Fundamentals
of
Digital Image Processing

Royal Microscopical Society
MRC National Institute for Medical Research

David Holburn University Engineering Department, Cambridge

Cambridge University
Engineering Department

Fundamentals
of
Digital Image Processing

Royal Microscopical Society
MRC National Institute for Medical Research

Web page:
`http://www-g.eng.cam.ac.uk/rms/ipworkshop`
For images, macros, filters and other downloads

David Holburn University Engineering Department, Cambridge

Cambridge University
Engineering Department

Digital Imaging

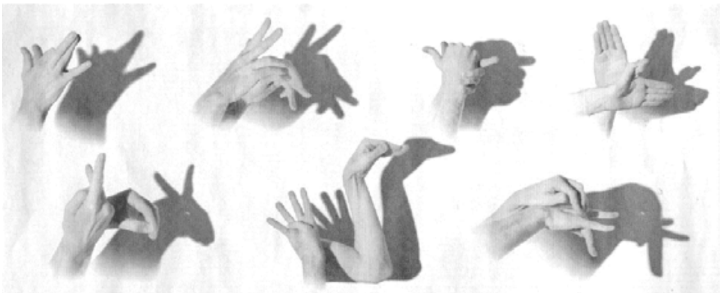

Digital Imaging has moved on a shade . . .

Cambridge University
Engineering Department

Why Digital Image Processing?

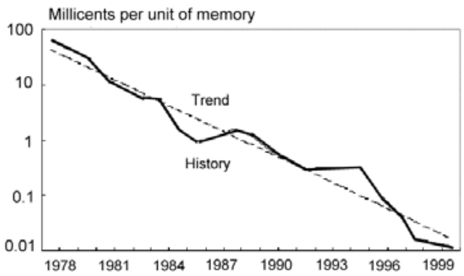

- This is the era of low-cost computer hardware
- Avoid human error
- Consistency, repeatability
- Allows diagnosis/analysis of images quantitatively
- Compensate for defects in imaging process (*restoration*)
- Certain techniques impossible any other way
- Speed & reduced specimen irradiation

Cambridge University
Engineering Department

Digital Images

- A natural image is a continuous, 2-dimensional distribution of brightness (or some other physical effect).
- Conversion of natural images into digital form involves two key processes, jointly referred to as digitisation:
  - Sampling
  - Quantisation
- Both involve loss of image fidelity i.e. *approximations*.

Cambridge University Engineering Department

Sampling

- Sampling represents the image by measurements at regularly spaced sample intervals. Two important criteria:-
  - Sampling interval
    - distance between sample points or pixels
  - Tessellation
    - the pattern of sampling points
- The number of pixels in the image is called the resolution of the image. If the number of pixels is too small, individual pixels can be seen and other undesired effects (e.g. aliasing) may be evident.

Cambridge University Engineering Department

Sampling

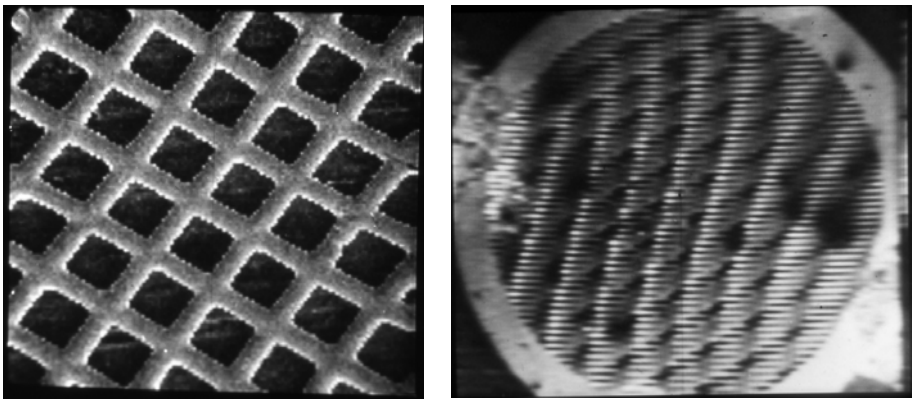

Metal grid in SEM – pixellation effects and moiré

Cambridge University
Engineering Department

Quantisation

- Quantisation uses an ADC (analogue to digital converter) to transform brightness values into a range of integer numbers, 0 to *M*, where *M* is limited by the ADC and the computer.

$$M = 2^m$$

where m is the number of bits used to represent the value of each pixel. This determines the number of grey levels.

- Too few bits results in steps between grey levels being apparent.

Cambridge University
Engineering Department

Quantisation

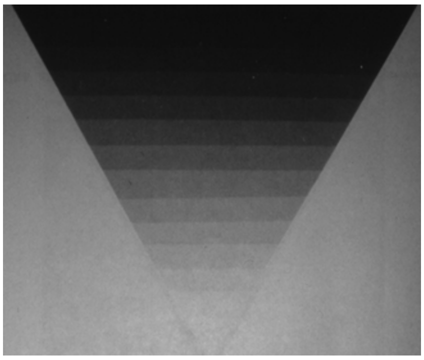

16 grey shades

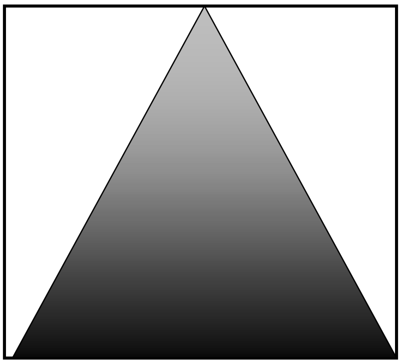

256 grey shades

Quantisation of grey levels:

Cambridge University
Engineering Department

Example

- For an image of 512 by 512 pixels, with 8 bits per pixel:

  Memory required = 0.25 megabytes

- Images from video sources (e.g. video camera) arrive at 25 images, or frames, per second:

  Data rate = 6.55 million pixels per second

- The capture of video images involves large amounts of data occurring at high rates.

Cambridge University
Engineering Department

Why Use a Framestore?

