CSCI435/CSCI935
Computer Vision: Algorithms and Systems

Image Quality & Enhancement

Lecturer: Prof Lei Wang
Room 3.219
Email: leiw@uow.edu.au
Web: https://scholars.uow.edu.au/lei-wang

Machine Vision Concept (review)

- Machine Vision is a multistage process where each previous stage affects performance of all following stages

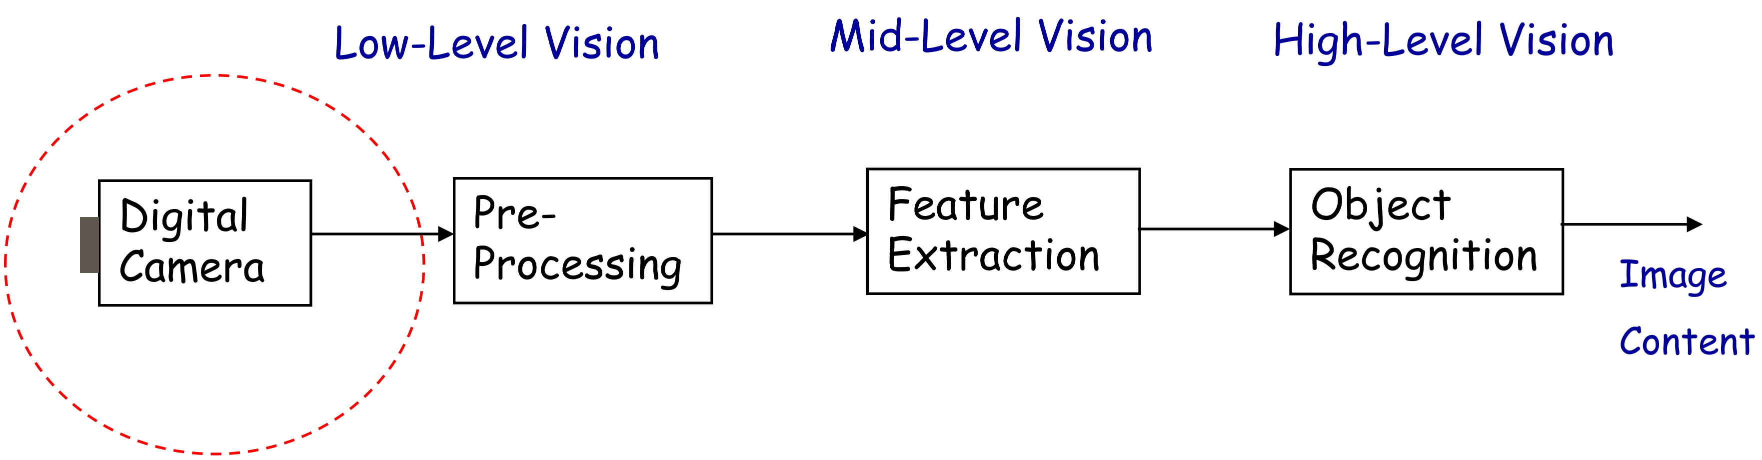

Single Sensor Digital Cameras (review)

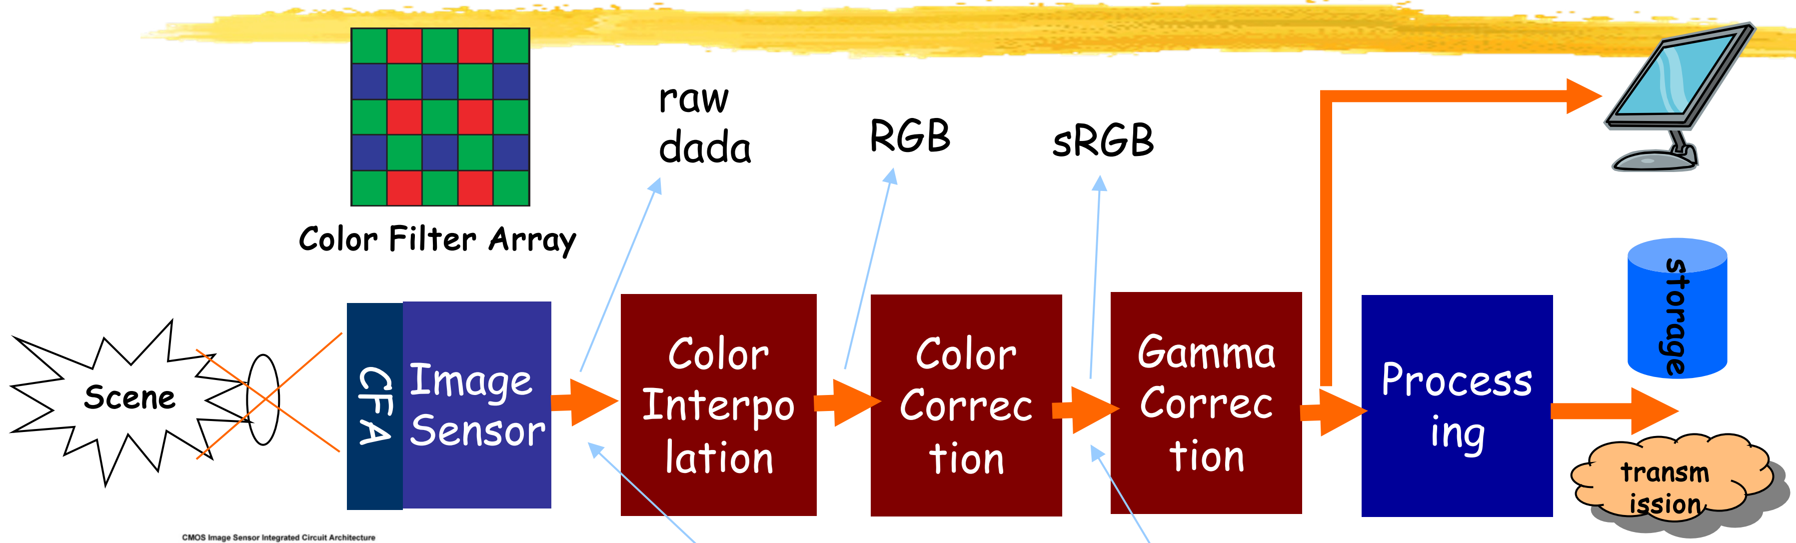

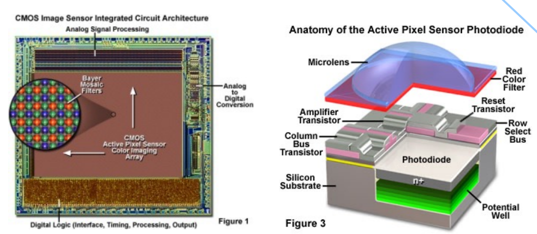

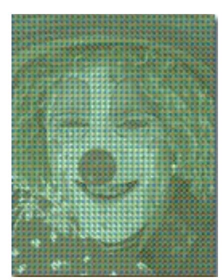

CCD - Charge-coupled devices
CMOS - Complementary
metal-oxide semiconductor

Machine Vision Concept (review)

- Machine Vision is a multistage process where each previous stage affects performance of all following stages

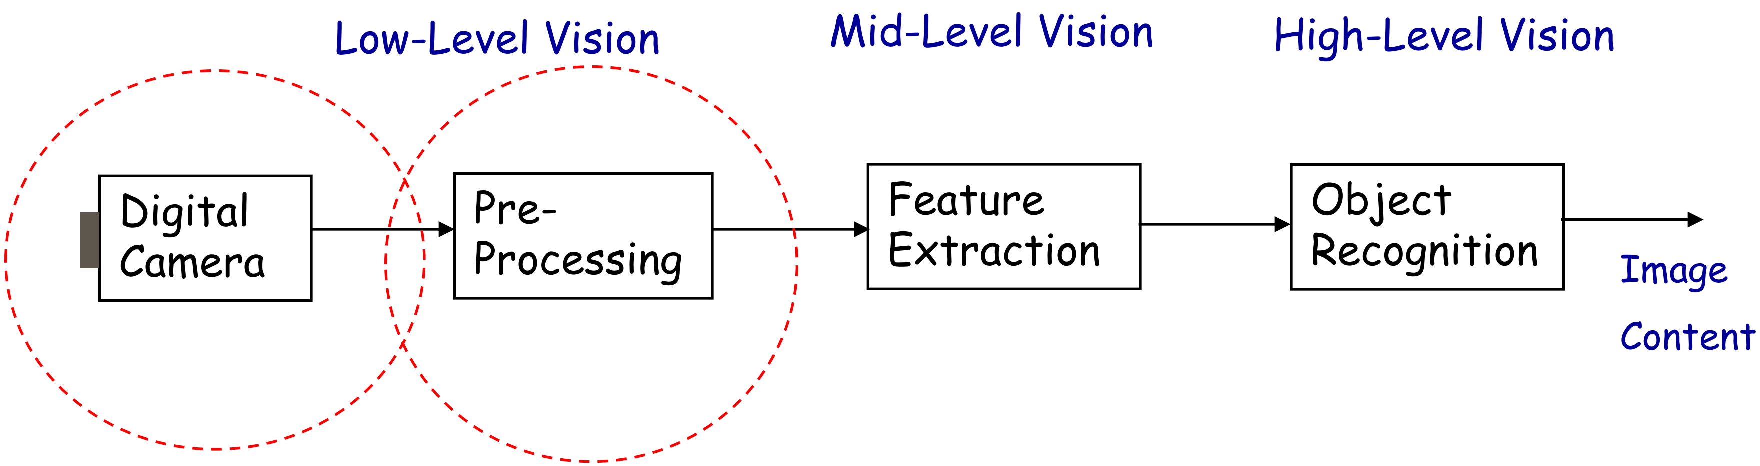

Image Enhancement

The objective of Image Enhancement is to produce an image that is more **suitable** than the original one

The word **suitable** has different meaning for different applications

Distortion & Criteria of Quality

- Image capturing process inevitably introduces distortions that degrade image quality
- Visual quality assessment is not sufficient for computer vision applications
- Image quality should be based on quantitative characteristics which in the end affect object recognition and measurement of parameters
- All distortions can be divided into two categories:
  - ✓ those which can be corrected by digital enhancement
  - ✓ those which cannot be corrected by digital enhancement and require optimisation of the image formation process

Sharpness

- Sharpness determine the amount of details that image can clearly reproduce
- Sharpness depends on several factors:
  - quality of the lens
  - focus accuracy
  - sensor resolution (sensor dimension in pixel)
  - CFA interpolation algorithm
  - blur due to handshaking
- Camera manufacturers are usually advertise only sensor resolution

Sharpness

- The measurable parameter of sharpness is how many black-white lines can be seen in vertical or horizontal direction

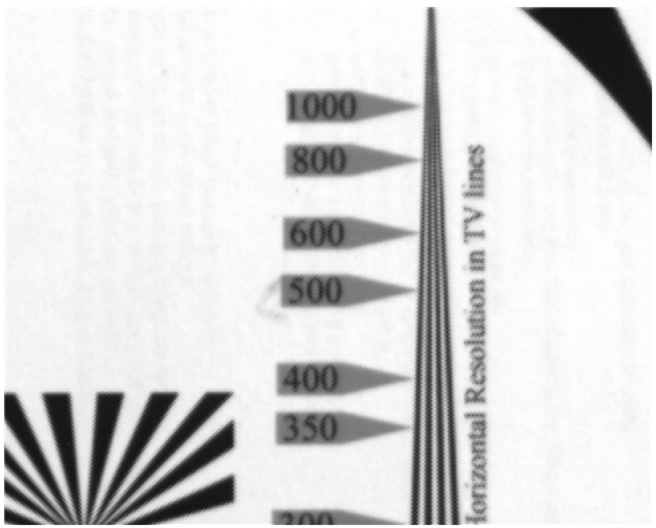

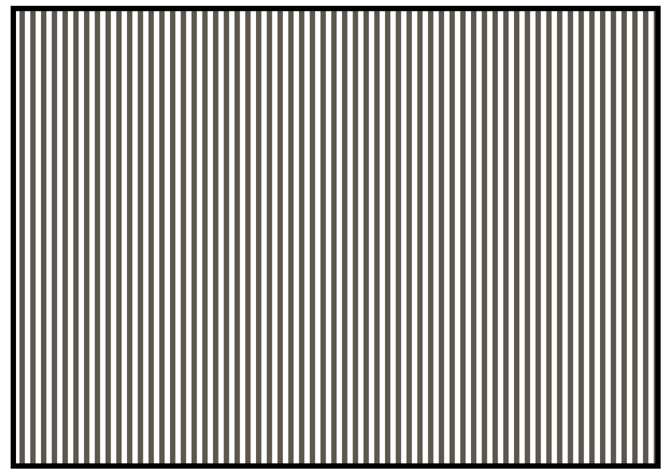

If a sensor size is 800 (cols)x600 (rows) pix, how many black-white lines can be clearly counted at best?

