

Image Processing for Microscopy

1. Scientific image processing fundamentals
2. Acquiring images
3. Considerations when taking electron microscopy images
4. Correcting defects (noise etc.)
5. Image enhancement in the spatial domain
(on image pixels)
6. Processing images in frequency space (in Fourier space)
7. Extraction of (quantitative) information:
Segmentation and thresholding
8. Image measurements
9. Correlation, classification, identification and matching

Software Practice: ImageJ [or Fiji]






<http://imagej.net/ImageJ2>

Image Processing Lectures



- Today: lecture
- December 7th : 2 hours of lectures
& 1 hour of introduction to ImageJ/Fiji, installation of Fiji on your computer !
- December 14th : 3 hours of practicals with Fiji, questions etc.

Reading Material




Main Resources

- Seeing the scientific image, article as supporting material on moodle
- The **Image Processing Handbook** by John Russ and Brent Neal, CRC Press - 7th edition – Can be found ONLINE at the EPFL library
- Chapter 4 & Chapter 7 in Scanning Electron Microscopy and X-Ray Microanalysis by Joseph Goldstein et al., Springer – 3rd Edition – Hardcopy at the Library & at CIME library
- Chapter 7 in Transmission Electron Microscopy by David Williams and Barry Carter, Springer – Hardcopy at the Library & at CIME Library
- Image Analysis in Earth Sciences, Microstructures and Textures of Earth Sciences by Renne Heilbronn & Steve Barett, Springer 2014 – ONLINE at the EPFL library



1. Image Processing Fundamentals


Do you use image processing?



Two main purposes:

1. **improving** the visual **appearance** of images for a human observer, including their printing and transmission

Do you use image processing?



Two main purposes:


1. improving the visual appearance of images for a human observer, including their printing and transmission
2. preparing images for the **measurement** and/or **real-time analysis** of the **features** and **structures** that they reveal

A word of caution

A very real concern for everyone involved in scientific imaging, is the question of what constitutes **proper** and **appropriate** processing, and what constitutes **unethical** or even **fraudulent** manipulation


Anything that alters an image so as to create a **false impression** on the part of the viewer is wrong

1. Always store a permanent copy of the original image along with relevant data on its acquisition: **OpenData, OpenScience** !
2. Carefully document whatever steps are taken to process the image and generally to report those steps when the processed image is published !
Image processing has to be reproducible !



Examples of what image processing is NOT: 1. Particles

Published in Spectrochimica Acta



Oops...



**"Your X-ray showed a broken rib,
but we fixed it with Photoshop."**



Image processing for Microscopy






Image processing for Microscopy




2. Acquiring Images: SEM (Scanning Electron Microscopy)

Schematic illustration of the scanning system of the SEM with the available signals




2. Acquiring Images: TEM (Transmission Electron Microscopy)

Available signals in the TEM



TEM: Projected image (BF, DF, HRTEM, Diff)
STEM: Detector (BF, DF, HAADF)

Electrons are not visible to the human eye



Accelerating voltage kV	Velocity m/s	Uncorrected λ nm	Relativistic λ nm
1	1.873e+7	0.03879	0.03877
5	4.163e+7	0.01735	0.01730
10	5.846e+7	0.01227	0.01221
25	9.049e+7	0.00776	0.00766
50	1.237e+8	0.00549	0.00536
100	1.644e+8	0.00388	0.00370
200	2.085e+8	0.00274	0.00251
300	2.328e+8	0.00224	0.00197
400	2.482e+8	0.00194	0.00164
1,000	2.821e+8	0.00123	0.00087

Human color receptor relative sensitivity




Image Processing in Microscopes





Image Display


An image is a 2 dimensional collection of pixel values that can be displayed or printed by assigning a shade of grey or color to every pixel value.




For meaningful results, the spatial calibration of the image must be known

Calibration

For meaningful results, the spatial calibration of the image must be known




Example: Scanning Tunnelling Microscope



Horizontal scale: lateral dimensions
Vertical scale: height (surface)

The information on the image: Histogram

A monochrome (grayscale) image and its **brightness histogram**, which plots the number of pixels with each brightness value. The counts in each peak correspond to the areas of the regions of the image




Features on a noisy background:

(a) High **signal** to **noise** (S/N)

(b) Medium signal to noise ratio

(c) Low signal to noise ratio

Noise




What is an image?