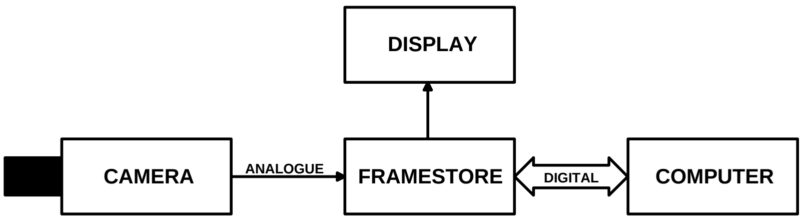

Cambridge University
Engineering Department

Framestore Structure

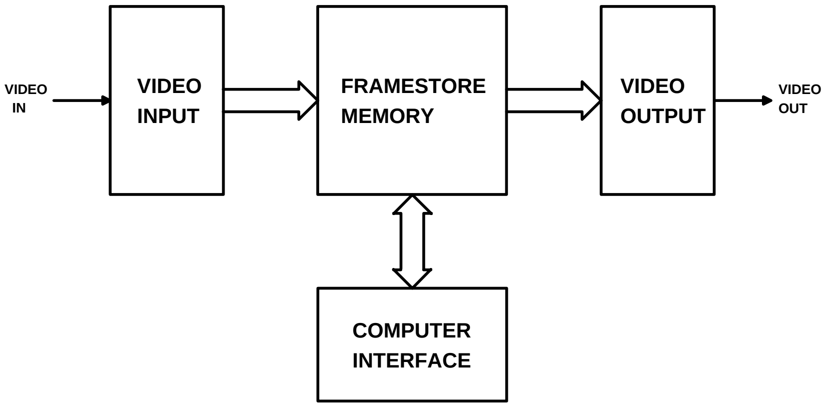

Cambridge University
Engineering Department

Framestore Memory Accesses

The framestore must be accessed in 3 ways:

- Capturing image, over 6.5 million accesses/sec.
- Displaying image, over 6.5 million accesses/sec.
- Computer access, over 1 million accesses/sec.

- The framestore must be able to be accessed over 14 million times per second .
- Conventional memories can only handle a maximum of 4-8 million accesses per second.

Cambridge University
Engineering Department

Basic operations

- The grey level histogram
- Grey level histogram equalisation
- Point operations
- Algebraic operations

Cambridge University
Engineering Department

The Grey-level Histogram

- One of the simplest, yet most useful tools.
- Can show up faulty settings in an image digitiser.
- Almost impossible to achieve without digital hardware.

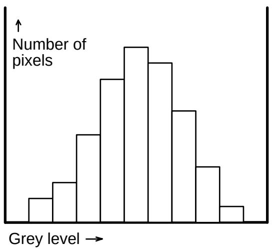

The GLH is a function showing for each grey level the number of pixels that have that grey level.

Cambridge University
Engineering Department

Correcting digitiser settings

- Inspection of the GLH can show up faulty digitiser settings

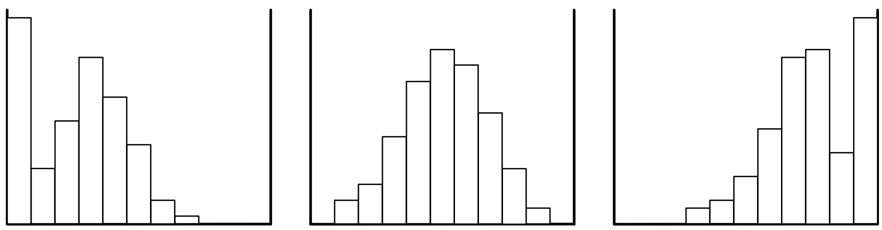

(a) Insufficient signal (b) Correct adjustment (c) Saturation

Cambridge University
Engineering Department

Image segmentation

- The GLH can often be used to distinguish simple objects from background and determine their area.

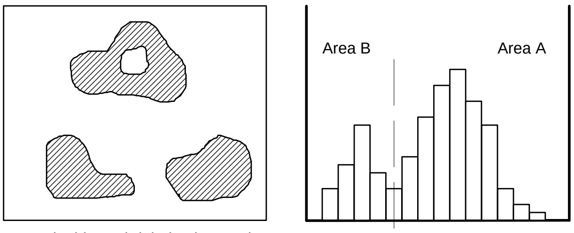

Dark objects, bright background

Cambridge University
Engineering Department

Grey Level Histogram Equalisation

- In GLH equalisation, a non-linear grey scale transformation redistributes the grey levels, producing an image with a flattened histogram.
- This can result in a striking contrast improvement.

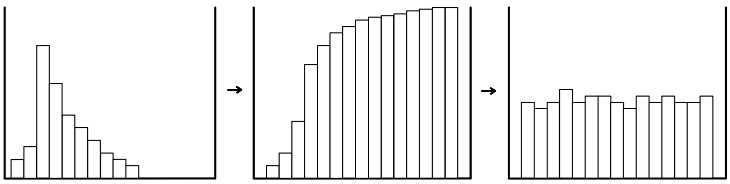

Original image

Cumulative GL histogram

Flattened histogram

Cambridge University
Engineering Department

Grey Level Histogram Equalisation

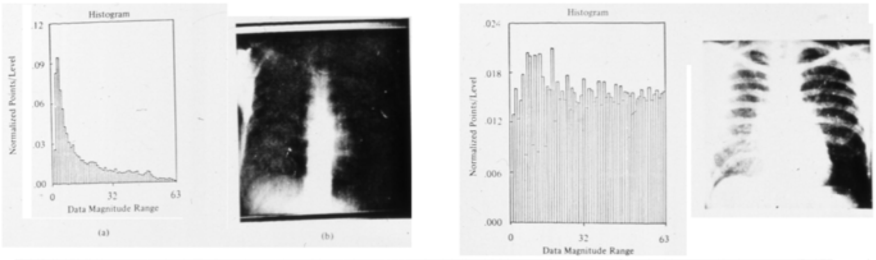

Figure 4.11 Pictorial example of histogram equalization (a chest x ray): (a) histogram of original chest x ray, (b) original chest x ray, (c) histogram of enhanced chest x ray, and (d) enhanced chest x ray.

Point Operations

- Point operations affect the way images occupy greyscale.
- A point operation transforms an input image, producing an output image in which each pixel grey level is related in a systematic way to that of the corresponding input pixel.

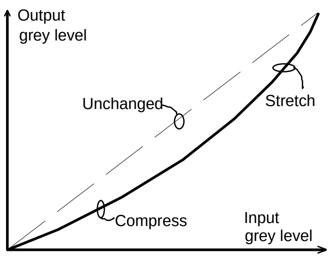

A point operation will never alter the spatial relationships within an image.