Sharpness

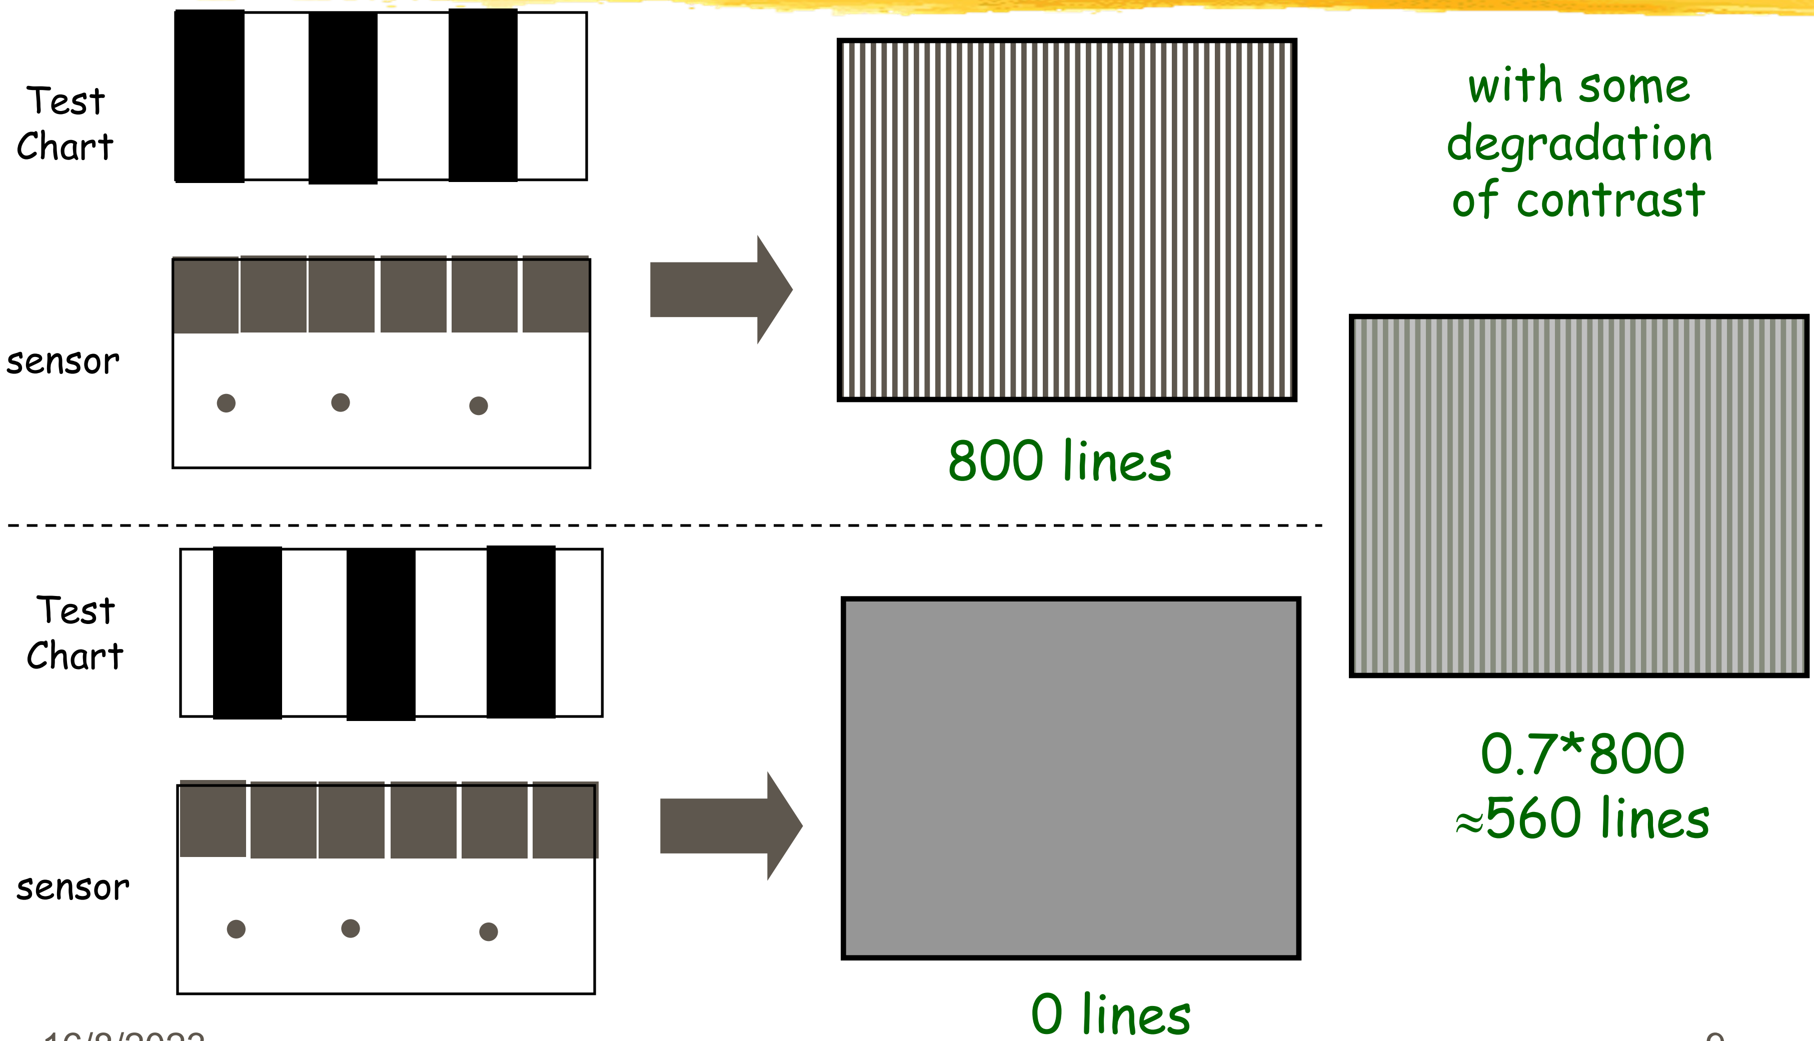
Test Chart
sensor
800 lines
with some degradation of contrast
Test Chart
sensor
0 lines
0.7*800
≈560 lines

Image Rescaling (Resizing)

- Some image analysis and object recognition operations do not require high resolution images
- Image size can be reduced through sub-sampling

Can a half size image be produced simply by discarding every second sample?

Image Rescaling (Resizing)

- Straightforward sub-sampling of an image will very likely violate requirements of the sampling theorem ($\Delta \leq 1/(2 * F_{max})$) introducing aliasing

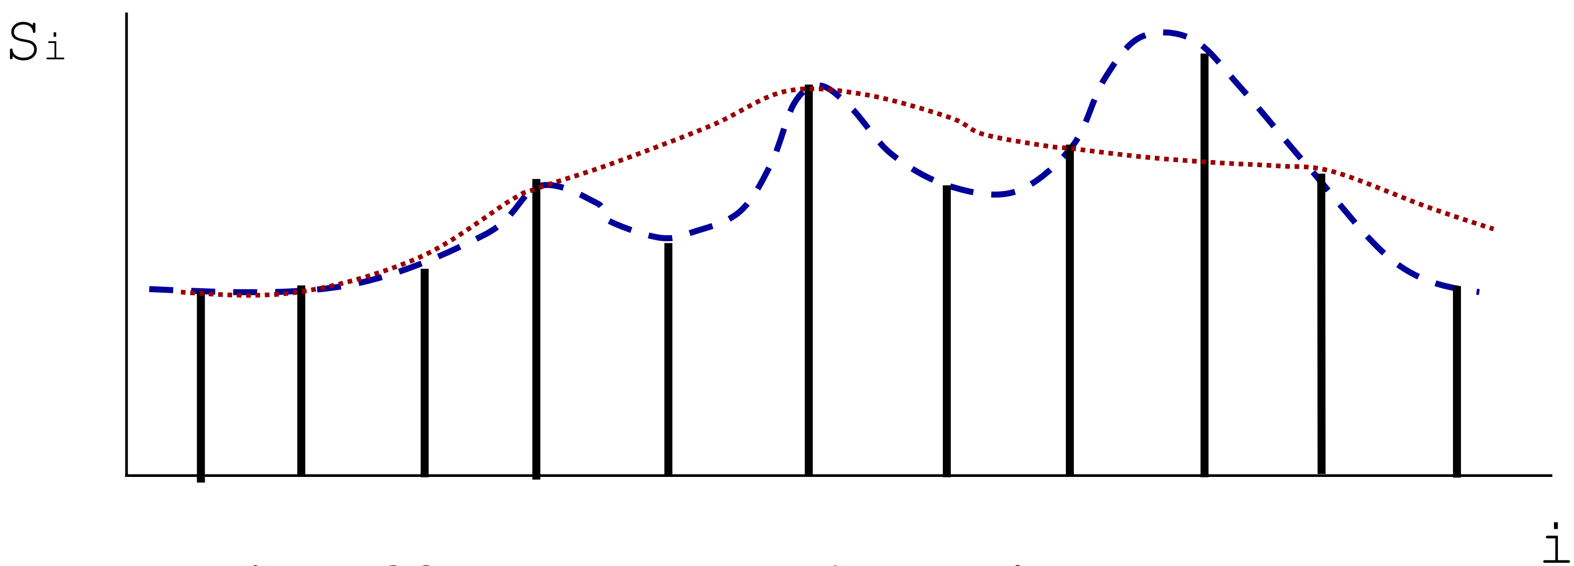

take off every second sample

Image Rescaling (Resizing)

- To satisfy the sampling theorem requirements, F_{max} of a sequence must be reduced to $F_{max}/2$ before down sampling

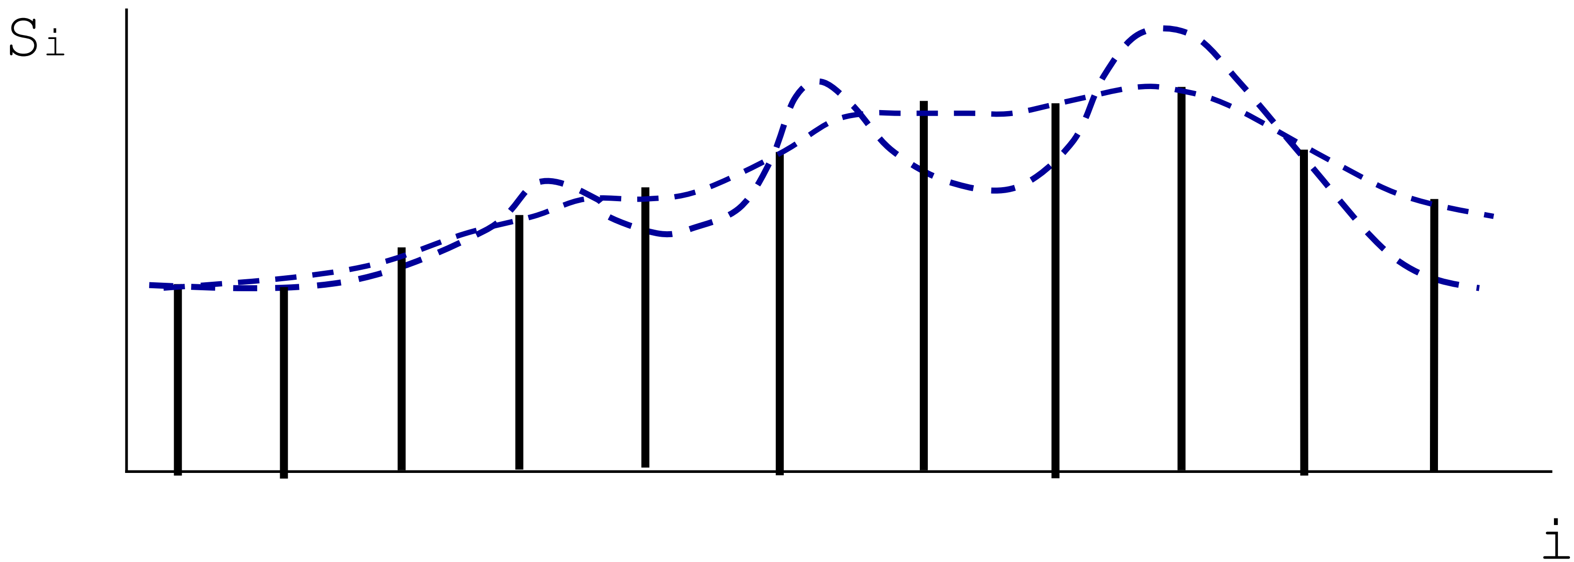

If $F_{max} = F_{max}/2$, down-sampling does not cause aliasing

Image Rescaling (Resizing)

- Image rescaling by factor of N is a two-stage process
  1. Limit the maximum frequency to F_{max}/N using an appropriate band limiting filter
  2. Decimate image by leaving only one sample out of N

Example: N = 2

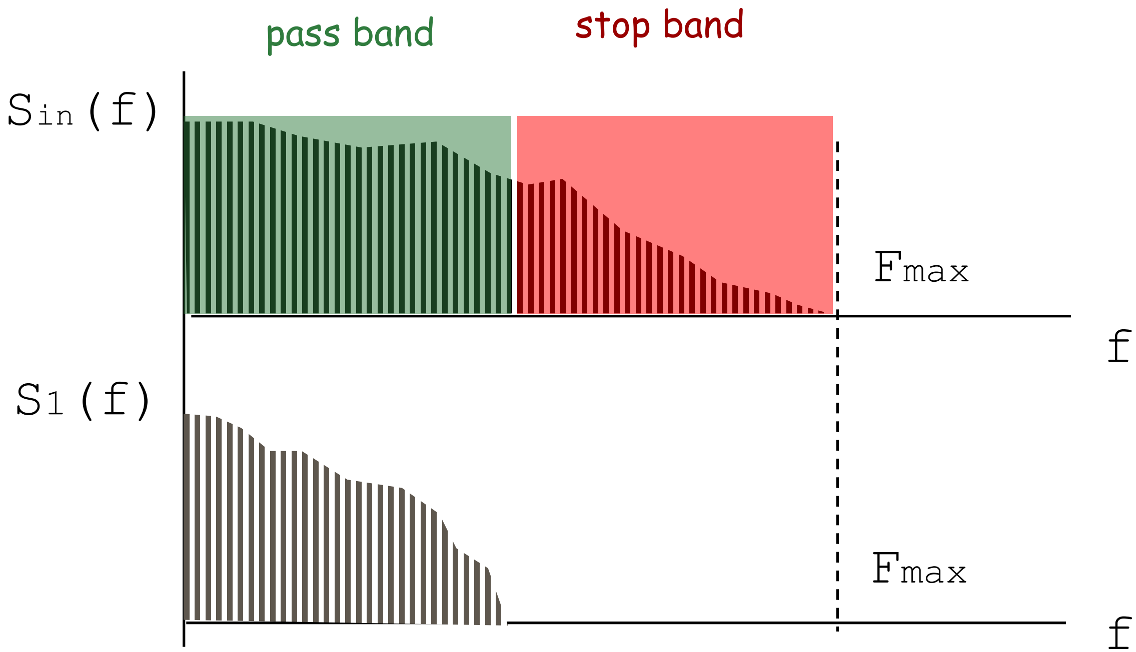

Band Limited Low-Pass Filter

- Band Limited filters are usually implemented as a linear operation that uses a weighted sum of input samples to produce an output sample

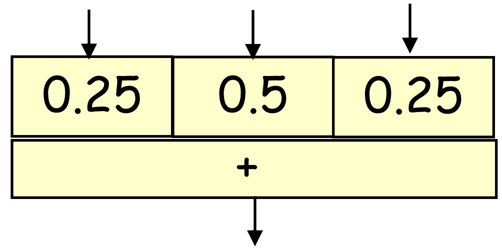

Example: A Low Pass Filter with three weights

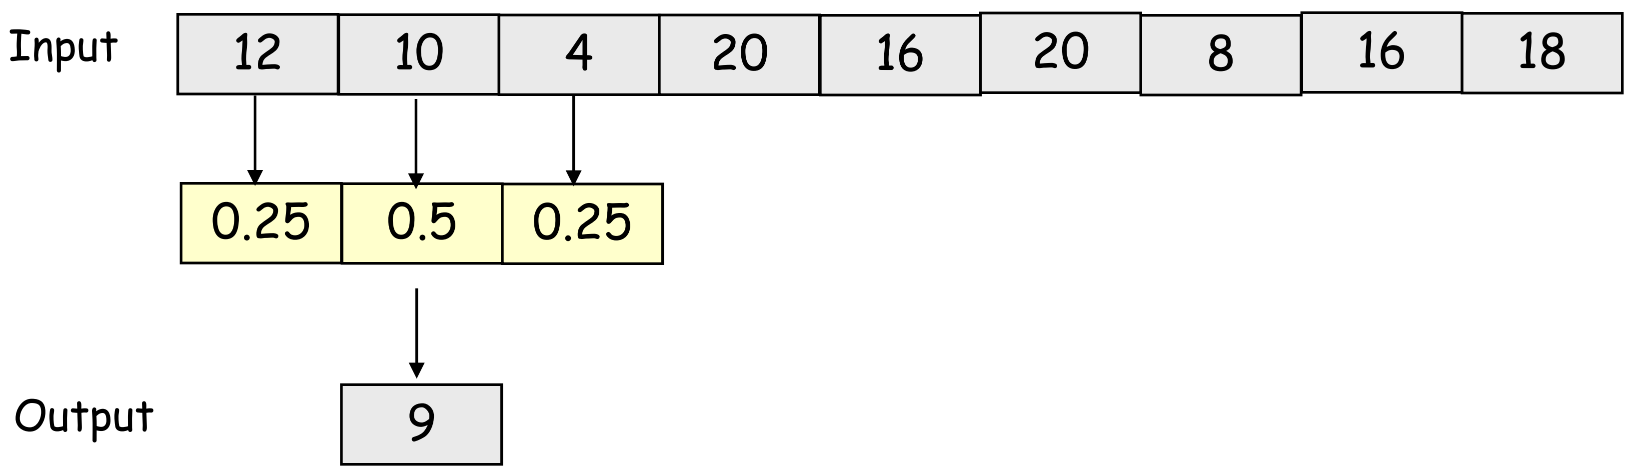

12*0.25 + 10*0.5 + 4*0.25 =9

Band Limited Low-Pass Filter

- Band Limited filters are usually implemented as a linear operation that uses a weighted sum of input samples to produce an output sample