Low intensity

$$I(x,y) = I_0 = 0$$



High intensity

$$I(x,y) = I_0 \neq 0$$



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

In order to get **information**, need variations in **intensity** from one point to another, i.e.,

Definition of CONTRAST

$$C = [I_2(x_2, y_2) - I_1(x_1, y_1)] / I_1(x_1, y_1)$$

Contrast vs. Intensity (Brightness)



contrast

High



Intensity/Brightness







Low





Contrast vs. Intensity (Brightness)

- Variations in Intensity with Position in the Object
- Use **Full Dynamic Range** of Scale: Don't Throw Away Information!
- Generally Better off with Lower Overall Intensity
 - Its easier to see small changes in intensity on a low background signal
 - Leads to Higher Contrast
- Can Boost Contrast Electronically






CBED



SAED



$$C = [I_2(x_2, y_2) - I_1(x_1, y_1)] / I_1(x_1, y_1)$$

$$C = \Delta I / I_0$$

Contrast vs. Intensity (Brightness)


Adjustment during live acquisition:
Use line scan and observe intensity profile

Bias:



Electronically shifts the signal up/down
Intensity, Brightness

Gain:

stretches the signal (contrast)



bias



gain




Image Formation: Line Scans

The information flow consists of the scan location in x-y space and a corresponding set of intensities from the detection device that monitor the beam-specimen interactions




Image Formation: Area Scans

The information flow consists of the scan location in x-y space and a corresponding set of intensities from the detection device that monitor the beam-specimen interactions





Image Formation: Parallel Illumination



All (x,y) values are acquired (recorded) at the same time (in parallel).

When

- Taking photographs
- Taking optical microscopy images
- Taking transmission electron microscopy images (shown is a TEM column with parallel illumination hitting a florescent screen)

Detectors vs. Cameras

Definitions (from Cambridge Dictionary)

Detector: a device used to find particular substances or things, or measure their level

Camera: a device for taking photographs or making films or television programs

Digital Camera: a type of camera that records images that can be look at on a computer



**in earlier days, images were taken in photographic film
analogic recording**

Basic difference

The signal from any electronic detector can be digitized and electronically manipulated prior to display, in a way that is impossible with analog images (images from film cameras or negatives)

Detector Properties: The Quantum Efficiency



To compare the properties of detection and recording devices, we use the concept of **detection quantum efficiency (DQE)**. For linear response detectors:

$$DQE = \frac{\left(\frac{S_{out}}{N_{out}} \right)^2}{\left(\frac{S_{in}}{N_{in}} \right)^2}$$


where S/N is the signal to noise ratio of the output or input signal

A perfect detector has a DQE of 1, but all practical detectors have a $DQE < 1$

- any detector adds a certain amount of (electronic) noise





weak signal amplified
noise amplified



too much amplified
signal cropped (cut)

Detector Properties... More


Linear response & Dynamic range




A detector generally meets the linearity condition for only a limited range of the input signal level. There are two effects that define the boundaries of this dynamic range

1. At very low levels of the input signal, the detector's output is largely dominated by noise
2. At high levels of the input signal, the detector's output signal no longer increases proportional to its input signal, and the detector therefore does not meet the linearity condition.

Linear vs Logarithmic Response



Secondary Electron Detectors in SEM



- Linear & fast response
- High gain
- Poor quantum efficiency

Everhart-Thornley Detector

Also Through The Lens (or In-Lens) Detector





Figure 1


Solid-State Detectors




For:

- Backscattered Electron detectors
- Xray Detectors
- Scanning Transmission Electron Microscopy Detectors (high angle annular dark field (HAADF), bright field, etc.

- Electrons amplified by high voltage and impact ionization
- High DQE
- Cannot be run too fast, cannot record too many electrons at the same time



Secondary Electron Detectors in ESEM





Amplification in native environment


- User has full control over the amplification conditions
- High noise
- Low DQE

Cameras: Charged Couple Device (CCD)

They are metal-insulator-silicon devices that store charge generated by light or electron beams




- Single read-out amplifier
- Transfers charge
- Slow
- High signal-to-noise ratio




“blooming” artefact:
overspill of charges

Cameras: Complimentary Metal-Oxide Semiconductors (CMOS)

The conversion of light photons to electrons functions the same way as in the CCD chip. The difference start with the way the signal is read out. Any pixel in the image that be read out directly, addressing a pixel by row and column as in a memory chip







- Each pixel has an amplifier
- Transfers voltage
- Fast
- Noisy



CCDs move photogenerated charge from pixel to pixel and convert it to voltage at an output node. CMOS imagers convert charge to voltage inside each pixel.