Cambridge University
Engineering Department

Examples of Point Operations

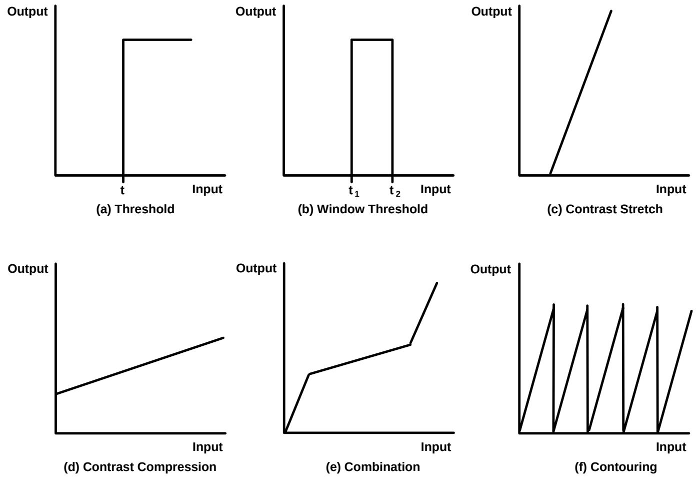

Cambridge University
Engineering Department

Examples of Point Operations

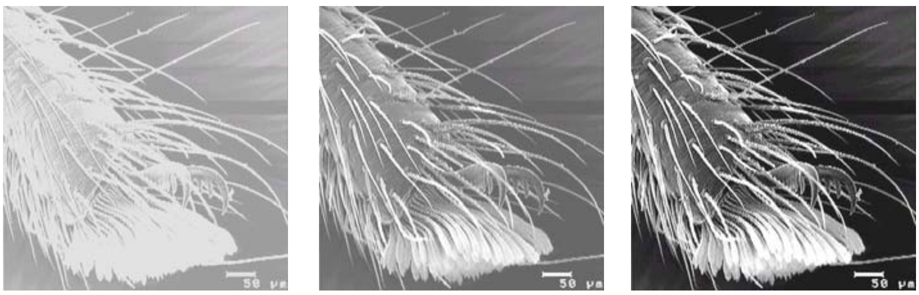

Cambridge University
Engineering Department

Algebraic operations

- A point form of operation (with >1 input image)
- Grey level of each output pixel depends only on grey levels of corresponding pixels in input images
- Four major operations:

– Addition:
$C = A + B$

– Multiplication:
$C = A \times B$

– Subtraction
$C = A - B$

– Division:
$C = A \div B$

Other operations can be defined that involve more than 2 input images, or using (for example) boolean or logic operators.

Cambridge University
Engineering Department

Applications of algebraic operations

- Addition
  - Ensemble averaging to reduce noise
  - Superimposing one image upon another
- Subtraction
  - Removal of unwanted additive interference (background suppression)
  - Motion detection

Cambridge University
Engineering Department

Applications of algebraic operations

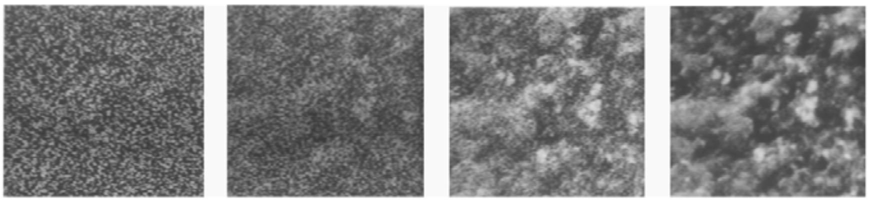

(a) 1 frame SNR=0.5 (b) 10 frames SNR=1.6 (c) 100 frames SNR=5.0 (d) 1000 frames SNR=16.0

Effect of frame averaging on Signal to Noise Ratio (SNR)

Applications (continued)

- Multiplication
  - Removal of unwanted multiplicative interference (background suppression)
  - Masking prior to combination by addition
  - Windowing prior to Fourier transformation
- Division
  - Background suppression (as multiplication)
  - Special imaging signals (multi-spectral work)

Cambridge University
Engineering Department

Look Up Tables

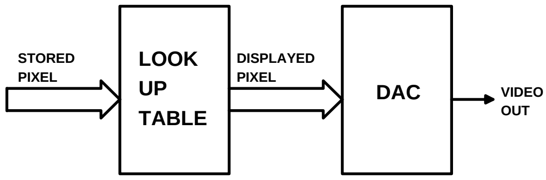

Cambridge University
Engineering Department

Pseudo-colour

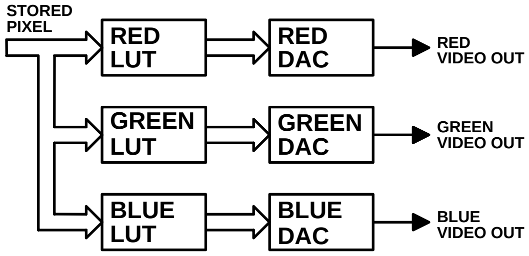

Cambridge University
Engineering Department

Using Pseudo-Colour

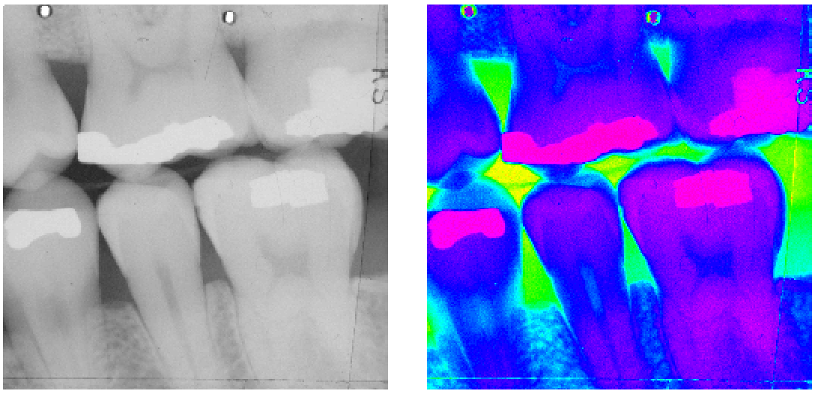

Pseudo-colour to accentuate subtle contrast variations

Cambridge University
Engineering Department

Noise Reduction

- An important noise reduction technique is frame averaging.
- A number of frames are averaged. The random noise is averaged out, resulting in a much improved image of the sample. Frame averaging may be written as:

$$y = \frac{1}{N}\sum_{n=1}^{N} x_n$$

where N is the number of images averaged, x are the images to be averaged, and y is the averaged image.

Cambridge University
Engineering Department

Frame Averaging contd.

- This has some disadvantages:
  - The averaged image is built up slowly. The display starts dark and gradually increases in brightness.
  - An output is only obtained once every N frames.