Example: A Low Pass Filter with three weights

0.25	0.5	0.25

Input

12	10	4	20	16	20	8	16	18

0.25	0.5	0.25

Output

9	9.5

10*0.25 + 4*0.5 + 20*0.25 =9.5

Band Limited Low-Pass Filter

- Band Limited filters are usually implemented as a linear operation that uses a weighted sum of input samples to produce an output sample

Example: A Low Pass Filter with three weights

0.25	0.5	0.25

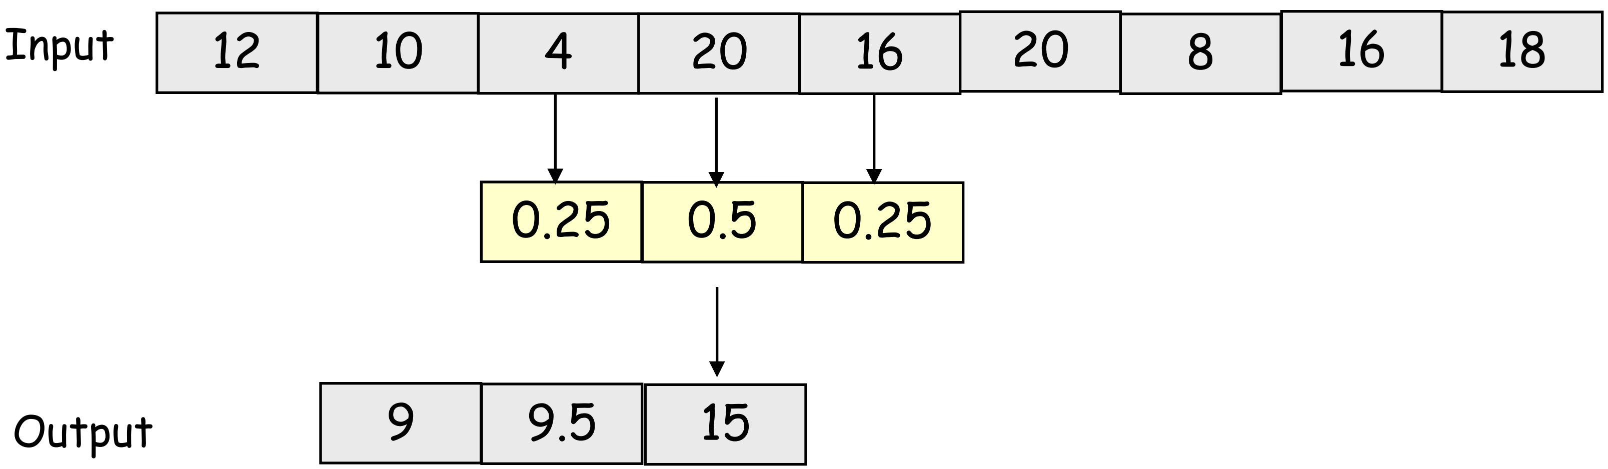

4*0.25 + 20*0.5 + 16*0.25 =15

Low-Pass Filter + Down-Sampling

- As some samples are to be discarded at the Stage 2, filtering and down-sampling processes can be combined to skip unneeded positions

Example: A Low Pass Filter and Down-sampling /2

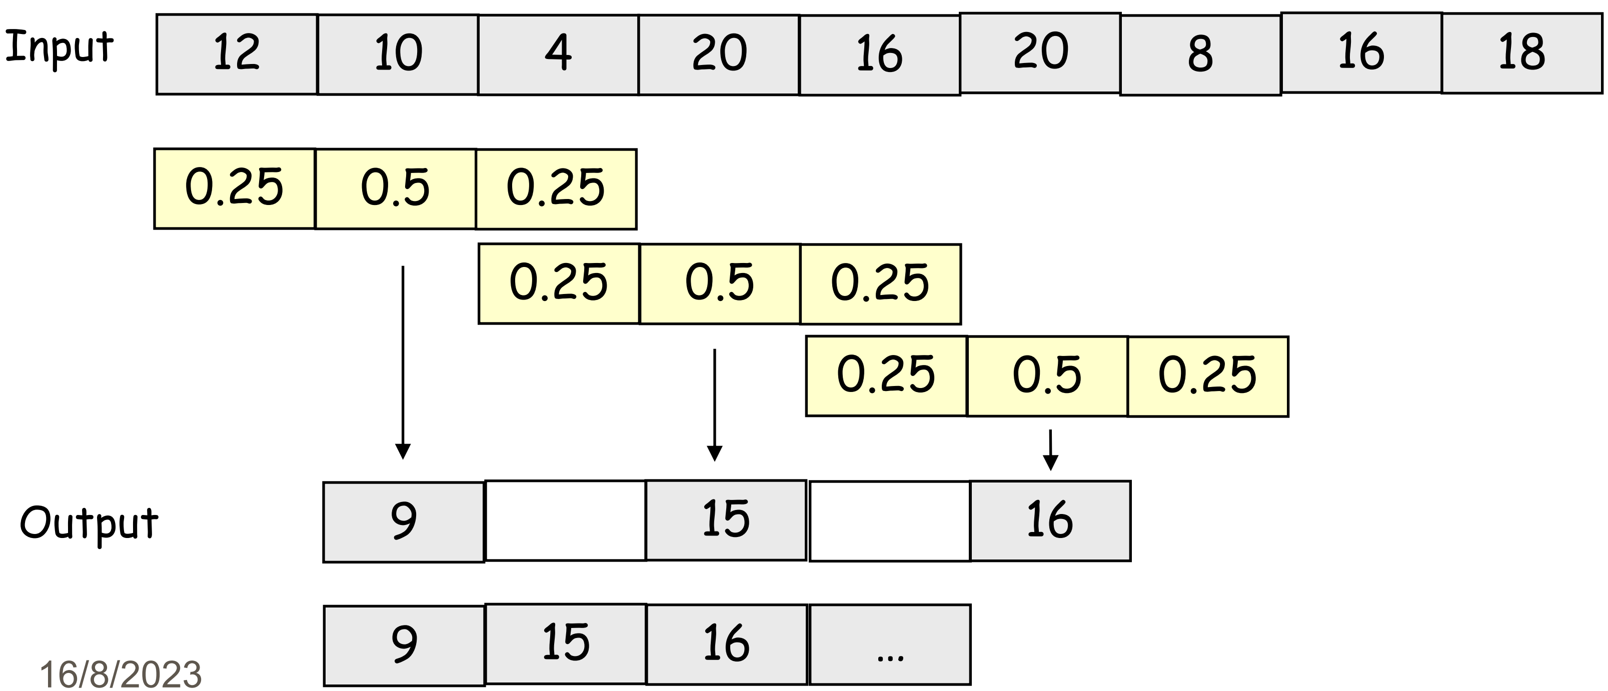

Convolution

- The process of filtering is mathematically described as a convolution of two sequences
  1. x(n) – an input sequence
  2. h(i) – a sequence of filter weights
- The output sequence y(n)

$$y(n) = \{x \otimes h\}(n) = \sum_k x(k)*h(n-k)$$

Example: h(-1) = 0.25 h(0) = 0.5 h(1)=0.25, other h(i) = 0

x(n) = { 12, 10, 4, 20, 16, 20, … }

y(1) = x(0)*h(1) + x(1)*h(0) + x(2)*h(-1) = 3+5+1 = 9

y(2) = ~~x(0)*h(2)~~ + x(1)*h(1) + . . .

y(3) =

2D Low-Pass Filter

- As images are 2D arrays of samples, they are filtered by 2D filters
- A 2D filter is a 2D array of weights (filter coefficients)

Example: `h(n, m) =`

0.11	0.11	0.11
0.11	0.11	0.11
0.11	0.11	0.11

To get one filtered pixel:

9 multiplication

1 addition

- A 2D filter moves across an image column-by-column, row-by-row to produce a processed output image

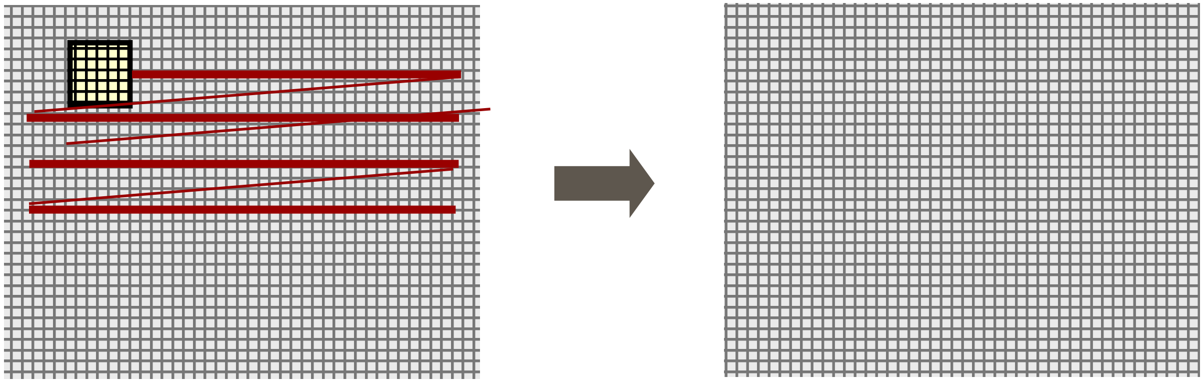

2D Low-Pass Filter

Example:

0.11	0.11	0.11	0.11	0.11	0.11
0.11	0.11	0.11	0.11	0.11	0.11
0.11	0.11	0.11	0.11	0.11	0.11
0.11	0.11	0.11	0.11	48	46
52	47	45	52	47	45
54	57	50	54	57	50

1. Multiply filter coefficients with the corresponding pixel samples and add the results together to produce a filtered sample
2. Place the filtered sample in the output image at the position corresponding to the current position of the filter centre
3. Move the filter to the next position

Example

Example

Down-sampled image

Colour Aliasing

- RGB components of image sensors with Bayer pattern are sampled at different rates. Which rate must be matched by lens sharpness?
- Camera manufacturers usually match the G rate even though this results in aliasing for R and B
- Aliasing cannot be removed, but its visibility can be reduced by CFA Interpolation

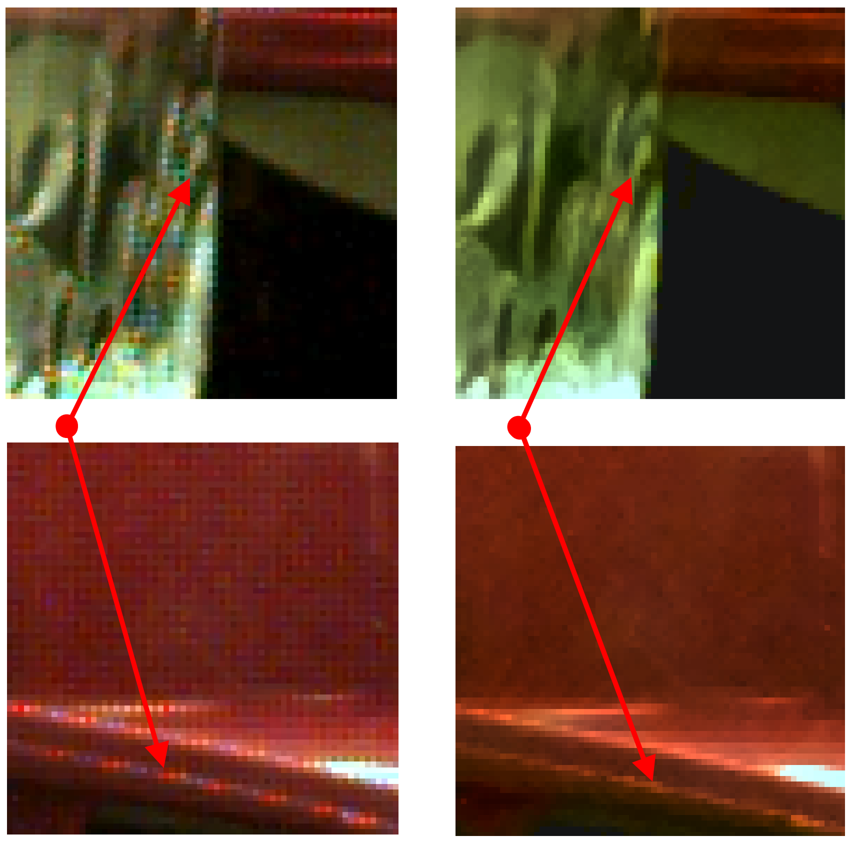

Focusing

- When a lens is focused on an object, theoretically, all objects at other distances S from the camera are out of focus

  1/S + 1/s = 1/f

- Practically, objects slightly in front and behind the object in focus are also appear reasonably sharp. This extra depth is called Depth of Field

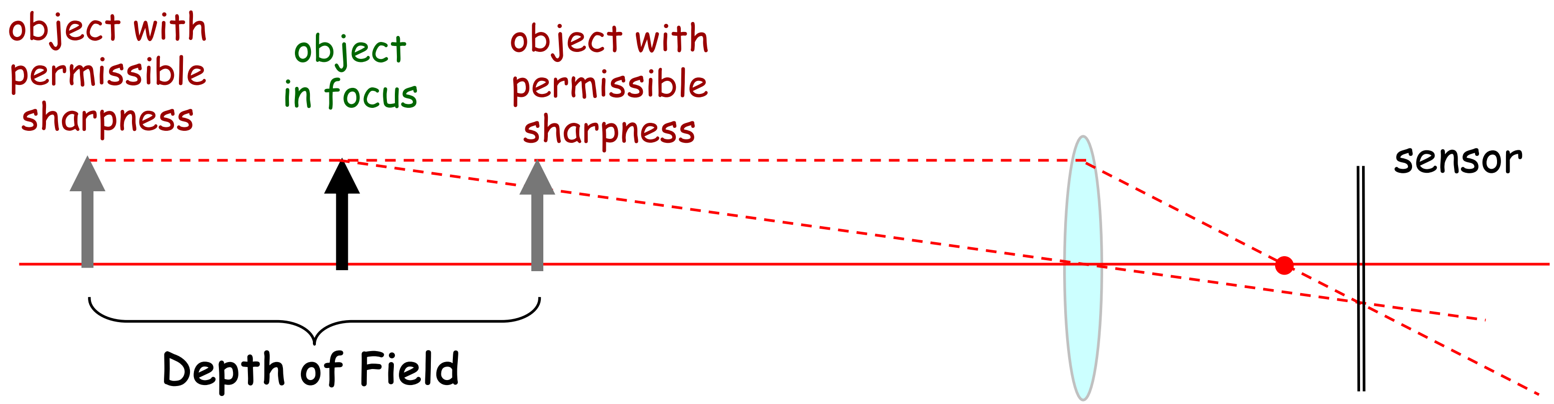

- All other objects will be blurred

Focusing

- Small depth of field is usually a desirable feature of art photos to emphasise an object of interest and isolate it from the background
- Computer Vision Cameras may require a wide depth of field to facilitate analysis of complex scenes
- The greater F, the greater depth of field

  $F = f/D$, where D can be adjusted by changing Iris

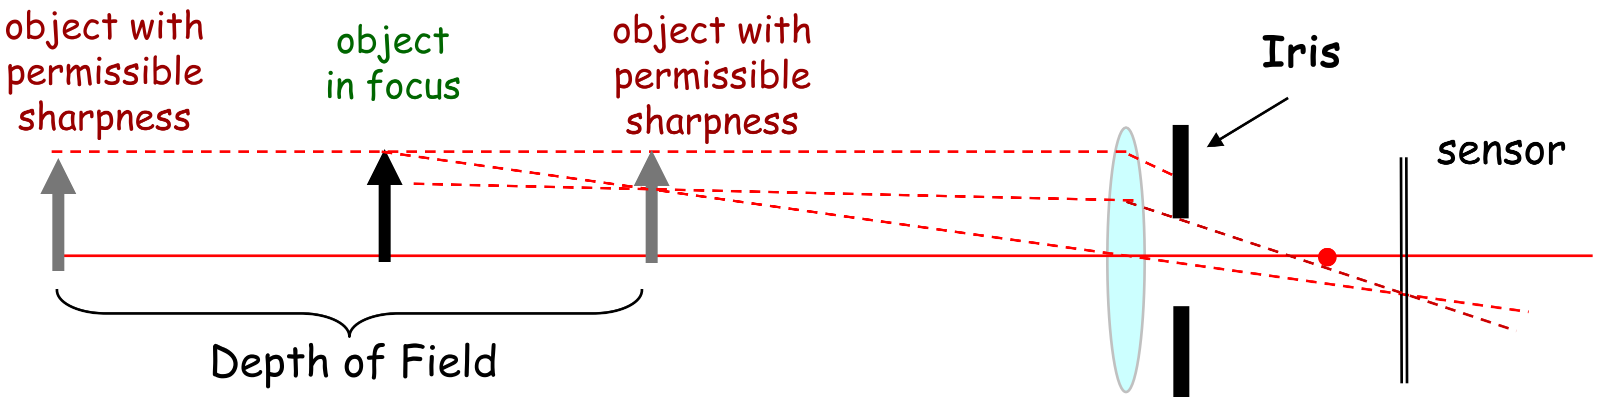

What is a side effect of increasing F ?

