Original image, of size 500×500 pixels. (b)–(f) Results of smoothing with square averaging filter masks of sizes $m = 3, 5, 9, 15,$ and $35,$ respectively. The black squares at the top are of sizes $3, 5, 9, 15, 25, 35, 45,$ and 55 pixels, respectively; their borders are 25 pixels apart. The letters at the bottom range in size from 10 to 24 points, in increments of 2 points; the large letter at the top is 60 points. The vertical bars are 5 pixels wide and 100 pixels high; their separation is 20 pixels. The diameter of the circles is 25 pixels, and their borders are 15 pixels apart; their intensity levels range from 0% to 100% black in increments of 20%. The background of the image is 10% black. The noisy rectangles are of size 50×120 pixels.




Smoothing non-linear filters



a b c

FIGURE 3.34 (a) Image of size 528×485 pixels from the Hubble Space Telescope. (b) Image filtered with a 15×15 averaging mask. (c) Result of thresholding (b). (Original image courtesy of NASA.)

Median filtering (denoising)



a b c

FIGURE 3.35 (a) X-ray image of circuit board corrupted by salt-and-pepper noise. (b) Noise reduction with a 3×3 averaging mask. (c) Noise reduction with a 3×3 median filter. (Original image courtesy of Mr. Joseph E. Pascente, Lixi, Inc.)

Overview

- What is an image?
 - quantization
 - color image
 - image pixels
- Quality
- Intensity transforms
- Histogram processing
- Filtering
- Smoothing
- Sharpening

Principle idea

- enhance / highlight transition in intensity
- how to find “transition”?
 - unsharp masking / highboost filtering
 - first / second order derivatives in 1D
 - multi-D:
 - gradient magnitude
 - Laplacian

Unsharp Masking / Highboost Filtering


- Blur image
- Subtract blurred image from original (--> mask)

$$g_{\text{mask}}(x, y) = f(x, y) - \bar{f}(x, y)$$

- add mask to original

$$g(x, y) = f(x, y) + k \cdot g_{\text{mask}}(x, y)$$

- unsharp masking: $k=1$
- highboost filtering: $k>1$



Unsharp Masking / Highboost Filtering

- Blur image
- Subtract blurred image from original (--> mask)

$$g_{\text{mask}}(x, y) = f(x, y) - \bar{f}(x, y)$$

- add mask to original

$$g(x, y) = f(x, y) + k \cdot g_{\text{mask}}(x, y)$$


- unsharp masking: $k=1$
- highboost filtering: $k>1$



a
b
c
d
e

FIGURE 3.40
(a) Original image.
(b) Result of blurring with a Gaussian filter.
(c) Unsharp mask. (d) Result of using unsharp masking.
(e) Result of using highboost filtering.

1st/2nd derivative - 1D example



a
b
c

FIGURE 3.36
Illustration of the first and second derivatives of a 1-D digital function representing a section of a horizontal intensity profile from an image. In (a) and (c) data points are joined by dashed lines as a visualization aid.

Multiple dimensions

- Laplacian $\Delta f = \nabla^2 f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$

- Approximation:

$$\frac{\partial^2 f}{\partial x^2} = f(x+1, y) + f(x-1, y) - 2f(x, y)$$

| | | | | | |
|----|----|----|----|----|----|
| 0 | 1 | 0 | 1 | 1 | 1 |
| 1 | -4 | 1 | 1 | -8 | 1 |
| 0 | 1 | 0 | 1 | 1 | 1 |
| 0 | -1 | 0 | -1 | -1 | -1 |
| -1 | 4 | -1 | -1 | 8 | -1 |
| 0 | -1 | 0 | -1 | -1 | -1 |

a b
c d


FIGURE 3.37

(a) Filter mask used to implement Eq. (3.6-6).

(b) Mask used to implement an extension of this equation that includes the diagonal terms.

(c) and (d) Two other implementations of the Laplacian found frequently in practice.

Laplacian



a
b c
d e

FIGURE 3.38

(a) Blurred image of the North Pole of the moon. (b) Laplacian without scaling. (c) Laplacian with scaling. (d) Image sharpened using the mask in Fig. 3.37(a). (e) Result of using the mask in Fig. 3.37(b). (Original image courtesy of NASA.)

Gradient magnitude

$$M(x, y) = \text{mag}(\nabla f) = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$


$$M(x, y) \approx \left|\frac{\partial f}{\partial x}\right| + \left|\frac{\partial f}{\partial y}\right|$$

| | | | |
|----|---|---|----|
| -1 | 0 | 0 | -1 |
| 0 | 1 | 1 | 0 |

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

Sobel

Gradient magnitude



a b

FIGURE 3.42
(a) Optical image of contact lens (note defects on the boundary at 4 and 5 o'clock).
(b) Sobel gradient.
(Original image courtesy of Pete Sites, Perceptics Corporation.)