Cambridge University Engineering Department

Kalman Averaging

- The Kalman averager overcomes these problems by calculating the average of the frames input so far. The image displayed starts noisy and gradually improves as more frames are averaged. The Kalman filter is calculated by:

For the first frame: $y_1 = x_1$

For the second frame: $y_2 = \frac{1}{2}(x_1 + x_2) = \frac{1}{2}y_1 + \frac{1}{2}x_2$

For the third frame: $y_3 = \frac{1}{3}(x_1 + x_2 + x_3) = \frac{2}{3}y_2 + \frac{1}{3}x_3$

In general, ith frame: $y_i = \frac{i-1}{i}y_{i-1} + \frac{1}{i}x_i$

where x is the input image and y is the averaged image.

Cambridge University Engineering Department

Recursive Averaging

- The most useful averaging technique for microscopes is recursive averaging. The current displayed image is a combination of the current input image and the previous displayed image. This may be written as:

$$y_i = ky_{i-1} + (1-k)x_i$$

where $0<k<1$, x is the input image, y is the averaged image. The constant k can be considered as a time constant. The longer the time constant, the more the noise is reduced.

Cambridge University
Engineering Department

Dynamic algebraic operations

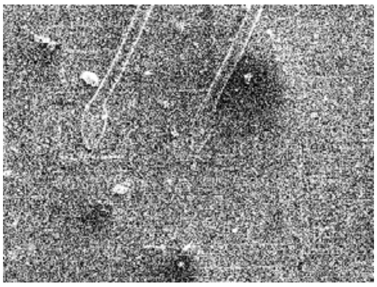

Recursive averaging A

Recursive averaging B

- Movies showing dynamic ensemble averaging

Cambridge University
Engineering Department

Background Shading Correction

- Background image q distorts the ideal microscope image p to give image x, the output from the camera. The distortion process is modelled by:

$$x = pq$$

To correct for the background distortion, the imaging system is uniformly illuminated. The ideal image p is now a constant C and the output from the camera x' is given by:

$$x' = Cq$$

Cambridge University
Engineering Department

Shading Correction contd.

- From this we can find the background image q:

$$q = \frac{x'}{C}$$

- To find an estimate of the ideal image p from the image x obtained from the camera we divide by q:

$$p = \frac{x}{q} = \frac{Cx}{x'}$$

Cambridge University
Engineering Department

Shading Correction contd.

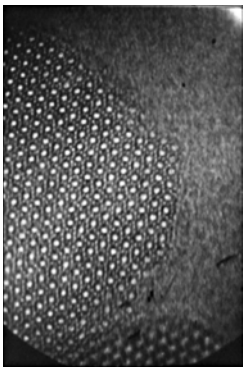

Original

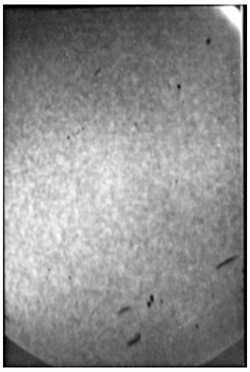

Background

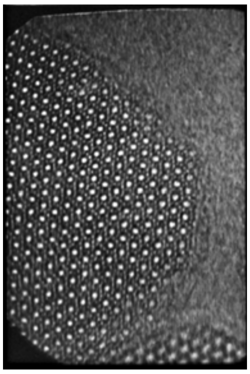

Compensated

Local Operations

- In a local operation, the value of a pixel in the output image is a function of the corresponding pixel in the input image and its neighbouring pixels. Local operations may be used for:-
  - image smoothing
  - noise cleaning
  - edge enhancement
  - boundary detection
  - assessment of texture

Local operations for image smoothing

- Image averaging can be described as follows:-

$$\frac{1}{4}\begin{array}{|c|c|c|}\hline 0 & 1 & 0 \\ \hline 1 & 0 & 1 \\ \hline 0 & 1 & 0 \\ \hline \end{array}$$

$h(x,y)$

The mask shows graphically the disposition and weights of the pixels involved in the operation.

Total weight = $\sum_{n} h(x,y) = 4$

Image averaging is an example of low-pass filtering.

Low Pass and Median Filters

- The low-pass filter can provide image smoothing and noise reduction, but subdues and blurs sharp edges.
- Median filters can provide noise filtering without blurring.

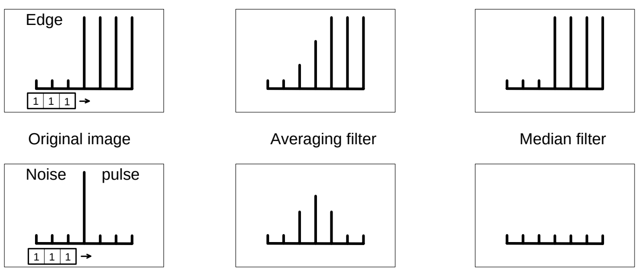

Low Pass and Median Filters

(a) Image with impulse noise 15 errors per line

(b) Median filtering of (a) with L = 3

(c) Median filtering of (a) with L = 5

(d) Median filtering of (a) with L = 7

Examples of 1-D median filtering for images corrupted by impulse noise

High Pass Filters

- Subtracting contributions from neighbouring pixels resembles differentiation, and can emphasise or sharpen variations in contrast. This technique is known as High Pass Filtering.

h1:

-1
1

h2:

1	-1

The simplest high-pass filter simulates the mathematical gradient operator:

$$G = \begin{bmatrix} df/dy \\ df/dx \end{bmatrix}$$

h1 gives the vertical, and h2 the horizontal component. The two parts are then summed (ignoring sign) to give the result.

High Pass Filters

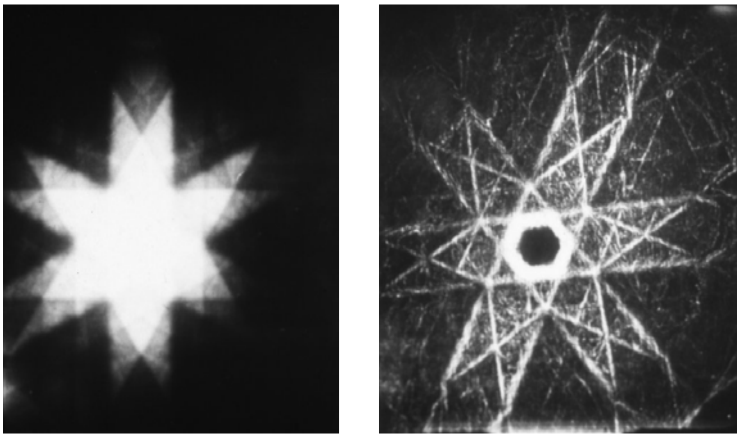

Electron channel patterning image in SEM (Si)

High Pass filtering for edge enhancement

Cambridge University
Engineering Department

Further examples of filters

- These masks contain 9 elements organised as 3 x 3. Calculation of one output pixel requires 9 multiplications & 9 additions. Larger masks may involve long computing times unless special hardware (a convolver) is available.