Out-of-Focus Blur

- Leads to losses in resolution for all objects which are outside the depth of field

- Out-of-Focus Blur affects images in a way similar to low-pass filtering
- Digital image processing can improve image sharpness by applying inverse filters, but as a side effect it introduces ringing artifacts

Motion Blur

- Leads to losses in resolution in the directions of fast motion

- Digital image processing can improve image sharpness, but as a side effect it introduces ringing artifacts

Saturation

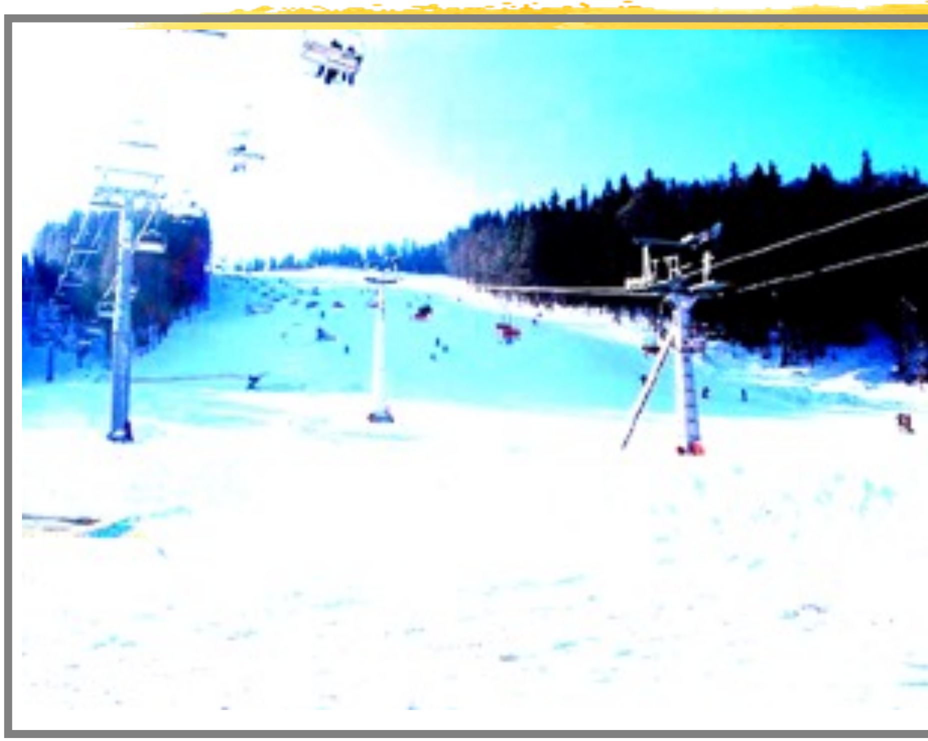

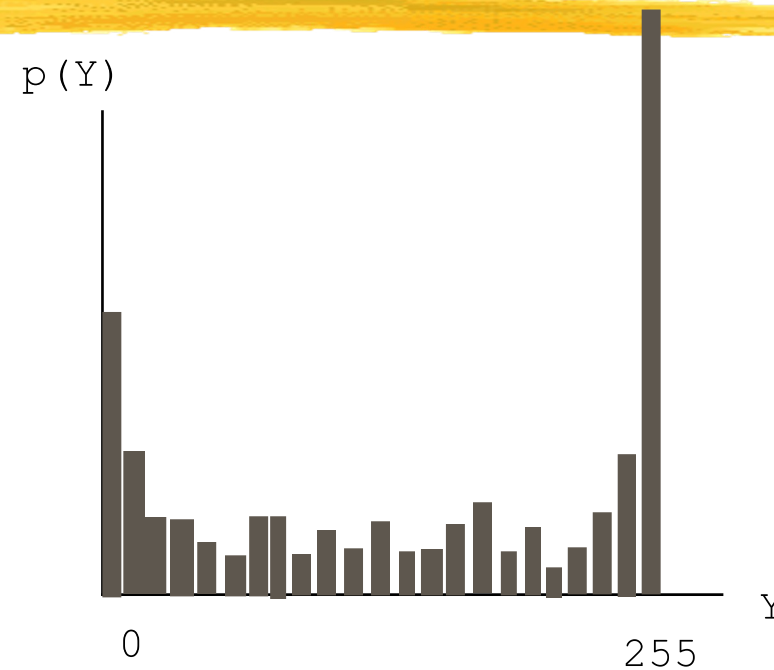

- Capturing of high dynamic range scenes may cause saturation
- Reaching maximum quantisation level 255 samples are clipped producing uniform areas with all details lost

Under-Exposure

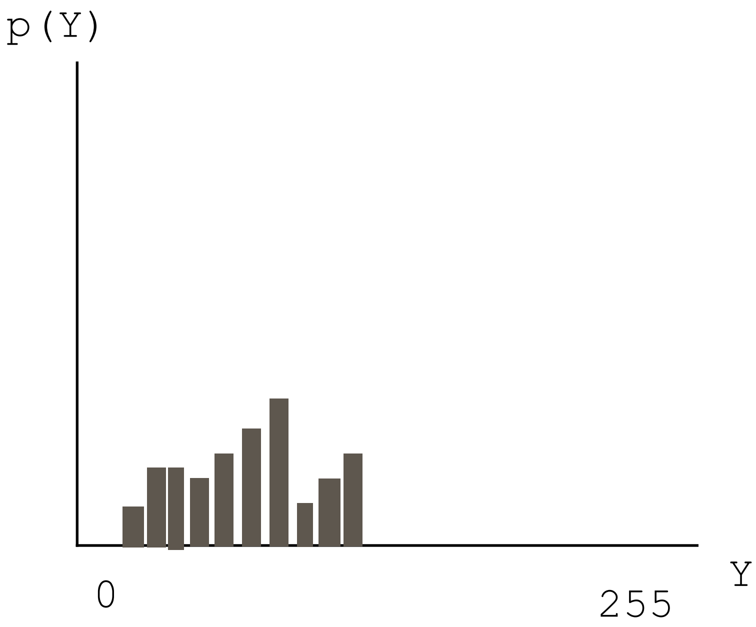

• Digital cameras with inefficient Auto-Exposure and Automatic Gain Control may capture images which do not utilise all quantisation levels leading to loss of details

Under-Exposure

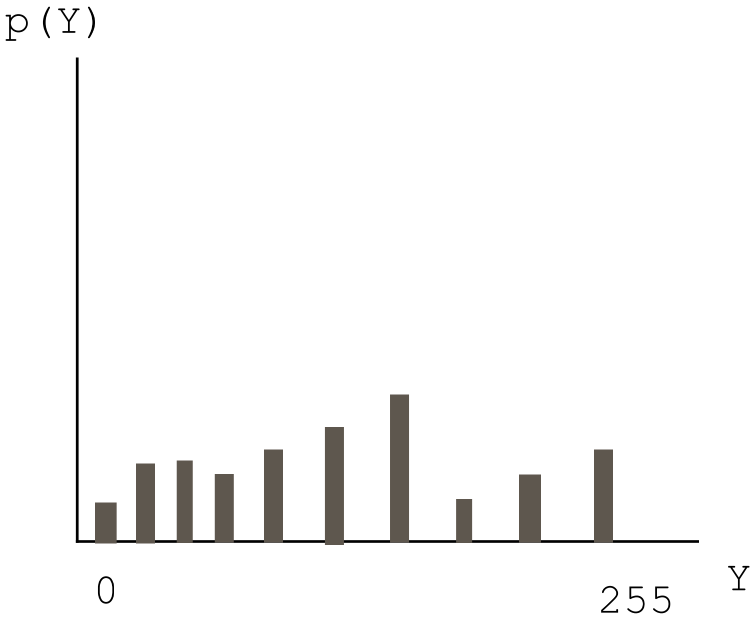

• Simple contrast enhancement changes only visual appearance of the image while the same number of quantisation levels is used

Histogram

- The histogram of an image with gray-levels in the range [0,L-1], where L is usually 256, is a discrete function

$$h(r_k) = n_k$$

where r_k is the k'th gray level

n_k is the number of pixels having gray level r_k

- Normalized histogram

$$p(r_k) = n_k / n$$

$$n = \sum_{k=0}^{L-1} n_k$$

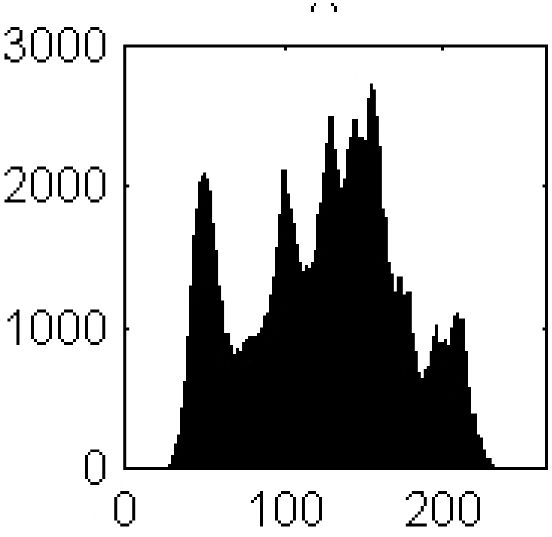

Histogram Equalization/Specification

- Any enhancement eventually change the distribution of color or intensity of the images.

- We could modify the image histogram directly to achieve the enhancement
  - Equalization
    - Try to make the histogram of the enhanced image flat, i.e. make use of all available color/intensity levels
  - Specification/matching
    - To modify histogram of the image to a desirable histogram that are pleasant to most viewers

Histogram Equalization

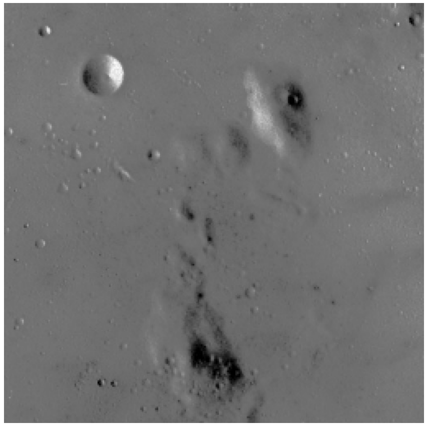

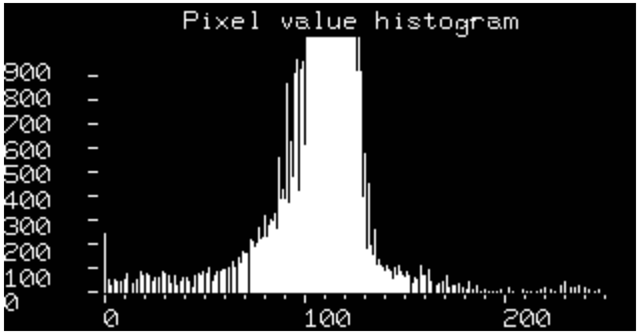

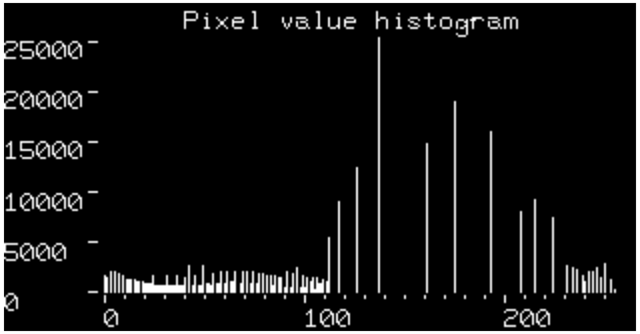

Histogram Equalization

- Consider the gray levels in an image as the realizations of *a random variable*, and let
  - $r \in [0,1]$ - the gray level before equalization
  - $s \in [0,1]$ - the gray level after equalization
- Given $p_r(r)$, histogram equalization is to transform r to s

$$s = T(r)$$

  - such that $p_s(s)$ is constant or close to constant for all s

Histogram Equalization…

If $T(r)$ satisfies the following conditions:

a) single-valued and monotonically increasing in the interval $r \in [0,1]$, and
b) $0 \le T(r) \le 1$ for $0 \le r \le 1$

$$p_s(s) = p_r(r)\left|\frac{dr}{ds}\right| \qquad (1)$$

$$s = T(r) = \int_0^r p_r(w)\,dw \qquad (2)$$

Histogram Equalization…

- For discrete version of the transform

$$s_k = (L-1)T(r_k)$$

$$= (L-1)\sum_{j=0}^{k} p_r(r_j)$$

$$= (L-1)\sum_{j=0}^{k} \frac{n_j}{n}$$

$$k = 0,1,2,\cdots,L-1$$

$$s_k \in [0, L-1]$$

$$r_k \in [0, L-1]$$

Histogram Equalization: Examples

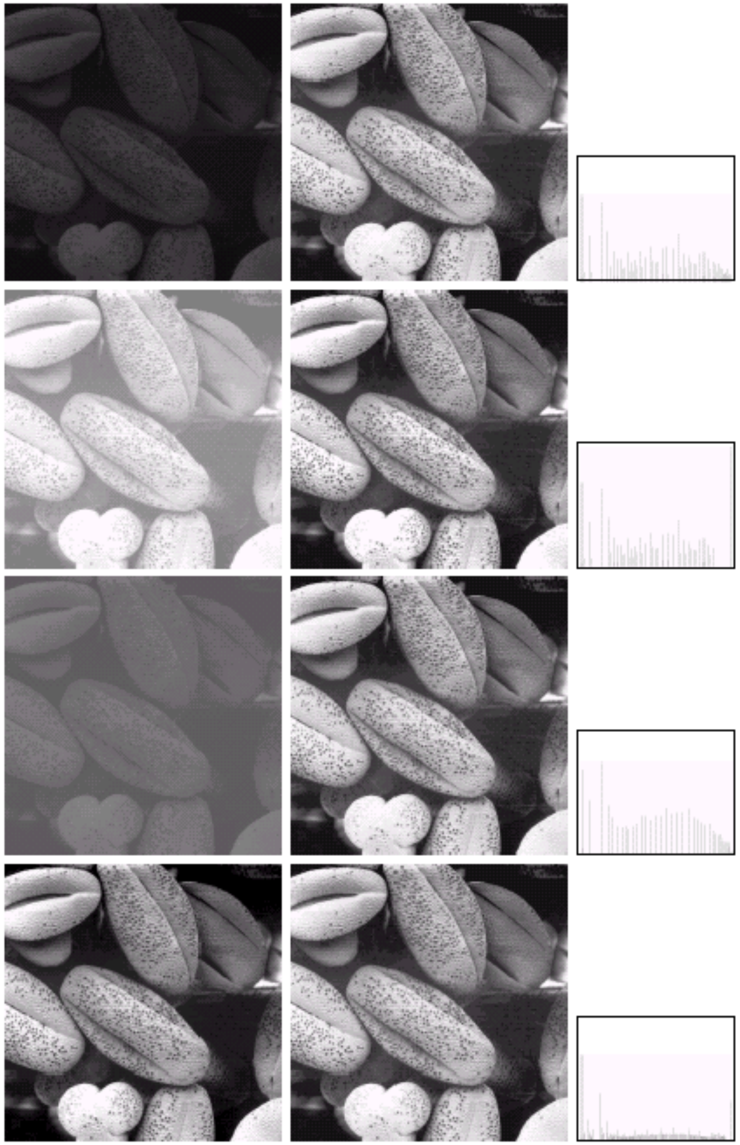

Impact of White Balance

Lighting

The effect of a light source on colour appearance is expressed in the **colour rendering index (CRI)**, on a scale of 0-100. Natural outdoor light has a CRI of 100 and is used as the standard of comparison for any other light source.