Read-out noise generated

Signal transfer in CCD/CMOS




Conversion of electrons into visible light
Scintillator, phosphorous screens

Image detection in electron microscopy


Conventionally

1. Convert electron to signal
2. Transfer signal
3. Detect signal with sensor
4. Electronically transfer signal and read out to form signal



Direct Electron Detection

1. -
2. -
3. Detect signal with sensor
4. Electronically transfer signal and read out to form signal
5. No spreading of signal
6. Single electron detection
7. Radiation damage !
8. Expensive (400-1000 kFr.)



3. Considerations when Taking Images, resolution

- **Microscope Resolution** is the minimum distance between two points or lines where they can be distinguished. Resolution is determined by the wavelength of the illuminating beam and the aperture of the collecting lens (numerical aperture.), aberrations in optics...
- **Computer image spatial resolution (Image Definition)** is measured (usually) in dots per inch (dpi) or pixels per inch.
- **Grayscale resolution (Dynamic Range)** is a binary measure of the number of possible gray values. Most image processing and image capture programs express grayscale in 8-bit format or 2^8 (256) levels of gray. For comparison, our eyes can distinguish only about 32 levels of gray.
- **Color resolution on a computer monitor** is usually a binary measure of 256 (2^8) possible values of each of R, G, and B. This results in a total of over 16 million colors ($256 \times 256 \times 256$), and is referred to as 24-bit color.
- **Printing resolution** is determined by the halftone screen frequency measured in lines per inch (lpi). This is a reflection of the size of the small, solid color dots in any printed image. A typical halftone screen frequency for newspapers is 65-85 lpi. Magazines print at 133 lpi, 600 dpi laser printers at 85 lpi, and high-quality publications print at 150-175 lpi.

Capturing Images: Image Digitalization


In order to get images into a form computers can **manipulate** they must be modified from their analog form (continuous gradients) into a digital form (step gradients) using the technique of **sampling**.

Sampling: A technique used to record analog information by recording periodic snapshots. If the sampling rate is fast enough, the eye cannot discern the gaps between each snapshot when they are played back. This is the principle behind motion pictures. Sampling is the key technique used to digitize analog information such as sound, photographs, and images.

Nyquist sampling $(f) = d/2$, where d =the smallest object, or highest frequency


The Nyquist Theorem states that in order to adequately reproduce a signal it should be periodically sampled at a rate of at least 2x the highest frequency you wish to record. With images, frequency is related to structure size. Small structures are said to have a high frequency. Thus, the imaging sample rate (or pixel) size should be 1/2 the size of the smallest object you wish to record.

Sampling Examples




Oversampling

A High Sampling Rate = much greater than 2X the highest frequency. This is 'Oversampling' that, while not "bad" will take time and create a large digital file



Nyquist Sampling Rate

The minimum sample rate that captures the "essence" of the analog information



Undersampled

Low sampling rate produces results that report false information about the analog data; which does not represent the original. This phenomenon is called aliasing

Red: real curve Blue: sampled curve shape

Sampling (pixel size $\leftrightarrow$ object size)

Chess board pattern 5x5 pixels




Image pixels size: 5 pixel / square




Image pixels size: 5 pixel / square




Image pixels size: 4 pixel / square




Image pixels size: 3 pixel / square




Image pixels size: 2 pixel / square





Image pixels size: 1 pixel / square
Sampling artefact (aliasing effect) : interference pattern

Going back to basic concepts

Image Processing: **Changing images into numbers and manipulating the numbers**

Grey scale image	«hill plot» of data
Zoomed area	Pixel values as numbers



Pixels, Bits and Bytes

Pixel is the **picture element**, a finite space where the analog signal is deposited to put it in digital form. It is common for the pixel to be **square**, equal dimensions in horizontal and vertical dimensions, as it helps with image processing afterwards

In the computer, numbers are stored as **bits and bytes**. A bit can assume the value 0 or 1, i.e., two different values. Eight bits in a row, **1 byte**, can assume $2^8 = 256$ different values between 0 and 255. Using 12 bits or 16 bits to record gray values, it is possible to represent $2^{12} = 4,096$ or $2^{16} = 65,536$ different gray values, respectively.

Helpful to remember:

1byte = 8 bits


1KB (kilobyte) = 1024 bytes
= 2^{10} bytes

1k x 1k (camera) = 1024 x 1024 pixels (= 1 Megapixels)


2k x 2k (camera) = 2048 x 2048 pixels (= 4 Megapixels)

binning 2 : reducing the total number of pixels by 2x

binning 3 : reducing the total number of pixels by 3x



Example: Pixel Size





256 Steps - appears visually to be continuous



128 Steps



64 Steps



32 Steps



16 Steps

Example: Gray Levels

32 grey levels



16 grey levels



8 grey levels



4 grey levels



4. Correcting Defects (post-processing)

Image processing operations and procedures that can be applied to correct some of the **defects** or shortcomings in as-acquired images that may occur due to **imperfect detectors**, **limitations of the optics**, inadequate or nonuniform **illumination**, or an undesirable viewpoint.