(a) Averaging (1/4 ×)

1	1	1
1	1	1
1	1	1

(b) Sobel

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

(c) Laplacian

0	1	0
1	-4	1
0	1	0

(d) High Pass

-1	-1	-1
-1	9	-1
-1	-1	-1

Cambridge University
Engineering Department

Frequency Methods

- Introduction to Frequency Domain
- The Fourier Transform
- Fourier filtering
- Example of Fourier filtering

Cambridge University
Engineering Department

Frequency Domain

- Frequency refers to the rate of repetition of some periodic event. In imaging, Spatial Frequency refers to the variations of image brightness with position in space.

- A varying signal can be transformed into a series of simple periodic variations. The Fourier Transform is a well known example and decomposes the signal into a set of sine waves of different characteristics (frequency and phase).

Cambridge University
Engineering Department

The Fourier Transform

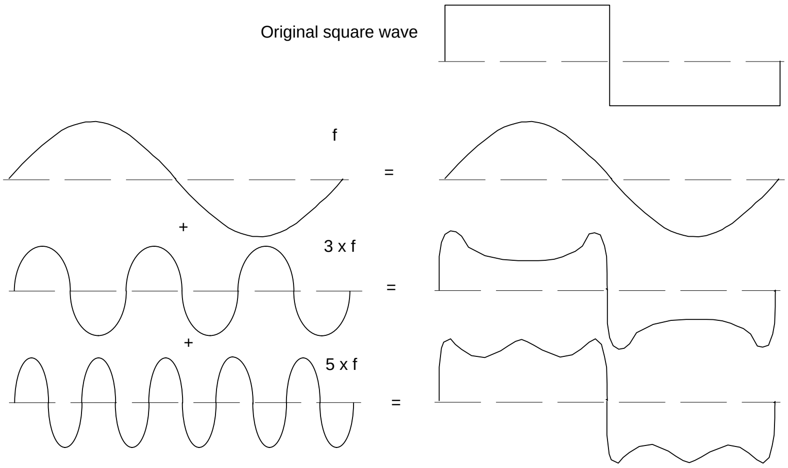

Amplitude and Phase

- The spectrum is the set of waves representing a signal as frequency components. It specifies for each frequency:
  - The amplitude (related to the energy)
  - The phase (its 'position' relative to other frequencies)

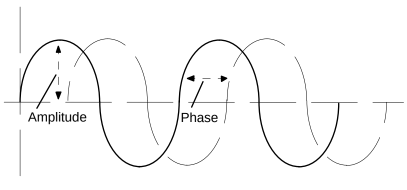

Amplitude and Phase

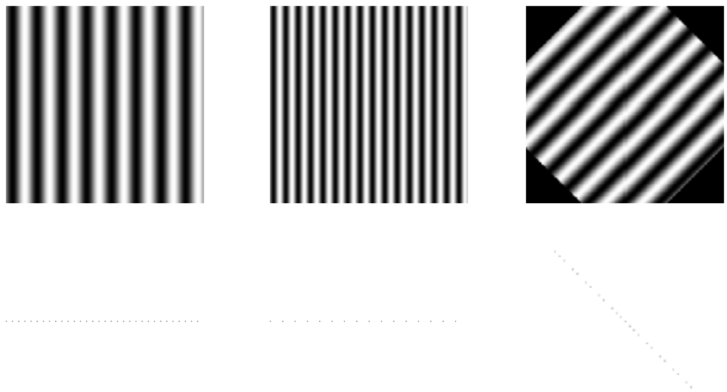

Fourier transforms of simple images

Cambridge University
Engineering Department

Fourier Filtering

- The Fourier Transform of an image can be carried out using:
  - Software (time-consuming)
  - Special-purpose hardware (much faster)
- using the Discrete Fourier Transform (DFT) method.
- The DFT also allows spectral data (i.e. a transformed image) to be inverse transformed, producing an image once again.

Cambridge University
Engineering Department

Fourier Filtering (continued)

- If we compute the DFT of an image, then immediately inverse transform the result, we expect to regain the same image.
- If we multiply each element of the DFT of an image by a suitably chosen weighting function we can accentuate certain frequency components and attenuate others. The corresponding changes in the spatial form can be seen after the inverse DFT has been computed.
- The selective enhancement/suppression of frequency components like this is known as Fourier Filtering.

Cambridge University
Engineering Department

Uses of Fourier Filtering

- Convolution with large masks (Convolution Theorem)
- Compensate for known image defects (restoration)
- Reduction of image noise
- Suppression of 'hum' or other periodic interference
- Reconstruction of 3D data from 2D sections
- Many others . . .

Cambridge University
Engineering Department

Image Analysis

Image analysis offers serious challenges …

Cambridge University
Engineering Department

Image Analysis

- Segmentation
  - Thresholding
  - Edge detection
- Representation of objects
- Morphological operations

Cambridge University
Engineering Department

Segmentation

The operation of distinguishing important objects from the background (or from unimportant objects).

- Point-dependent methods
  - Thresholding and semi-thresholding
  - Adaptive thresholding
- Neighbourhood-dependent
  - Edge enhancement & edge detectors
  - Boundary tracking
  - Template matching

Cambridge University Engineering Department

Point-dependent methods

- Operate by locating groups of pixels with similar properties
- **Thresholding**: assign a threshold grey level which discriminates between objects and background. This is straightforward if the image has a bimodal grey-level histogram.

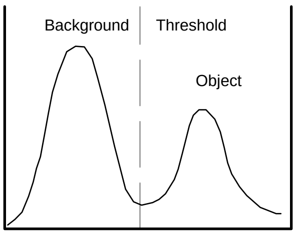

$$f_t = \begin{cases} 255 & f \geq t \\ 0 & f < t \end{cases}$$ (thresholding)

$$f_t = \begin{cases} f & f \geq t \\ 0 & f < t \end{cases}$$ (semi-thresholding)

Cambridge University Engineering Department

Thresholding example

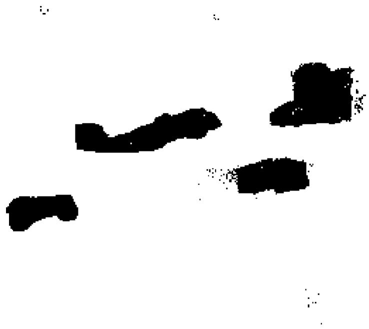

Segmentation of amalgam fillings in dental X-ray using thresholding

Cambridge University
Engineering Department

Adaptive thresholding

- In practice the GLH is rarely bimodal, owing to:-
  - Random noise - use LP/median or temporal filtering
  - Varying illumination
  - Complex images - objects of different sizes/properties

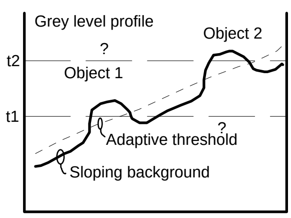

Background correction (subtract or divide) may be applied if an image of the background alone is available. Otherwise an adaptive strategy can be used.

