Medical Image Processing

2 Image Representation

Prof. Dr. Marcin Grzegorzek

Research Group for Pattern Recognition
Institute for Vision and Graphics
University of Siegen, Germany

Table of Contents

Overview

Overview

Digital Image Representation

A (greyscale) digital image is a matrix of numbers.

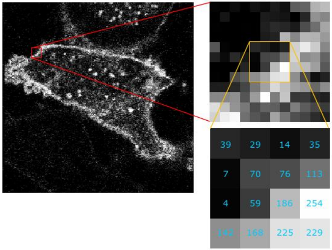

Each number is called a *pixel value* (pixel = picture element).

Usually the minimum number is shown black and the maximum white.

Notation

Notation for image pixels in this course:

The pixel values of image I are denoted I_i, with $i \in 1, 2, \ldots, n$, where n is the number of pixels in the image.

If spatial position matters, the pixel at row r and column c is denoted $I_{r,c}$ or $I[r, c]$.

Another convention (used e.g. in *ImageJ*) is to use x for c and y for r and reverse the order of the coordinates: $I_{x,y}$.

Notation (continued)

The first (topmost) row is row 0 and row numbers increase towards the bottom.
The first (leftmost) column is column 0 and column numbers increase towards the right.

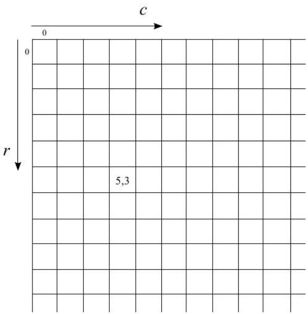

Digital Images - Spatial Resolution

The size of the image is the number of rows and columns of the matrix.

With n rows and m columns there are $n \times m$ pixels.

The (physical) *spatial resolution* is the number of pixels per length unit in the real world (e.g. pixels per $\mu\mathrm{m}$).

For cameras, the number of rows/columns or total number of pixels is given instead (e.g. 5 megapixels).

Digital Images - Bit Depth

The *bit depth* or *radiometric resolution* is the number of bits (0s and 1s) used to represent each pixel value.

Bits	**range**	**Notes**
1	0–1	Binary image
8	0–255	Typical greyscale image
12	0–4095	High quality greyscale
16	0–65535	Very high quality greyscale
32	(0.0-1.0)	Floating point format
8+8+8	3 × 0–255	“24 bit True Colour” (monitor)

Overview

Image Processing

Image ***processing*** takes an input image and produces an output image.

There is *no interpretation* of the image.

Examples: Noise reduction; visual enhancement; thresholding; edge detection; compression; texture detection...

Most image processing techniques are generic (not application specific).

Example Processing

Correct range of image and suppress noise.

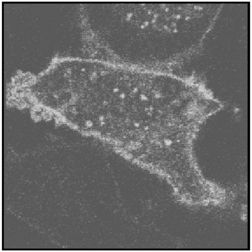

$\Rightarrow$

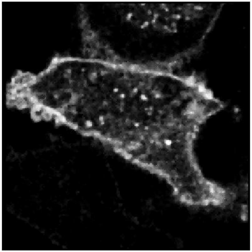

Image Analysis

Image ***analysis*** (IA) *interprets* images.

Examples: Find, count and measure objects; fit model to image data.

Will (almost) always use the spatial information in the image.

Often specific to a certain application or type of application.

IA systems are often created by IA experts in collaboration with application domain specialists.

Image Analysis Chain

Raw images

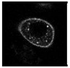

⇓ preprocessing, e.g. noise removal

Preprocessed images

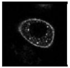

⇓ segmentation / detection

Segmented images / objects

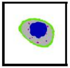

⇓ object level analysis

Information

???

Exemplary Image Analysis Project

Determining the thickness of subcutaneous fat at multiple sites from CT images.

Course project by Fintan J. McEvoy, KU-LIFE / IMHS.

The application is coded as a macro in ImageJ.

Exemplary Image Analysis Project

CT Image (abdomal slice) of a pig:

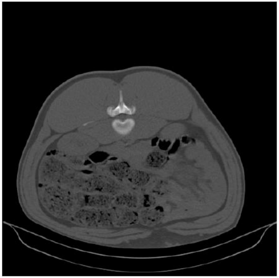

Exemplary Image Analysis Project

Step 1: Remove table and center the subject area.

Exemplary Image Analysis Project

Step 2: Edge detection along line to find transitions.