Noise

- Digital camera is an electronic device that besides digital, contains analog components
- Receiving and processing target signals they are affected by random fluctuation of electric current, random variation of electronic circuit parameters and interferences from other electronic components
- Noise is a random process and thus it cannot be eliminated in a similar way to other factors affecting image quality such as: *resolution, defocus, colour fidelity, under-exposure*, etc
- It is not possible to reduce noise without employing statistics and the theory of probability

Sources of Noise

Photon Shot Noise – pixel size is so small that the number of absorbed photons fluctuates

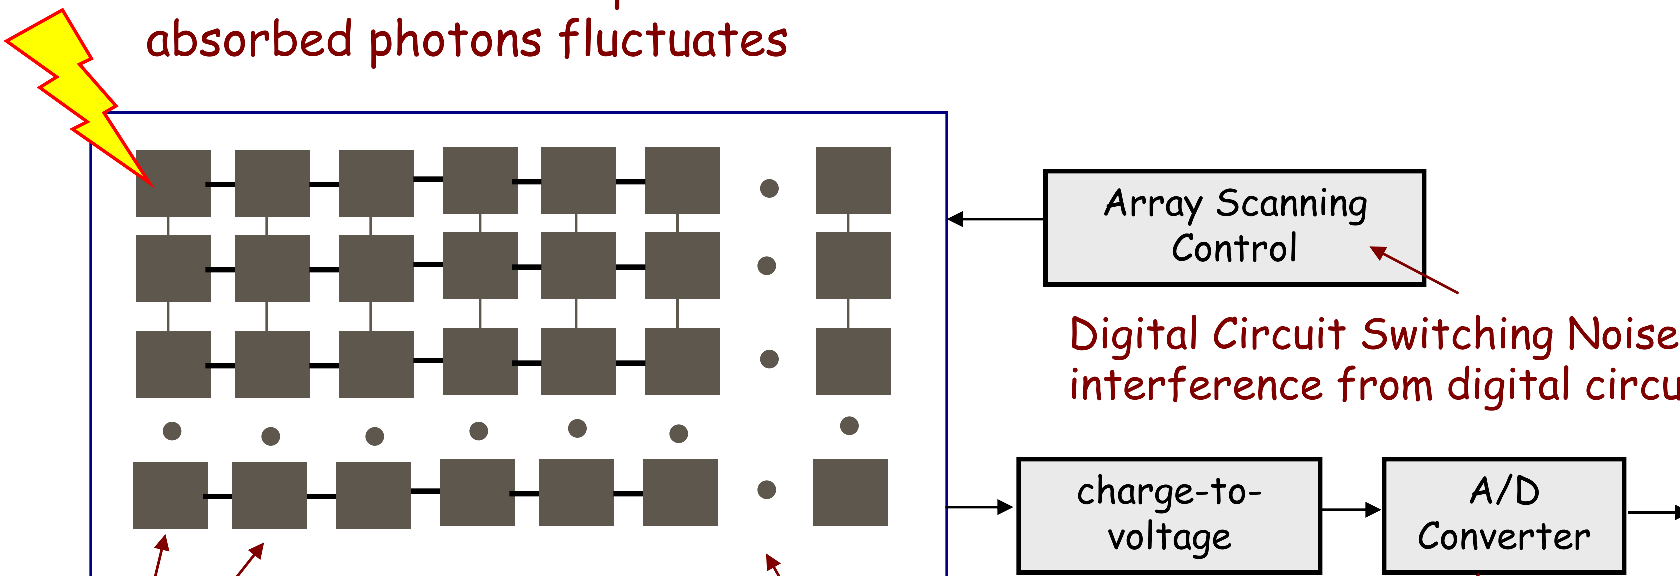

Digital Circuit Switching Noise – interference from digital circuits

Fixed Pattern Noise – different pixels have slightly different sensitivity

Thermal Noise – temperature dependent fluctuation of electrical current

Quantisation Noise – rounding errors due to quantisation

Parameters of Noise

- Noise is a random process and thus, it can be described only statistically

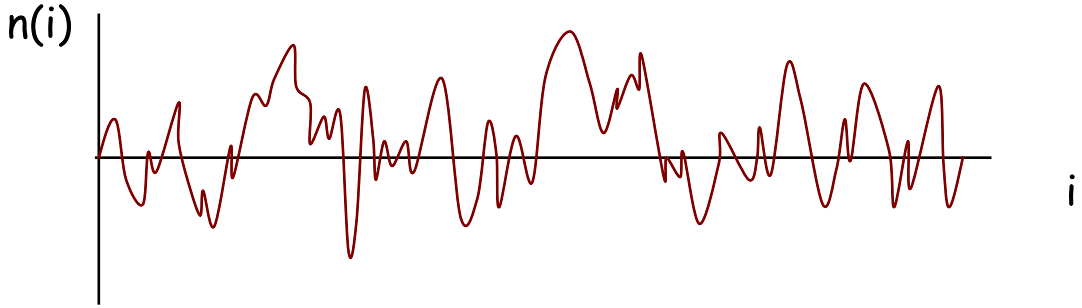

Probability Density

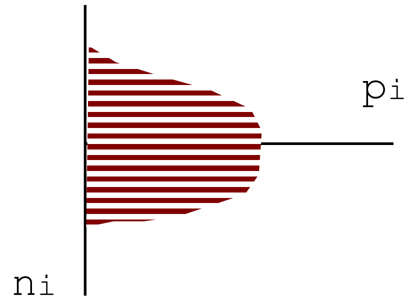

- Other statistical parameters also characterise noise

The mean: $\overline{n} = \frac{1}{n} * \Sigma n_i$ or $\overline{n} = \Sigma p_i * n_i$

The standard deviation: $\sigma = \sqrt{\frac{1}{n} * \Sigma (n_i - \overline{n})^2}$ or $\sigma = \Sigma p_i * (n_i - \overline{n})^2$

Noise Probability Density

- Different types of noise have different probability densities
- There are two the most common Probability Densities

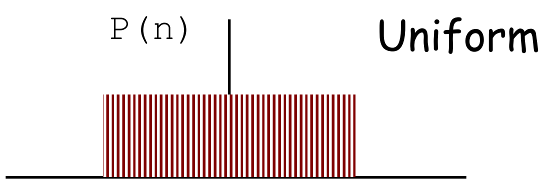

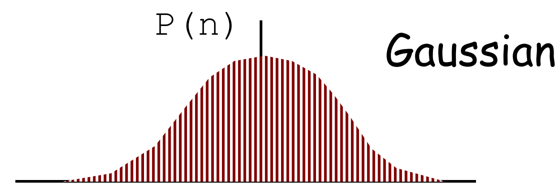

Q: A uniform area of a digital image with luminance Y = 200 is affected by noise. What is the probability that a luminance sample affected by noise is visually different from the uniform background? Consider the zero-mean uniform and the Gaussian distributions.

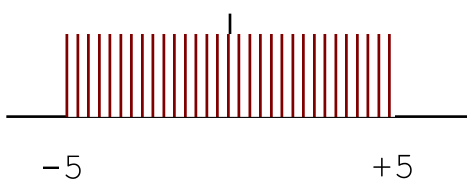

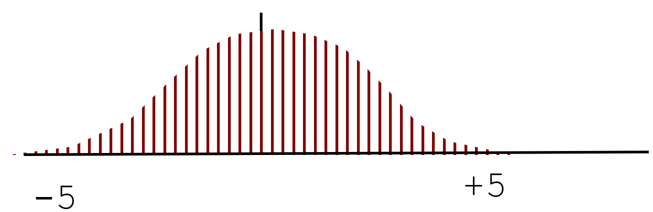

According to Webber's law the visibility threshold is 2%. 2% of 200 is ±4

Uniform: P(|n|>4) = 20%

Quantisation Noise

Quantised samples

$s_q(i) = \mathrm{int}(s(i)/q)$

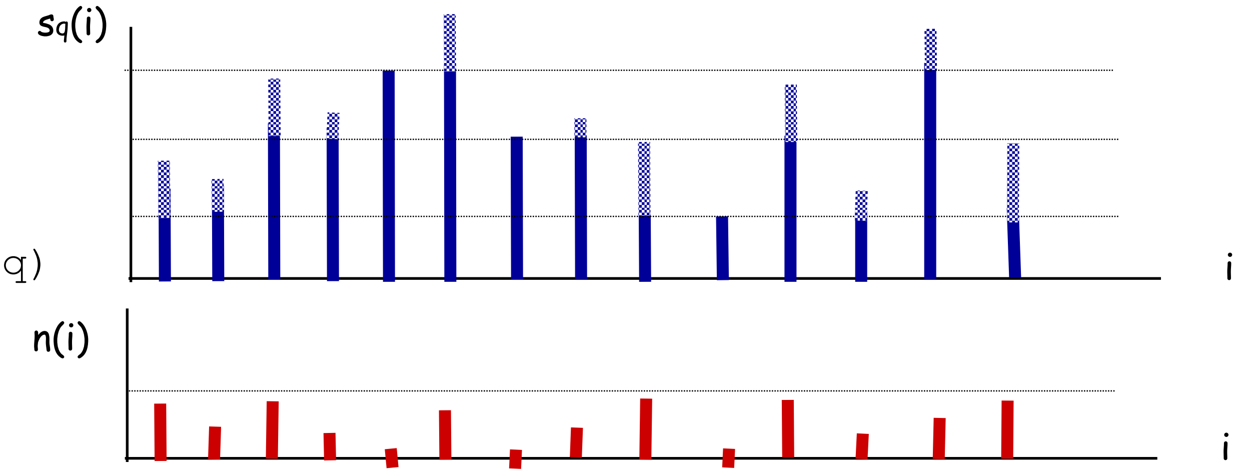

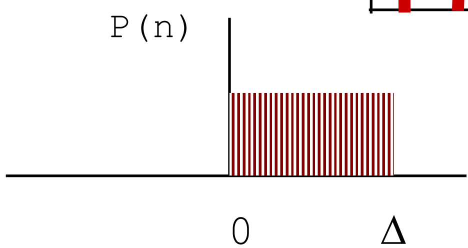

- Quantisation noise has a uniform probability density in the range $[0\ldots\ \Delta]$, or $[-\ \Delta/2\ldots\ \Delta/2]$ if rounding used instead of truncation

Signal to Quantization Noise Ratio (SQNR)

- For a quantization accuracy of *n* bits per sample, in the worst case, the SQNR is

$$\Delta = \frac{2X_{\max}}{M} \quad M = 2^n \quad \sigma_d^2 = \frac{\Delta^2}{12} \quad \sigma_x^2 = (2X_{\max})^2/12$$

$$SQNR = 10\log_{10}(\frac{\sigma_x^2}{\sigma_d^2}) = 10\log_{10} M^2 = 6.02n(dB)$$

What should be the bits per sample if SQNR=60db ?

60db => 10bit/sample

Central Limit Theorem

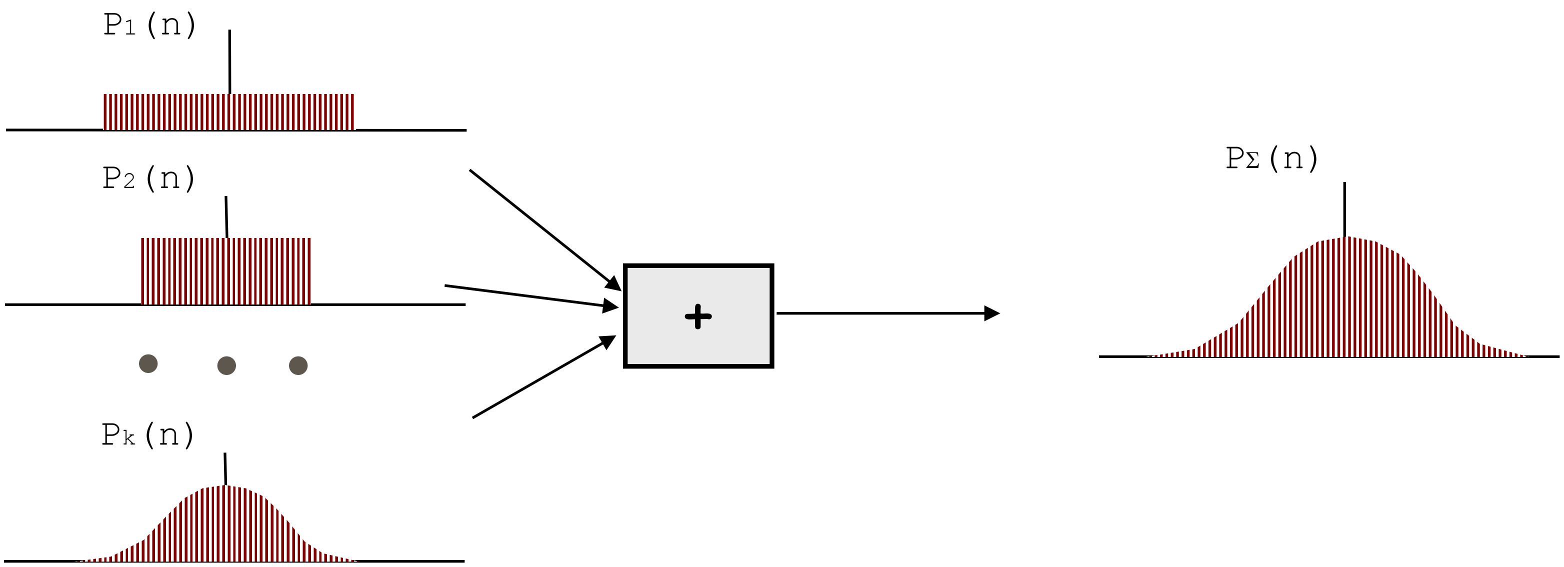

- The sum of a large number of independent random processes will have approximately Gaussian probability density
- As image noise comes from several independent sources its distribution is approximately Gaussian with $\overline{n} = 0$