These are corrections applied **after the image has been digitized and stored, and therefore are unable to deliver the highest quality result that could be achieved by optimizing or correcting the acquisition process in the first place. Nothing can replace a “perfect” image !**



However, acquiring the optimum quality image can be impractical – reasons

- low dose imaging – noisy
- intensity variations – irregularities of surfaces
- outliers – x-rays, dead pixels


Color Adjustments: Look Up Table (LUT)

Image processing refers changing all or some of the pixel values in an image, usually with the aim of making some features of the image more easily “visible”

The easiest way would be to change the color used to represent each pixel value – the **Look-Up Table (LUT)**
= **remapping of grey scale levels**



default greyscale




increased contrast


Color Adjustments

Both linear and non-linear scales can be used. A non-linear mapping between the pixel value and displayed color can reveal unexpected features

Always have a idea of what you expect to see – do some research beforehand



default greyscale




zebra greyscale


Color Adjustments

Changing the LUT is **reversible**, as it is only mapping between pixel values and display colors that is changed

Taking a differential – replacing each pixel with the value of the local differential of the surface with respect to some direction – is **irreversible** in the sense that integrating doesn't necessarily get you your original image back



greys $\rightarrow$ z values




greys $\rightarrow$ $\frac{\partial z}{\partial x}$ values

Color Adjustments

All images in microscopy are taken in grey levels, however, you can also apply color to them if needed

There are many color spaces (chromaticity diagrams) – the most commonly used is Red – Green – Blue (RGB)



The neutral greyscale axis is running diagonally from black to white and has equal amounts of red, green, and blue


Noise in Images

The term noise refers to **random** fluctuations in pixel values that arise from the characteristics of the image acquisition.

This includes


- the **statistical** variations in the number of electrons that are **emitted** (and hitting the specimen),
- the number of electrons each sample **produces**, plus thermal electrons,
- the effects of capacitance and other **electronic** sources of variation in **amplification** and **digitization** process.

De-noising is often carried out using a **kernel filter** which uses a $n \times n$ matrix of elements. The kernel matrix is applied to every pixel in the image in turn




Neighborhood Averaging

The most common way to accomplish neighborhood averaging is to replace each pixel with the average of itself and its neighbors – sum of the pixel values in a region multiplied by an array of numeric weights



Neighborhood Averaging vs Binning


Neighbourhood averaging with 3x3 kernel



Binning 2

a

166	89	72	89
102	49	39	51
70	34	31	62
44	29	24	44



102	62
44	40

Reduced number of pixels

Binning 4

a

166	89	72	89
102	49	39	51
70	34	31	62
44	29	24	44



62

Reduced number of pixels

Neighborhood averaging

3x3 binning
(=resampling)

Kernel of 3x3

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & \color{red}{1} & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

Kernel of 7x7

$$\begin{bmatrix} 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & \color{red}{1} & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 \end{bmatrix}$$




Figure 4.14 Noise reduction by averaging. (a) original image; (b) averaging in 3×3 blocks; (c) 3×3 averaging around each pixel; (d) 7×7 averaging around each pixel.

Apparent features -> artifacts





Figure 4.19 Artifacts due to smoothing. Applying an 11 pixel wide smoothing kernel to the test pattern in (a) produces apparent lines between the original features (b).

Gaussian Smoothing

This is a set of weights that approximates the profile of a Gaussian function along any row, column, or diagonal through the center

$$G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} \cdot \exp\left(\frac{-(x^2 + y^2)}{2\sigma^2}\right)$$




increasing sigma


i.e. smoothing proceeds with a weighted average convolution

$$\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$


3x3 kernel



(a)



(b)



(c)


5x5 kernel

7x7 kernel

Gaussian Smoothing

This is a set of weights that approximates the profile of a Gaussian function along any row, column, or diagonal through the center

$$G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} \cdot \exp\left(\frac{-(x^2 + y^2)}{2\sigma^2}\right)$$



increasing sigma

In ImageJ for the Gaussian Blur filter -> 'Radius' means the radius of decay to $\exp(-0.5) \sim 61\%$ i.e, the standard deviation, sigma (in pixels).