Cambridge University
Engineering Department

Neighbourhood-dependent operations

- Edge detectors
  - Highlight region boundaries.
- Template matching
  - Locate groups of pixels in a particular group or configuration (pattern matching)
- Boundary tracking
  - Locate all pixels lying on an object boundary

Cambridge University
Engineering Department

Edge detectors

- Most edge enhancement techniques based on HP filters can be used to highlight region boundaries - e.g. Gradient, Laplacian. Several masks have been devised specifically for this purpose, e.g. Roberts and Sobel operators.
- Must consider directional characteristics of mask
- Effects of noise may be amplified
- Certain edges (e.g. texture edge) not affected

Cambridge University
Engineering Department

Template matching

- A template is an array of numbers used to detect the presence of a particular configuration of pixels. They are applied to images in the same way as convolution masks.

-1	-1	-1
-1	8	-1
-1	-1	-1

This 3x3 template will identify isolated objects consisting of a single pixel differing in grey-level from the background.

Other templates can be devised to identify lines or edges in chosen orientations.

Cambridge University Engineering Department

Boundary tracking

- Boundary tracking can be applied to any image containing only boundary information. Once a single boundary point is found, the operation seeks to find all other pixels on that boundary. One approach is shown:-

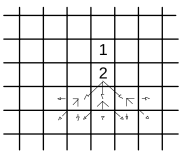

• Find first boundary pixel (1);

• Search 8 neighbours to find (2);

• Search in same direction (allow deviation of 1 pixel either side);

• Repeat step 3 till end of boundary.

Cambridge University Engineering Department

Connectivity and connected objects

- Rules are needed to decide to which object a pixel belongs.
- Some situations easily handled, others less straightforward.
- It is customary to assume either:
  - 4-connectivity
    - a pixel is regarded as connected to its four nearest neighbours
  - 8-connectivity.
    - a pixel is regarded as connected to all eight nearest neighbours

	1	
2	X	0
	3	

4-connected pixels

3	2	1
4	X	0
5	6	7

8-connected pixels

Cambridge University Engineering Department

Connected components

Object - Shaded

Background - Blank

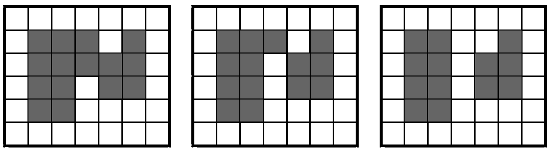

Connectivity	A. Objects found	B. Objects found	C. Objects found
4	1	2	2
8	1	1	2

Results of analysis under 4- or 8- connectivity

- A hidden paradox affects object and background pixels

Cambridge University Engineering Department

Line segment encoding

- Objects are represented as collections of chords
- A line-by-line technique
- Requires access to just two lines at a time

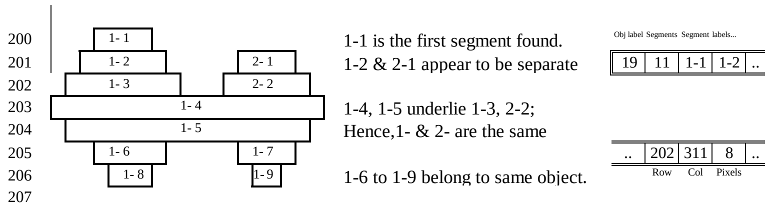

- Data compression may also be applied
- Feature measurement may be carried out simultaneously

Cambridge University
Engineering Department

Representation of objects

Object membership map (OMM)

An image the same size as the original image

- Each pixel encodes the corresponding object number,
  e.g. all pixels of object 9 are encoded as value 9
- Zero represents background pixels

Example OMM

	1	1	1	1		2	2	2	2										
	1	1	1			2	2		2										
						2													
												3	3						
									3	3	3	3	3						
			4	4						3	3	3	3						
		4	4	4						3			3	3			5	5	
		4		4									3			5	5	5	
		4																	

- ♦ requires an extra, full-size digital image
- ♦ requires further manipulation to yield feature information

Cambridge University
Engineering Department

Representation of objects

A compact format for storing object information about an object

- Defines only the position of the object boundary
- Takes advantage of connected nature of boundaries.
- Economical representation; 3 bits/boundary point
- Yields some feature information directly

  - Choose a starting point on the boundary (arbitrary)
  - One or more nearest neighbours must also be a boundary point
  - Record the direction codes that specify the path around the boundary

3	2	1
4	■	0
5	6	7

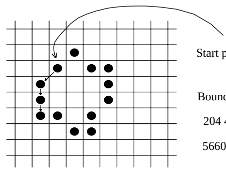

Cambridge University
Engineering Department

Size measurements

Area

- A simple, convenient measurement, can be determined during extraction.
- The object pixel count, multiplied by the area of a single pixel.
  - Determined directly from the segment-encoded representation
  - Additional computation needed for boundary chain code.

Simplified C code example

```
a = 0;                                   // Initialise area to 0
x = n; y = n;                            // Arbitrary start coordinates
for (i=0; i<n; i++)
switch (c[i])                            // Inspect each element
{                                        // 0246 are parallel to the axes
  case 0: a -= y;    x++;      break;
  case 2:            y++;      break;
  case 4: a += y;    x--;      break;
  case 6:            y--;      break;
}
printf ("Area is %10.4f\n",a);
```

Cambridge University
Engineering Department

Integrated optical density (IOD)

Determined from the original grey scale image.

IOD is rigorously defined for photographic imaging
In digital imaging, taken as sum of all pixel grey levels over the object:

$$IOD = \sum_{y=0}^{y=y\max} \sum_{x=0}^{x=x\max} f(x,y).o(x,y)$$

where:

$$o(x,y) = \begin{cases} 1 & (x,y) \text{ lies within the object} \\ 0 & \text{elsewhere} \end{cases}$$

- may be derived from the OMM, LSE, or from the BCC.
- IOD reflects the mass or weight of the object.
- Numerically equal to area multiplied by mean object grey level.