Noise Model

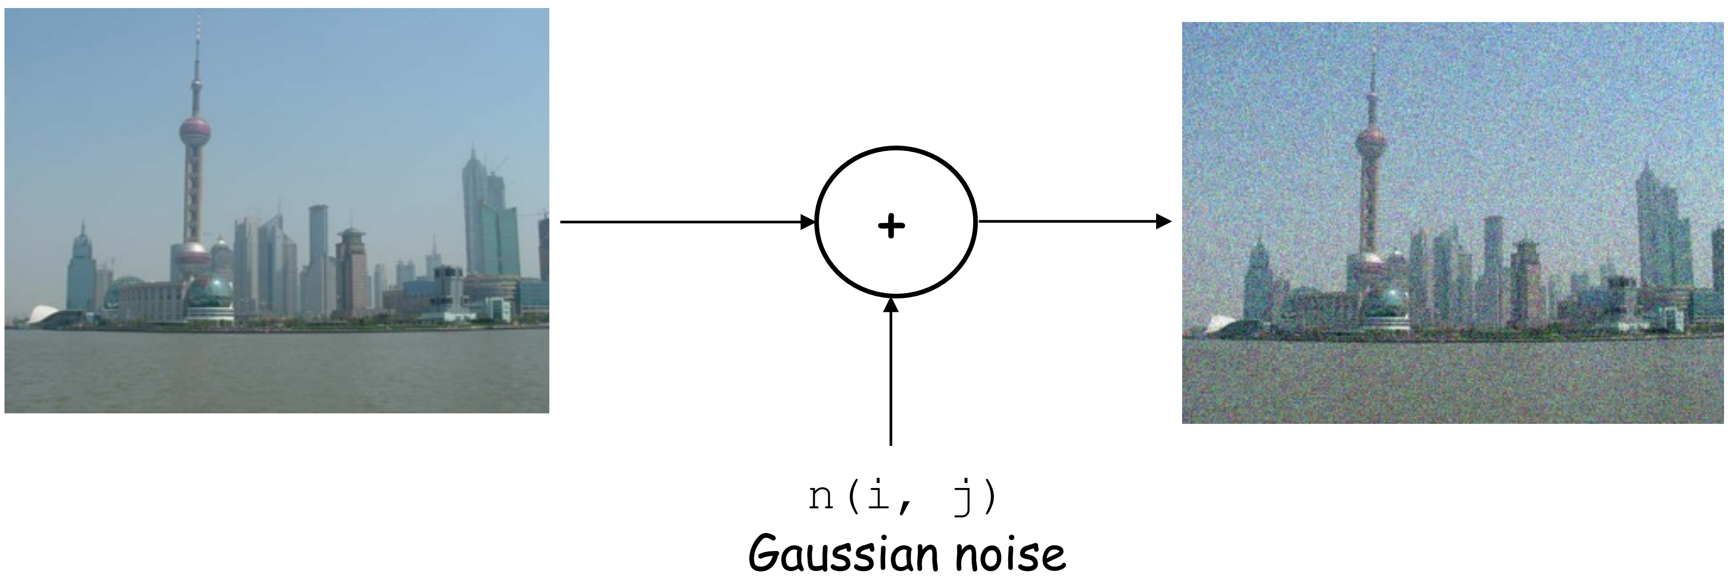

- Image quality analysis, image processing and image enhancement algorithms are based on the model that noise is additive to image samples with 0-mean Gaussian distribution

Quiz

- Images of Macbeth Colour Checker are used for measuring colour fidelity of the colour correction. The images are affected by noise with the standard deviation above the visibility threshold. Each colour patch is 50x50 pix. Gamma correction is disabled.

  Can the noise affect the measurements?

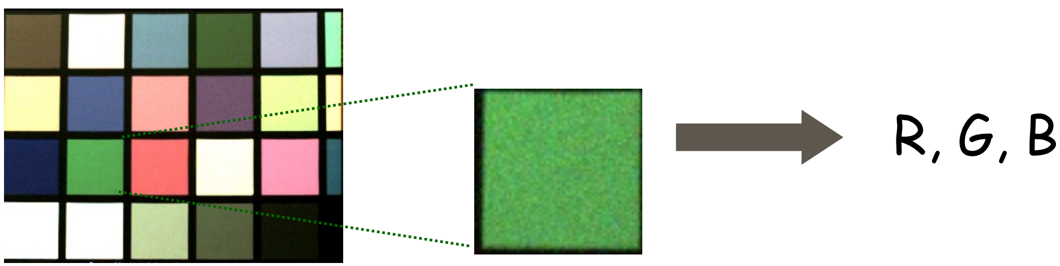

$$G + \overline{n} = G$$

- For green plane:

$$G_n = \frac{1}{2500}\sum_{50}\sum_{50} S_n(i,j) = \frac{1}{2500}\sum_{50}\sum_{50}(S(i,j)+n(i,j)) = G + \frac{1}{2500}\sum_{50}\sum_{50} n(i,j)$$

Quiz

- Images of Macbeth Colour Checker are used for measuring colour fidelity of the image processing chain that included Gamma Correction. The images are affected by noise with the standard deviation above the visibility threshold. Each colour patch is 50x50 pix.

Can the noise affect the measurements?

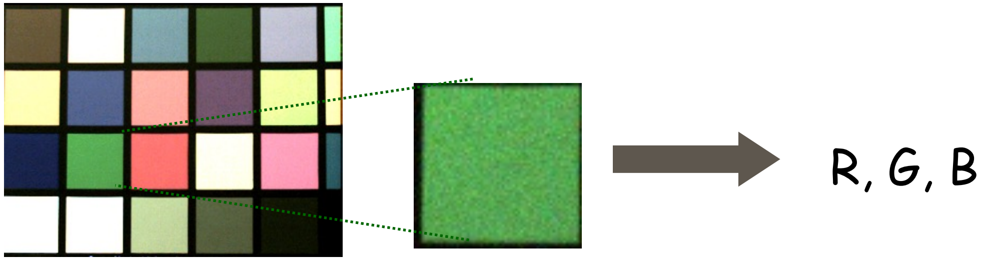

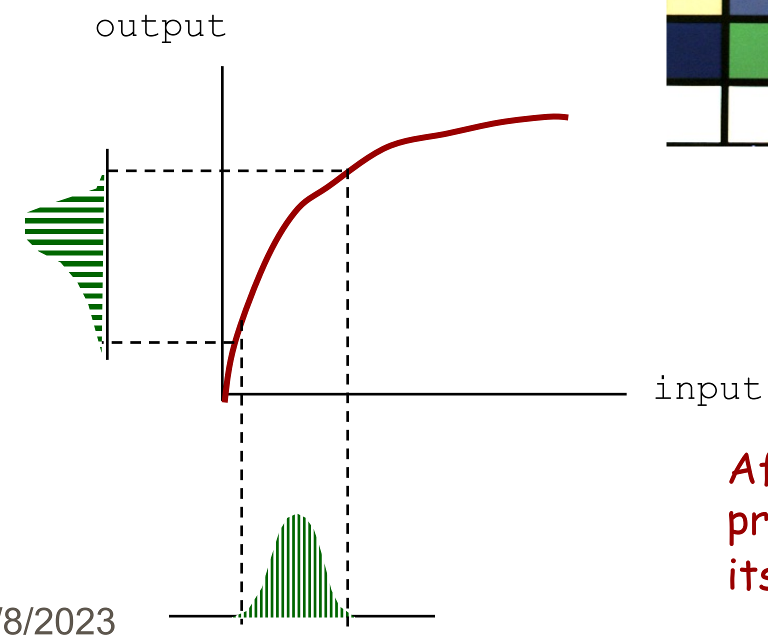

After the Gamma Correction stage, noise probability density is not symmetrical and its mean is not equal to 0. Thus $G_n \neq G$

Gamma Correction and Noise

- Non-linearity of Gamma Correction leads to noise amplification in image regions with low illuminance

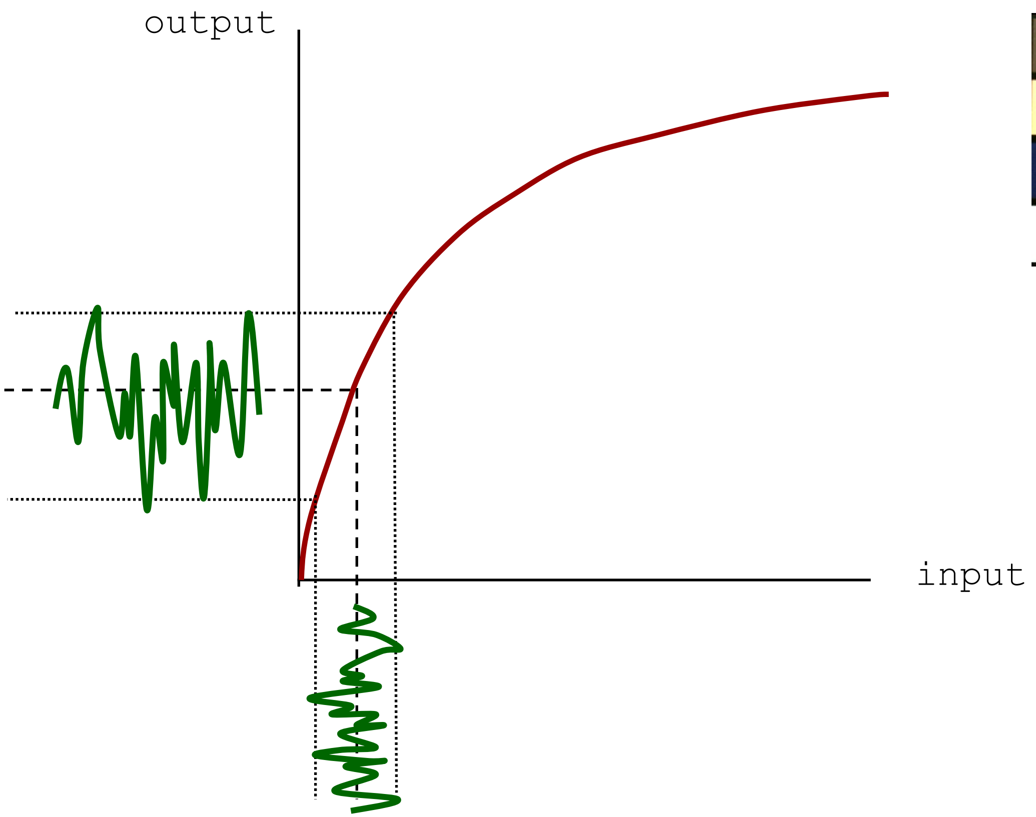

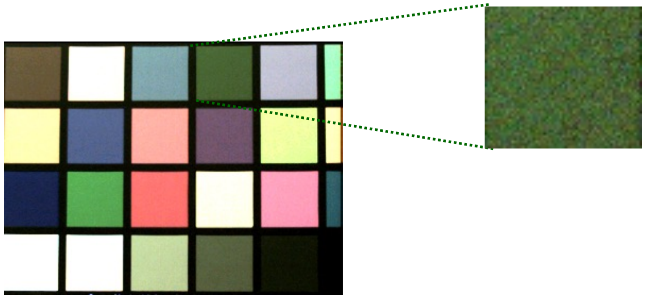

Colour Correction and Noise

- Improving accuracy of colour reproduction Colour Correction amplifies noise

$$\begin{bmatrix} R' \\ G' \\ B' \end{bmatrix} = \begin{bmatrix} 1.28 & -0.11 & -0.12 \\ -0.12 & 1.41 & -0.13 \\ -0.14 & -0.15 & 1.36 \end{bmatrix} * \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

- As Colour Correction $\mathbf{T}$ is a linear operation and noise $\mathbf{n}$ is additive, we can analyse noise propagation independently

$\mathbf{I' = T*(I + n) = T*I + T*n}$, where $\mathbf{n}$ – noise vector

$$\begin{bmatrix} n_r' \\ n_g' \\ n_b' \end{bmatrix} = \begin{bmatrix} 1.28 & -0.11 & -0.12 \\ -0.12 & 1.41 & -0.13 \\ -0.14 & -0.15 & 1.36 \end{bmatrix} * \begin{bmatrix} n_r \\ n_g \\ n_b \end{bmatrix}$$

Colour Correction and Noise

- To simplify analysis, we make two assumptions:
  1. Noise in three colour channels R, G, B is statistically independent
  (explain why this is not strictly correct for cameras with Bayer pattern sensors)
  2. Noise standard deviation σ in R, G, B channels is the same
  (explain why this is not strictly correct for cameras with Bayer pattern sensors)

$$\sigma(n_g') = \sqrt{(0.12*\sigma(n))^2 + (1.41*\sigma(n))^2 + (0.13*\sigma(n))^2}$$

$$\sigma(n_g') = 1.42*\sigma(n)$$

noise deviation is 42% higher

Exercise: What is the noise amplification in other colour channels?
Which colour channel noise is more visually noticeable?

Image Spectrum

- Basic steps of image processing in frequency domain

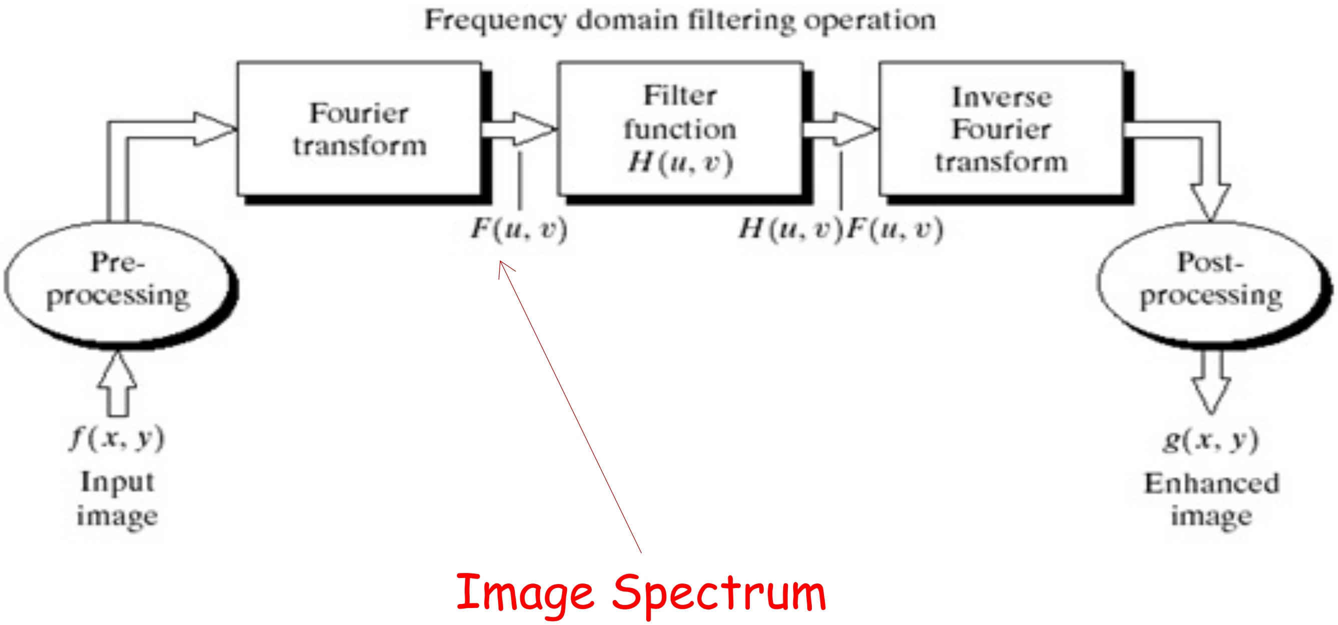

Fourier Transform

- A digital image f(x,y) of size MxN, its Fourier Transform is defined as

$$\Im[f(x,y)] = F(u,v) = \frac{1}{MN}\sum_{x=0}^{M-1}\sum_{y=0}^{N-1} f(x,y)e^{-j2\pi(ux/M+vy/N)}$$

$$j = \sqrt{-1} \quad e^{j\theta} = \cos\theta + j\sin\theta$$

$$u = 0,1,\cdots M-1 \quad v = 0,1,\cdots N-1$$

- Similarly, the inverse transform is

$$\Im^{-1}[F(u,v)] = f(x,y) = \frac{1}{MN}\sum_{x=0}^{M-1}\sum_{y=0}^{N-1} F(u,v)e^{j2\pi(ux/M+vy/N)}$$