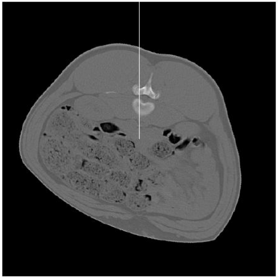

Repeat for n different angles.

Exemplary Image Analysis Project

Result:

Measurements exported in text file; can be imported by other software.

Exemplary Image Analysis Project

An IA system must be robust to variations in input.

(Position of subject on table; position of slice; ...)

This system can actually handle sheep as well as pigs!

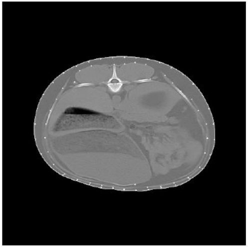

Overview

Image Acquisition Concerns

Most important prerequisite for succesful image analysis:
The information must be there!

Sufficient spatial resolution (and sharpness) is just one concern.

Radiometric Range

You want to assure the radiometric range is correctly calibrated.

Check the *histogram*!

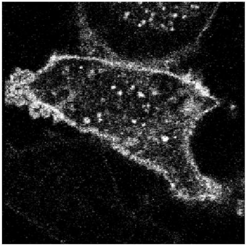

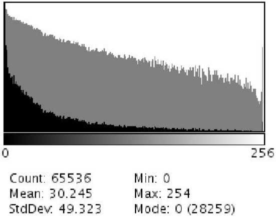

Example of almost perfect calibration.

Retain Radiometric Resolution

High quality cameras (and instruments) often have 12-bit radiometric resolution internally.

You may have to save in a special “raw” mode to avoid truncation to 8 bits.

8 bits are OK for computer monitors and print (even a little overkill).

For image analysis / processing we generally want to retain all the bits!

Image File Formats

Instrument software often has its own proprietary image file format.

If you want to use other software you must export/save in a standard format.

Popular formats: TIFF, PGM, PNM, DICOM...

I recommend using TIFF.

Image Compression

To save disk space, images may be compressed.

With *lossless* compression all information is retained.

With *lossy* compression only the “appearance” is retained.

Never use lossy compression!!!

JPEG is lossy! Don’t use it except for holiday snapshots.

Use TIFF instead. (Though not JPEG compressed TIFF...)

TIFF vs. JPEG (Grayscale)

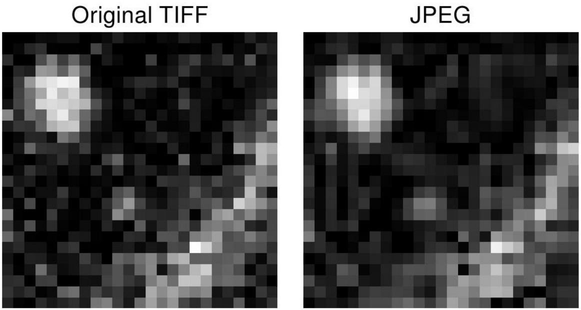

TIFF vs. JPEG (Colour)

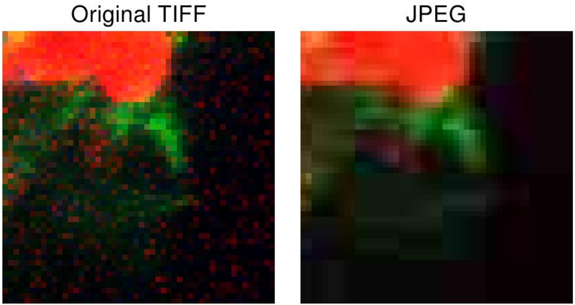

Overview

Pixel-wise Transformations (I)

In a *pixel-wise transformation*, pixel values are changed independently of their neighbourhood:

$$O_i = f(I_i)$$

Example: Pixel values scaling (normalization)

$$O_i = \frac{O_{\max}}{I_{\max} - I_{\min}}(I_i - I_{\min})$$

Example: Thresholding

$$O_i = \begin{cases} 255 & \text{if } I_i > T, \\ 0 & \text{otherwise.} \end{cases}$$

Example: Thresholding

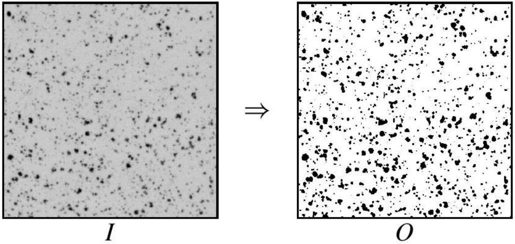

Pixel-wise Transformations (II)