Cambridge University
Engineering Department

Length and width

Straightforwardly computed during encoding or tracking.

- Record coordinates:
  - minimum x
  - maximum x
  - minimum y
  - maximum y
- Take differences to give:
  - horizontal extent
  - vertical extent
  - minimum boundary rectangle.

Cambridge University
Engineering Department

Perimeter

May be computed crudely from the BCC simply by counting pixels
More accurately, take centre-to-centre distance of boundary pixels
For the BCC, perimeter, P, may be written

$$P = N_E + \sqrt{2}N_O$$

where:-

- N_E is the number of even steps
- N_O is the number of odd steps

taken in navigating the boundary.

- Dependence on magnification is a difficult problem
- Consider area and perimeter measurements at two magnifications:
  - Area will remain constant
  - Perimeter invariably increases with magnification
- Presence of holes can also affect the measured perimeter

Cambridge University
Engineering Department

Number of holes

Hole count may be of great value in classification.

A fundamental relationship exists between:-

- the number of connected components C (i.e. objects)
- the number of holes H in a figure
- and the *Euler number*:-

$$E = C - H$$

A number of approaches exist for determining H.

- Count special motifs (known as *bit quads*) in objects.

  These can give information about:-
  - Area
  - Perimeter
  - Euler number

Cambridge University
Engineering Department

Bit-quad codes

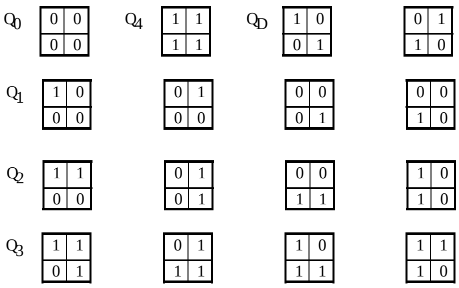

$$A = \tfrac{1}{4}n(Q_1) + \tfrac{1}{2}n(Q_2) + \tfrac{7}{8}n(Q_3) + n(Q_4) + \tfrac{3}{4}n(Q_D)$$

$$P = n(Q_2) + \tfrac{1}{\sqrt{2}}\left(n(Q_1) + n(Q_3) + 2n(Q_D)\right)$$

$$E = \tfrac{1}{4}\left(n(Q_1) - n(Q_3)\right) - 2n(Q_D)$$

For 1 object alone, $H = 1 - E$

Disposition of the 16 bit-quad motifs

Equations for A, P and E

Cambridge University
Engineering Department

Derived features

For example, shape features

Rectangularity

Ratio of object area A to area A_E of minimum enclosing rectangle

- Expresses how efficiently the object fills the MER
- Value must be between 0 and 1.
- For circular objects it is
- Becomes small for curved, thin objects.

Aspect ratio

The width/length ratio of the minimum enclosing rectangle

- Can distinguish slim objects from square/circular objects

Cambridge University
Engineering Department

Derived features (cont)

Circularity

- Assume a minimum value for circular shape
- High values tend to reflect complex boundaries.

One common measure is:

$$C = P^2/A$$

(ratio of perimeter squared to area)

- takes a minimum value of 4π for a circular shape.
- Warning: value may vary with magnification

Cambridge University
Engineering Department

Derived measurements (cont)

***Boundary energy* is derived from the curvature of the boundary.**
Let the instantaneous radius of curvature be $r(p)$,p along the boundary.
The curvature function *K(p) is* defined:

$$K(p) = \frac{1}{r(p)}$$

This is periodic with period *P*, the boundary perimeter.
The average energy for the boundary can be written:

$$E = \frac{1}{P}\int_0^P \left|K(p)\right|^2 \mathrm{dp}$$

A circular boundary has minimum boundary energy given by:

$$E_0 = \left(\frac{2\pi}{P}\right)^2 = \left(\frac{1}{R}\right)^2$$

where R is the radius of the circle.

Cambridge University
Engineering Department

Erosion and Dilation

- **Erosion** and **dilation** and **skeletonisation** alter the diameters of structures in the image
- **Erosion** converts object pixels lying at the boundary of the object into background pixels – objects appear to be *eroded* or eaten away
- **Dilation** converts background pixels lying at the boundary of each object into object pixels – objects appear to be *dilated* or to grow
- **Skeletonisation** (or *medial axis transform)* converts all the objects into stick-like skeletons which represent the central axes of the limbs of each object

Cambridge University
Engineering Department

Erosion and dilation (cont)

- Used in combination, erosion and dilation offer ways of *filtering out* unwanted objects
- Diameters of wanted objects can be nearly preserved

- **Opening**: (**erode** then **dilate**)
  - Removes thin objects

- **Closing**: (**dilate** then **erode**)
  - Agglomerates clusters
  - Fills holes

Cambridge University
Engineering Department

Erosion and dilation (cont)

- Binary image

After erosion

Cambridge University
Engineering Department

Erosion and dilation (cont)

- Eroded image

After dilation (opening)

Cambridge University
Engineering Department

Skeletonisation

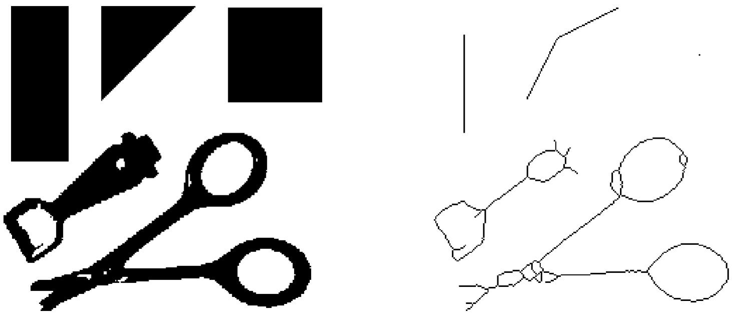

- A binary image before and after skeletonisation

Cambridge University
Engineering Department

3D Measurements

- Most imaging systems are 2D; many specimens are 3D.
- How can we extract the information?
- Photogrammetry - standard technique for cartography

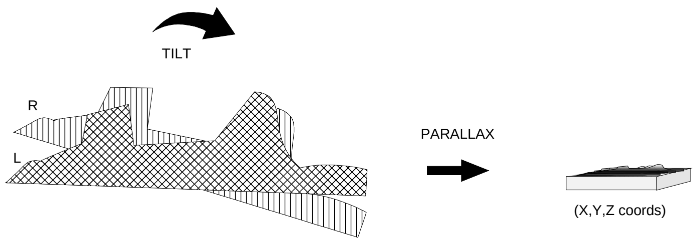

- Either the specimen, or the electron beam can be tilted.