Fourier Transform

- Fourier spectrum (or magnitude), phase and power spectrum of f(x,y)

$$|F(u,v)| = \left[R^2(u,v) + I^2(u,v)\right]^{1/2}$$

$$\Phi(u,v) = \tan^{-1}\left[\frac{I(u,v)}{R(u,v)}\right]$$

$$P(u,v) = R^2(u,v) + I^2(u,v)$$

Digital Fourier Transform Illustration

Fourier Transform

- Centred Fourier Spectrum

$$\Im[f(x,y)(-1)^{x+y}] = F(u - M/2, v - N/2)$$

$$u = 0,1,\cdots M-1 \quad v = 0,1,\cdots N-1$$

- F(0,0) moved to the centre of the MxN area

- F(0,0) represents the average gravel level of the image, known as dc component of the spectrum

An example

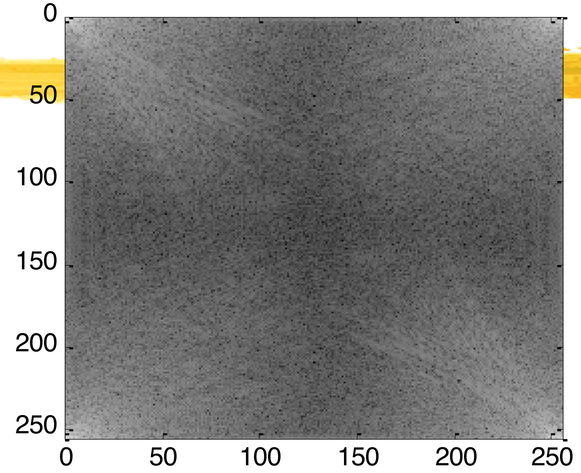

$\log_{10}P$

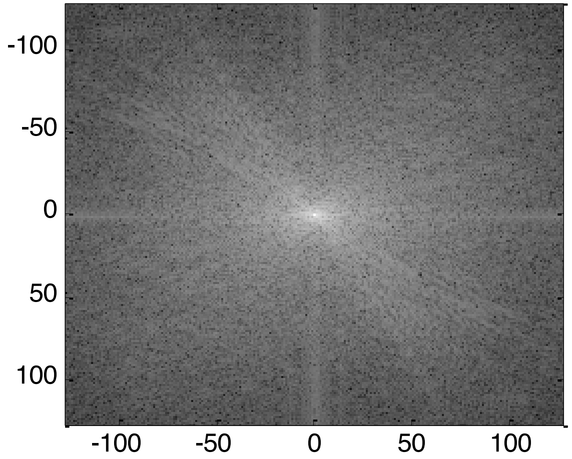

Centred $\log_{10}P$

- Digital Image Spectrum should not be confused with the spectrum of visible light (400nm – 700nm) !

Properties of Fourier Transform

- If f(x,y) is real, its Fourier Transform is conjugate symmetric, i.e.

$$F(u,v) = F^*(-u,-v)$$

- Convolution of f(x,y) with a kernel h(x,y) is equivalent the multiplication of their Fourier Transform, i.e.

$$f(x,y) \otimes h(x,y) = F(u,v)H(u,v)$$

$$f(x,y) \otimes h(x,y) = \frac{1}{MN}\sum_{m=0}^{M-1}\sum_{n=0}^{N-1} f(m,n)h(x-m,y-n)$$

Noise Spectral Properties

- Dominant components of natural images have smooth variation of luminance and colour, thus low spatial frequencies
- A typical spectrum of natural images

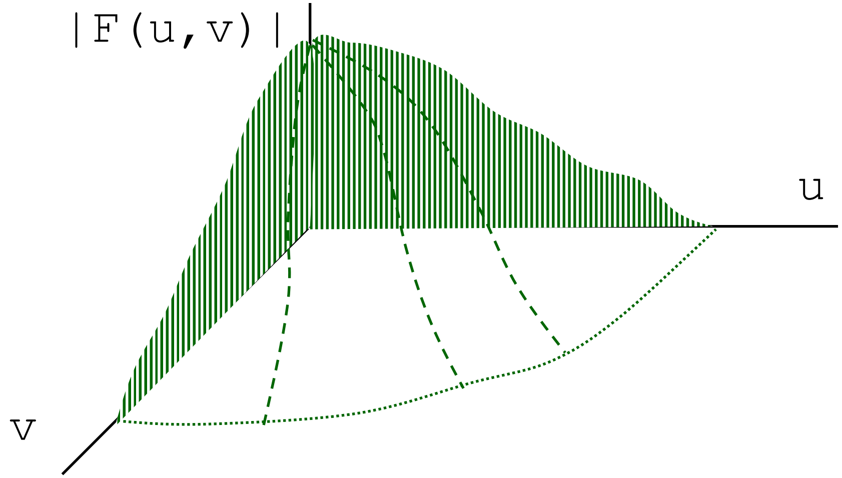

• The dominant frequencies are usually low

• High frequencies are usually not very powerful

Exercise: How would you explain the fact that diagonal spectrum components are usually weaker than horizontal and vertical once?

Noise Spectral Properties

- Spatial spectrum of image white noise (i.e. zero means, finite variance) is 2D (horizontal - vertical)

  e.g. white Gaussian noise
- Spectrum is uniform

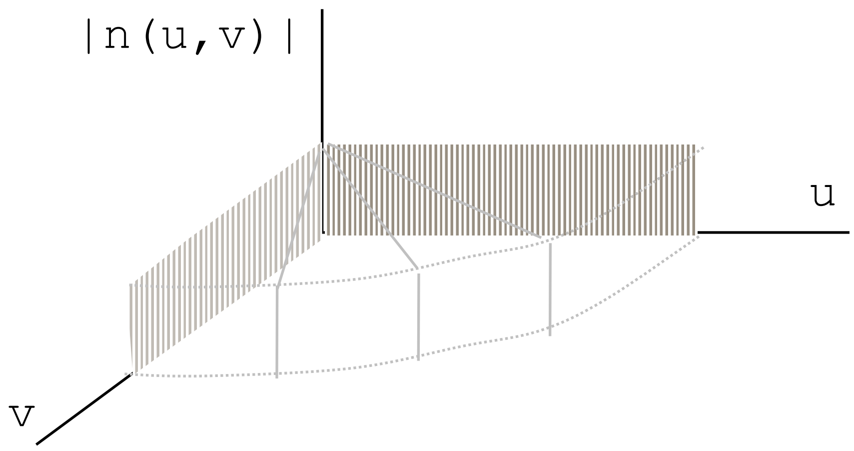

Image and Noise Spectral Properties

- Linearity of spectral representation

  If $F(u, v)$ is the spectrum of image $f(x, y)$

  $N(u, v)$ is the spectrum of noise $n(x, y)$

  Then, the spectrum of $f(x, y) + n(x, y)$ is equal to $F(u, v) + N(u, v)$

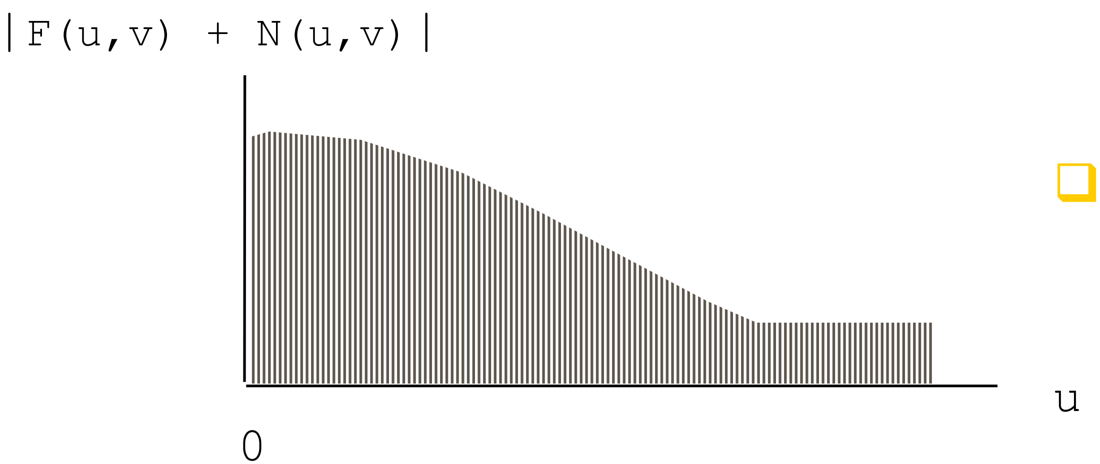

- Once the spectrums are added, their separation is not a trivial task

Filtering in Frequency Domain

- Filtering in the Frequency domain is straightforward:
  - Compute F(u,v) of image f(x,y)
  - Design filter function H(u,v)
  - Calculate

$$G(u,v) = F(u,v)H(u,v)$$

  - Calculate inverse Fourier Transform of G(u,v) to obtained filtered (processed image)

$$\Im^{-1}[G(u,v)]$$

Types of Filters

- Depending what part of the image spectrum is passed and what part of the spectrum is suppressed all filters can be divided into several groups
- 1D examples of possible frequency responses