Sometimes the pixel coordinates may be used in the computation:

$$O_{r,c} = f(I_{r,c}, r, c)$$

Example: Illumination correction

$$O_{r,c} = \frac{I_{r,c}}{\text{Illumination}(r, c)}$$

Example: Illumination Correction

Known (modelled) illumination:

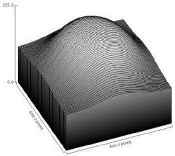

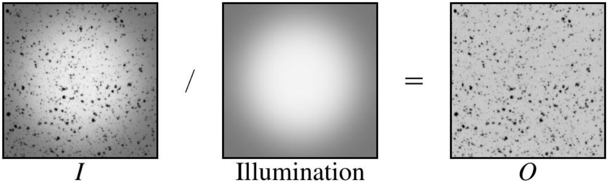

Pixel-wise Transformations (III)

It is also a pixel-wise transformation if the computation uses pixel values from the *same position* in multiple images (or bands):

$$O_{r,c} = f(I_{r,c}, J_{r,c}, K_{r,c}, r, c)$$

Example: Background subtraction

$$O_{r,c} = I_{r,c} - B_{r,c}$$

or (background normalization)

$$O_{r,c} = \frac{I_{r,c}}{B_{r,c}}$$

Example: Background Subtraction

Subtracting the background:

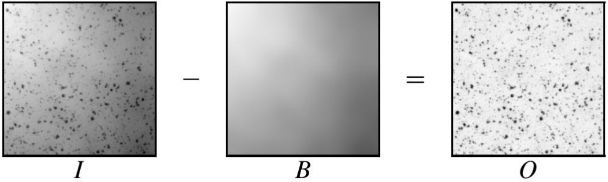

Example: Background Normalisation

Dividing with the background:

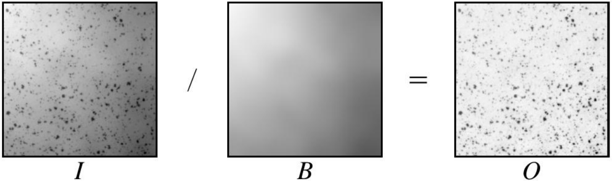

Overview

Visual Enhancement

Visual enhancement is done to assist human interpretation.

Typically you will want to maximize contrast.

Always keep the original images!

Always do processing / analysis on the original images...

In publications: Mention you have used enhancement.

Normalisation

Normalization linearly scales pixel values to use the full available range.

Original

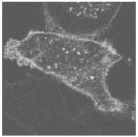

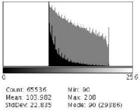

Normalized

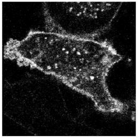

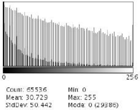

Histogram Equalisation

Histogram equalization re-maps pixel values to make the accumulated histogram approximately linear.

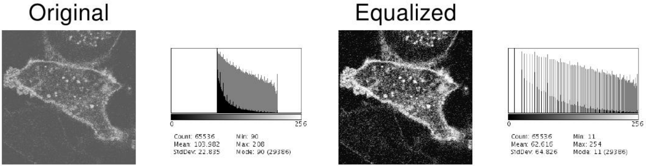

Look-up Tables

Visual enhancement can be done by applying a *look-up table* (LUT) mapping pixel values to displayed colour.

$$I_i \underset{\text{LUT}}{\longmapsto} O_i$$

Changing the LUT does not change the image pixel values, only the appearance of the image.

Inversion

Sometimes it is useful to invert the grey values (especially in print).

Original

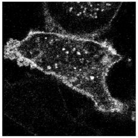

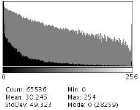

Inverted

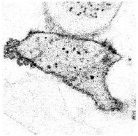

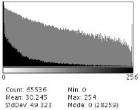

Colourisation

Sometimes colour may be useful (and/or eye-catching).

Original

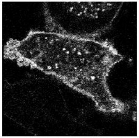

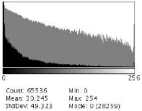

Colourized

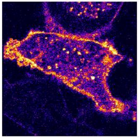

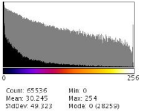

Humans can distinguish less than 100 grey levels but thousands of colours...