Cambridge University
Engineering Department

Visualisation of height & depth

Seeing 3D images requires the following:-

- Stereo pair images
  - Shift the specimen (low mag. only)
  - Tilt specimen (or beam) through angle α
- Viewing system
  - lens/prism viewers
  - mirror-based stereoscope
  - twin projectors
  - anaglyph presentation (red & green/cyan)
  - LCD polarising shutter, polarised filters
- Stereopsis - ability to *fuse* stereo-pair images
- 3D reconstruction (using projection, Fourier, or other methods)

Cambridge University
Engineering Department

Anaglyph Stereo Images

Cambridge University
Engineering Department

Anaglyph Stereo Images

Cambridge University
Engineering Department

Anaglyph Stereo Images

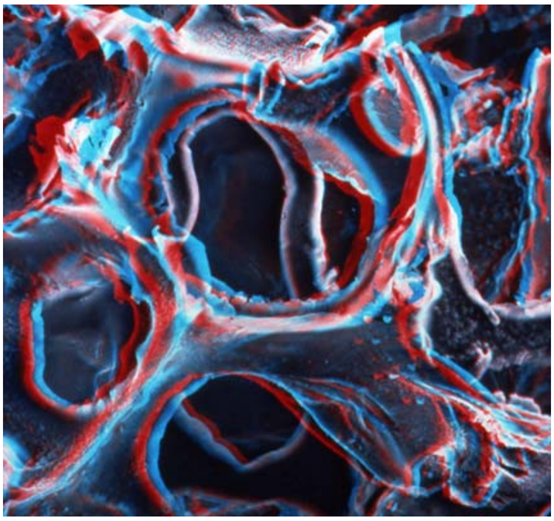

Cambridge University
Engineering Department

Anaglyph Stereo Images

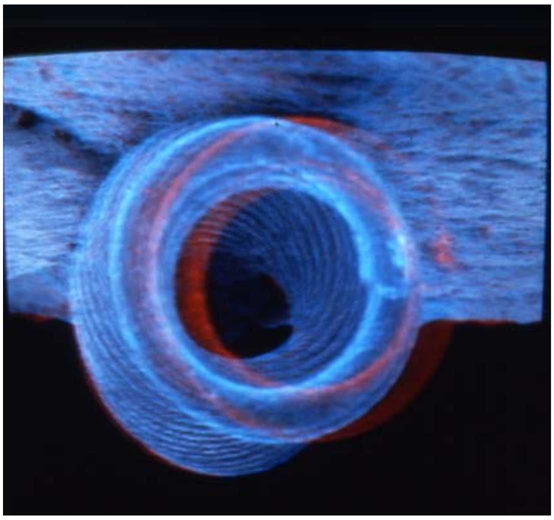

Cambridge University
Engineering Department

Other Applications

- Image restoration (compensate instrumental aberrations)
- Lattice averaging & structure determination (esp. TEM)
- Automatic focussing & astigmatism correction
- Analysis of diffraction (and other related) patterns
- 3D measurements, visualisation & reconstruction
- Analysis of sections (stereology)
- Image data compression, transmission & access
- Desktop publishing & multimedia

Cambridge University
Engineering Department

Remote Microscopy

- Modern SEMs are fully computer-controlled instruments
- Networking to share resources - information, hardware, software
- The Internet explosion & related tools
- Don’t Commute --- Communicate!

Cambridge University
Engineering Department

Remote Microscopy with NetSEM

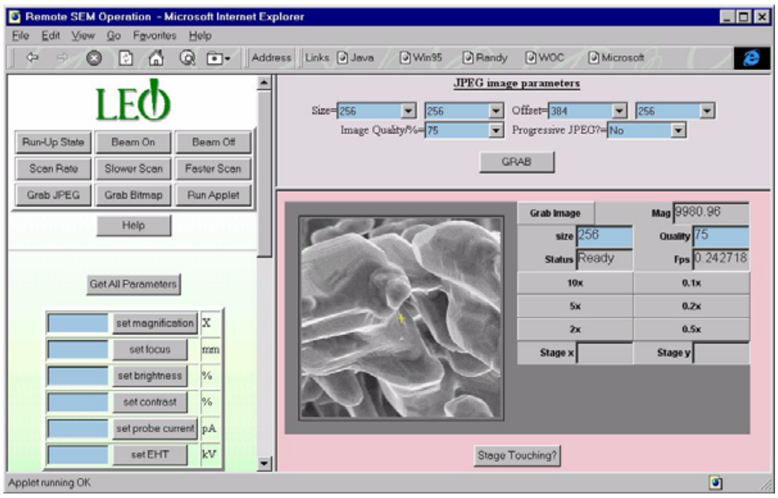

Cambridge University
Engineering Department

Automated Diagnosis for SEM

- Fault diagnosis of SEM
  - Too much expertise required
  - Hard to retain expertise
  - Verbal descriptions of symptoms often ambiguous
  - Geographical dispersion increases costs.
- Amenable to the **Expert System** approach.
  - A computer program demonstrating expert performance on a well-defined task
  - Should explain its answers, reason judgementally and allow its knowledge to be examined and modified

Cambridge University
Engineering Department

An Expert System Architecture

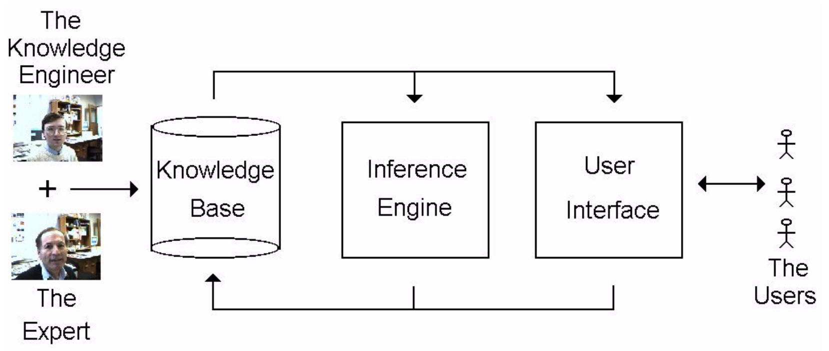

Cambridge University
Engineering Department

Remote Diagnosis

- Stages in development
  - Knowledge acquisition from experts, manuals and service reports
  - Knowledge representation --- translation into a formal notation
  - Implementation as custom expert system
  - Integration of ES with the Internet and RM
- Conclusions
  - RM offers accurate information and SEM control
  - ES provides engineer with valuable knowledge
  - ES + RM = Effective Remote Diagnosis