In practice, the Gaussian distribution is effectively zero more than about three standard deviations from the mean, and so the kernel is truncated at this point

For example, for a Gaussian with sigma equal to 1, the kernel is approximated by a 7x7 matrix

7x7 smoothing kernel with Gaussian weights

0.0	0.00307	0.00676	0.00860	0.00676	0.00307	0.0
0.00307	0.01106	0.02355	0.02991	0.02355	0.01106	0.00307
0.00676	0.02355	0.04937	0.06350	0.04937	0.02355	0.00676
0.00860	0.02991	0.06350	0.08152	0.06350	0.02991	0.00860
0.00676	0.02355	0.04937	0.06350	0.04937	0.02355	0.00676
0.00307	0.01106	0.02355	0.02991	0.02355	0.01106	0.00307
0.0	0.00307	0.00676	0.00860	0.00676	0.00307	0.0

For example, for a Gaussian with sigma equal to 1, the kernel is approximated by a 7x7 matrix

Discrete approximations of Gaussian kernels

3x3, 5x5, 7x7

1/16

1	2	1
2	4	2
1	2	1

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

1/1003

0	0	1	2	1	0	0
0	3	13	22	13	3	0
1	13	59	97	59	13	1
2	22	97	159	97	22	2
1	13	59	97	59	13	1
0	3	13	22	13	3	0
0	0	1	2	1	0	0

One dimensional convolutions of Gaussian blur




Figure 4.17 Performing a Gaussian smooth by applying two one-dimensional convolutions.





Figure 4.18 Gaussian smoothing of the image from **Figure 4.14a**: (a) standard deviation = 1.0 pixels; (b) standard deviation = 3.0 pixels.

Neighborhood Averaging vs Median

The most common way to accomplish neighborhood **averaging** is to replace each pixel with the average of itself and its neighbors. Calculated value !

Median: replace the central pixel value by the **central value** of an **ordered array** of all values in the neighbourhood. Replace with a existing (most central) value.




Neighborhood **Ranking** – Median Filtering


Smoothing filters reduce random noise, but the underlying assumption is that all of the pixels in the neighborhood represent multiple samples of the same value -> that they belong to the same structure or object – but what happens at the edges or boundaries?

- Ranking of pixel according to their brightness level in their neighborhood
- The median value is used as the new value for the central pixel

The median filter can dramatically reduce the noise if it is “speckle” noise.



Salt and pepper
(black and white pixels)



Random noise

Neighborhood **Ranking** – Median Filtering


Smoothing filters reduce random noise, but the underlying assumption is that all of the pixels in the neighborhood represent multiple samples of the same value -> that they belong to the same structure or object – but what happens at the edges or boundaries?

- Ranking of pixel according to their brightness level in their neighborhood
- The median value is used as the new value for the central pixel

The median filter can dramatically reduce the noise if it is “speckle” noise.




Salt and pepper & median




Random noise & median

Gaussian Smoothing vs Median Filtering


SEM image




Gaussian
sigma=2 pixels



Median
radius=2pixels



Median
radius=15pixels



Median filters preserve **edges** in general better than gaussian or averaging filters!

Iterative Median (several steps)




Figure 4.24 Iterated median filtering: (a) original noisy image; (b) iterated median using a single large (7 pixel radius) neighborhood; (c) iterated median using a sequence of 3, 5, and 7 pixel radius neighborhoods.

- Fine details are erased and large regions take on the same brightness
- This is also called contouring
- First one to improve signal-to-noise and second one (larger) to demarcate the boundaries

Hybrid Median (or corner-preserving median)








Figure 4.30 Diagram of neighborhood pixels used in the 3×3 hybrid median filter. Both groups include the central pixel and are ranked separately. The median of each group and the central pixel are then ranked again to select the final median value.

- Multistep ranking median
- Preserves lines and corners that are erased or rounded by the conventional median

- a) Original
- b) Radius 2 median: rounded corners
- c) Radius 2 hybrid median: preservation of corners
- d) Radius 4 hybrid median



Several Other Defect/Noise Removal Methods

- Conditional neighborhoods: excludes pixels from the summation or from the ranking if they are very different from the original central pixel
- Interpolation: filling the region of defects using adjacent values – does not always match the surrounding
- Maximum entropy: alter the pixel brightness values to maximize the entropy in the image based on prior knowledge on the image formation process
- Maximum likelihood: it assumes that the boundaries between regions should be sharp – whichever region has the lowest variance is taken to represent the region to which the central pixel should belong and it is assigned the mean value from that region
- Also, adjustments form non-uniform illumination, background correction and image alignment can be performed