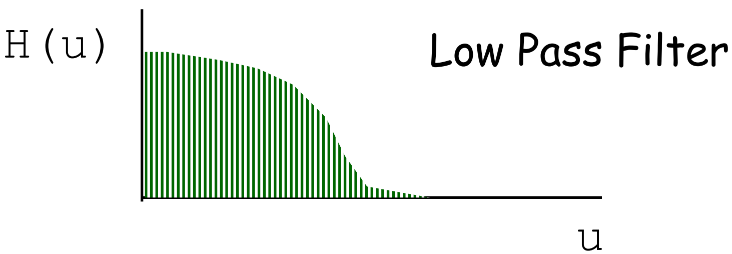

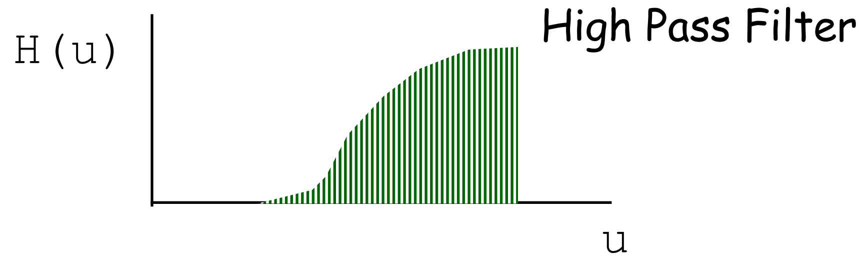

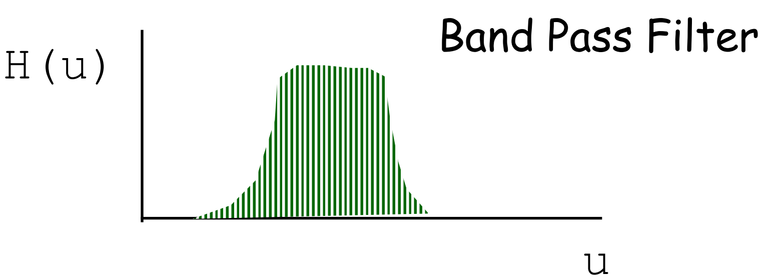

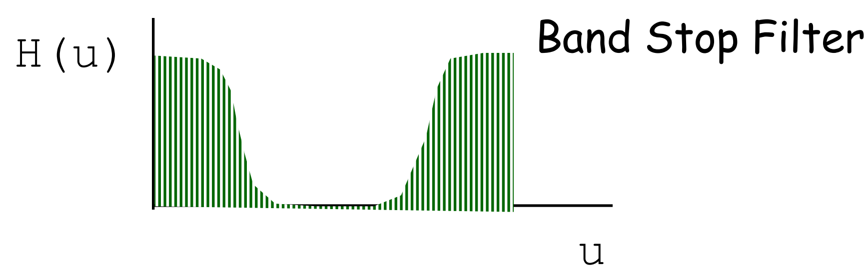

An example

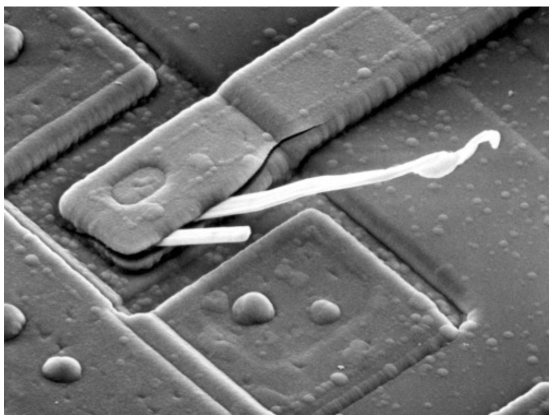

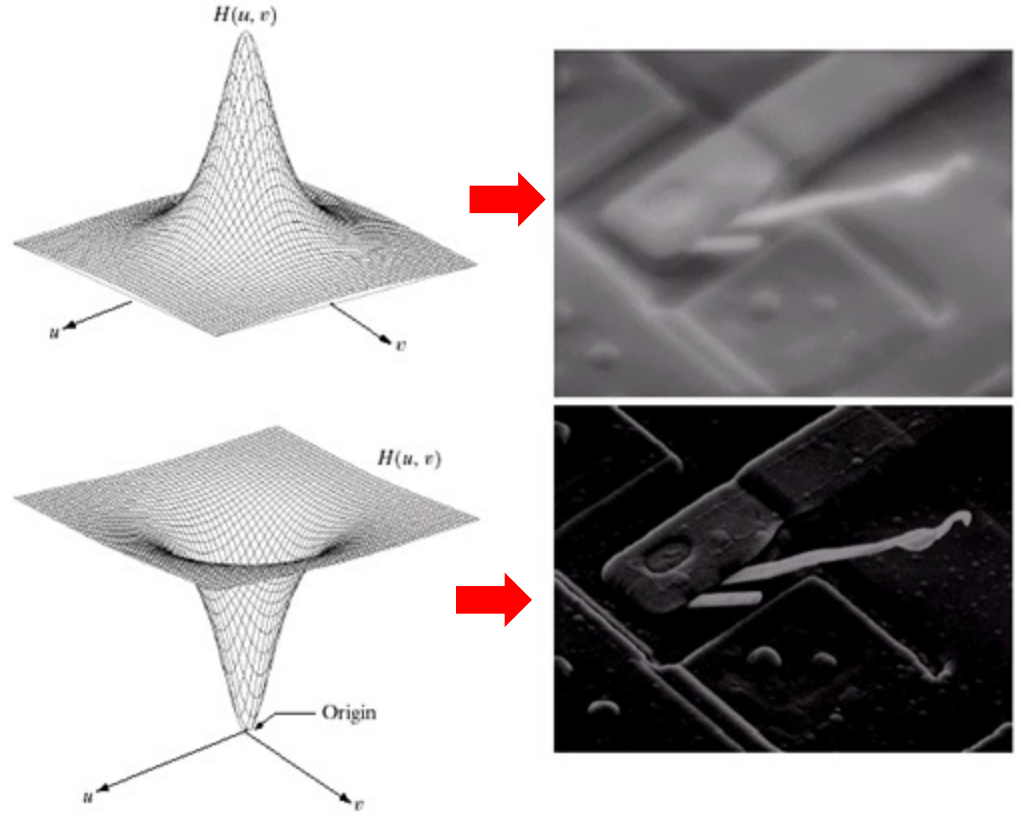

a b
c d

Ideal Lowpass Filters

$$H(u,v)=\begin{cases}1 & D(u,v)\le D_0\\ 0 & D(u,v)>D_0\end{cases}$$

$$D(u,v)=[(u-M/2)^2+(v-N/2)^2]^{1/2}$$

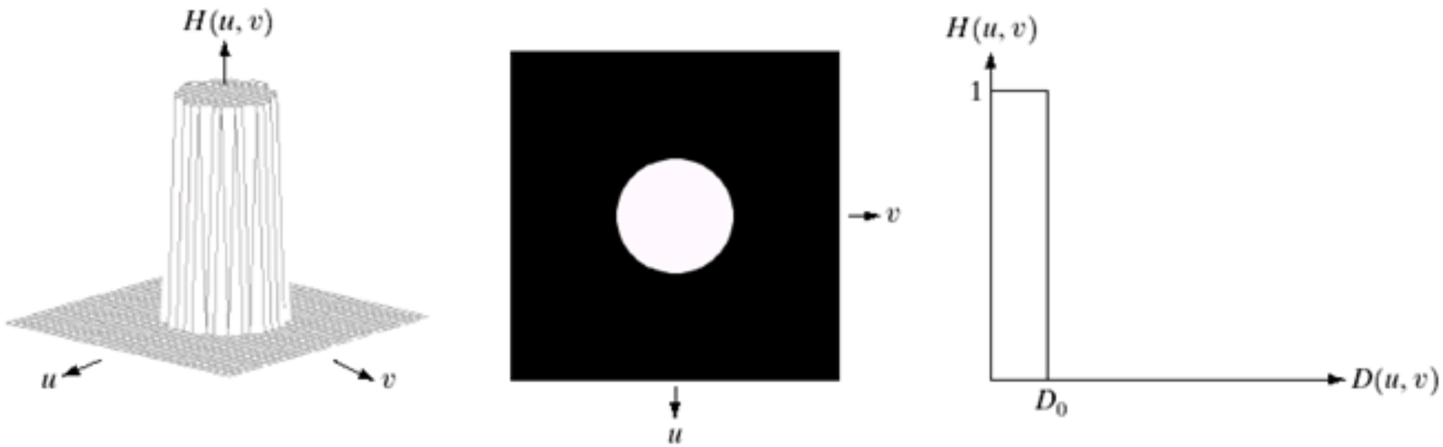

An Example

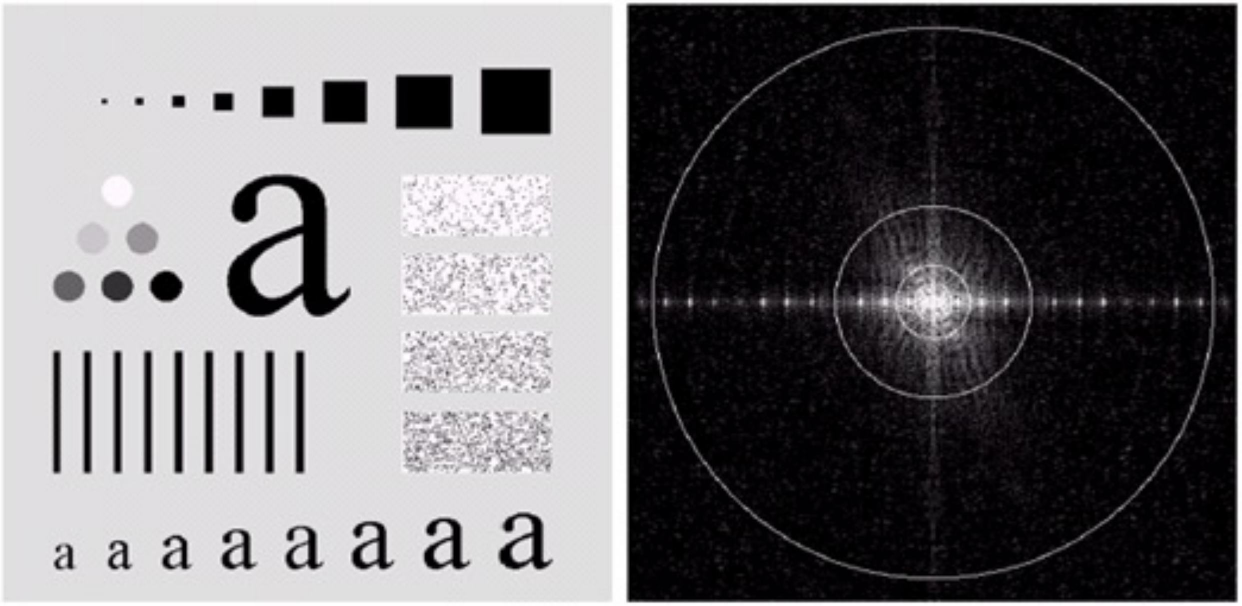

An image of size 500x500 pixels and its Fourier spectrum. The superimposed circles have radii values of 5, 15, 30,80 and 230, which enclose 92.0, 94.6, 96.4, 98.0 and 99.5% of the image power, respectively

An Example …

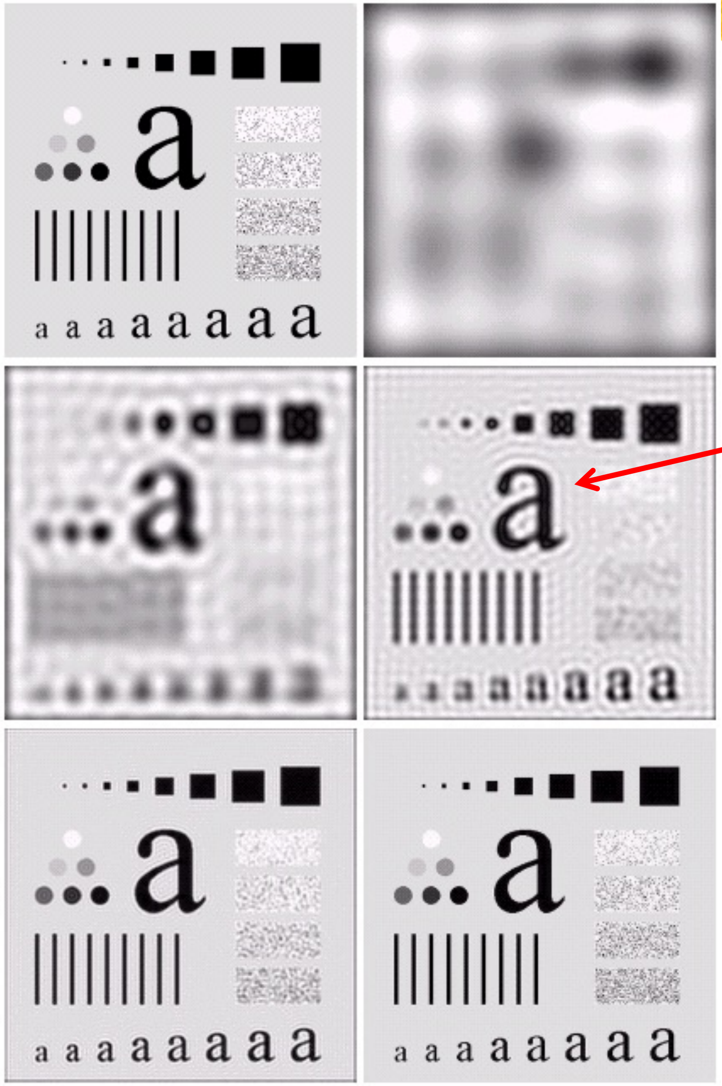

Ringing artefacts

Butterworth Lowpass filters

$$H(u,v) = \frac{1}{1+[D(u,v)/D_0]^{2n}}$$

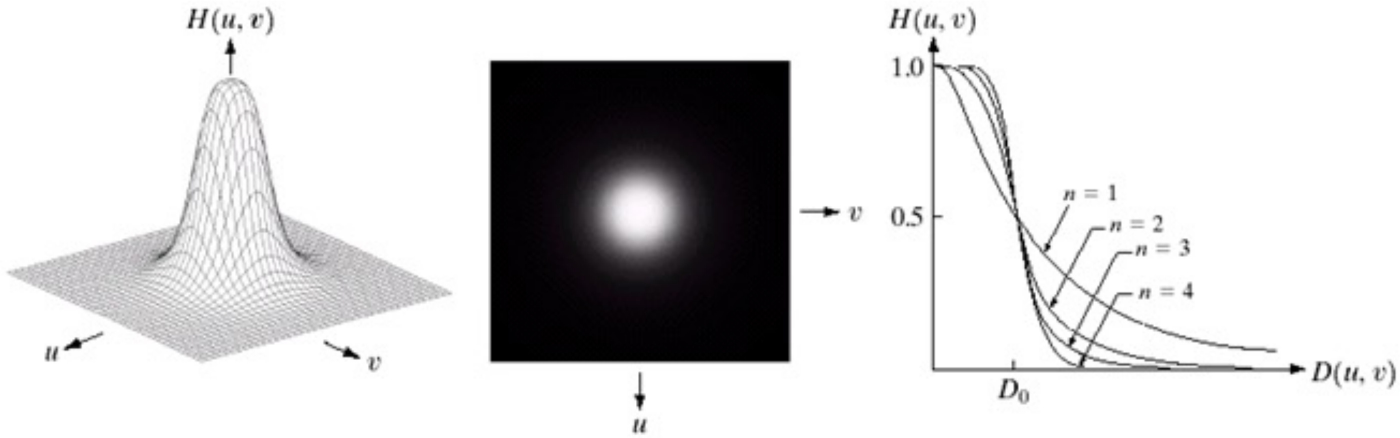

Gaussian Lowpass Filters

$$H(u,v) = e^{-D^2(u,v)/2D_0^2}$$

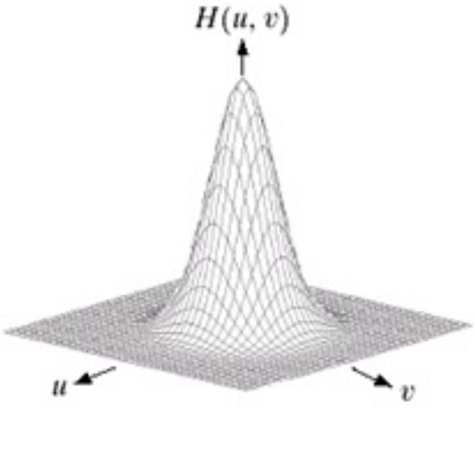

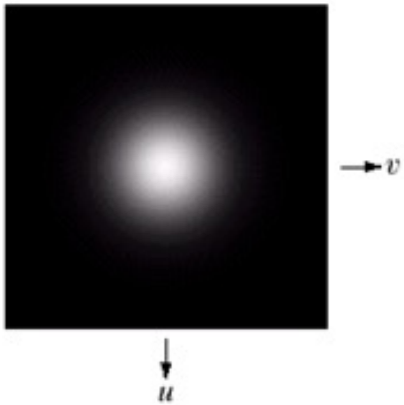

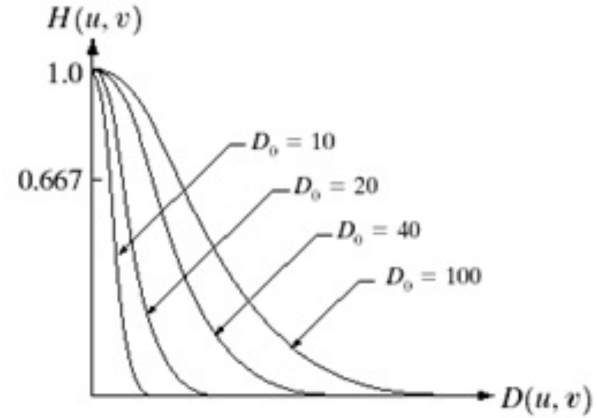

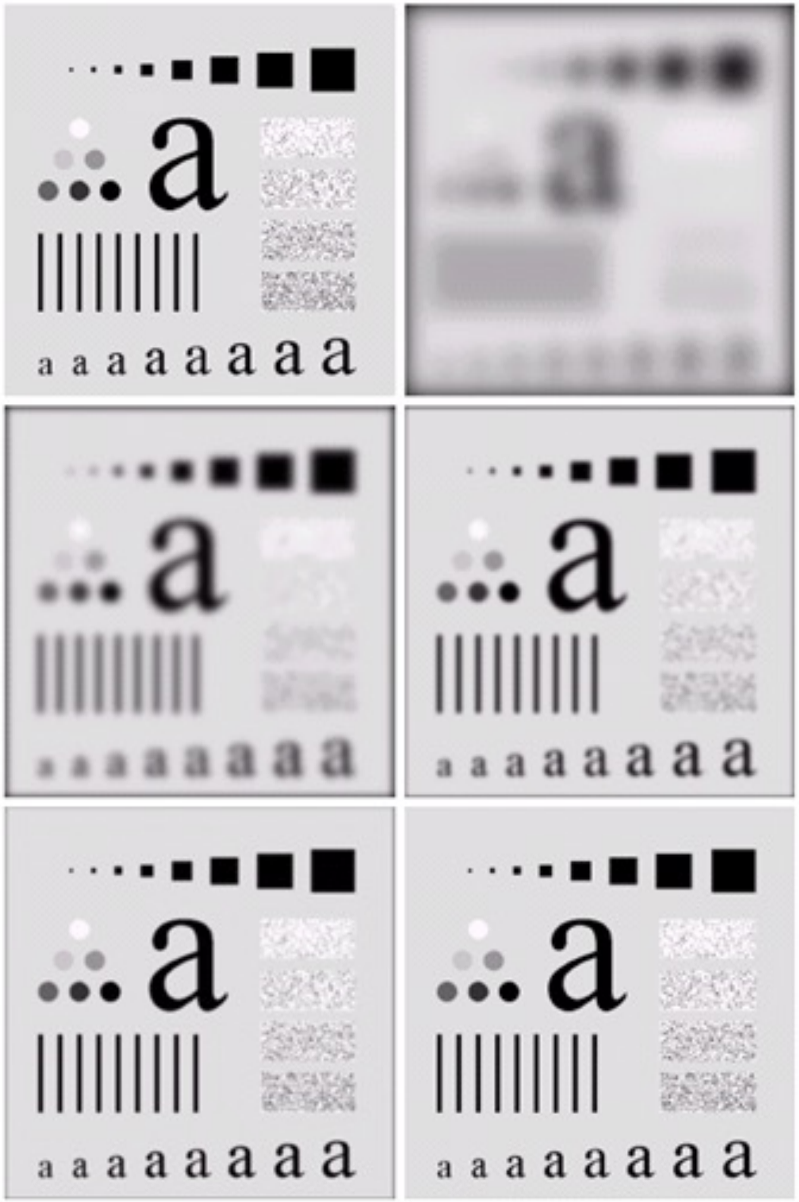

Results of filtering with Gaussian LPF with D_0=5,15,30,80,235, respectively

Highpass Filter

- In general,

$$H_{hp}(u,v) = 1 - H_{lp}(u,v)$$

- Butterworth highpass filters

$$H(u,v) = \frac{1}{1 + [D_0 / D(u,v)]^{2n}}$$

- Gaussian highpass filters

$$H(u,v) = 1 - e^{-D^2(u,v)/2D_0^2}$$

Low Pass Filter (LPF) in Spatial Domain (review)

- As a typical natural image spectrum is dominated by low frequency components
- There are several commonly used LPF kernels which combine reasonable noise suppression properties with low implementation complexity

0.11	0.11	0.11
0.11	0.11	0.11
0.11	0.11	0.11

0.07	0.12	0.07
0.12	0.25	0.12
0.07	0.12	0.07

- As the sum of all filter weights is equal to 1, the filters do not affect the average luminance
- Other odd size filters (5x5, 7x7) can also be used

Spatial vs frequency domain filtering

$$H(u) = Ae^{-u^2/2\sigma^2}$$

$$h(x) = \sqrt{2\pi}\sigma A e^{-2\pi^2\sigma^2 x^2}$$

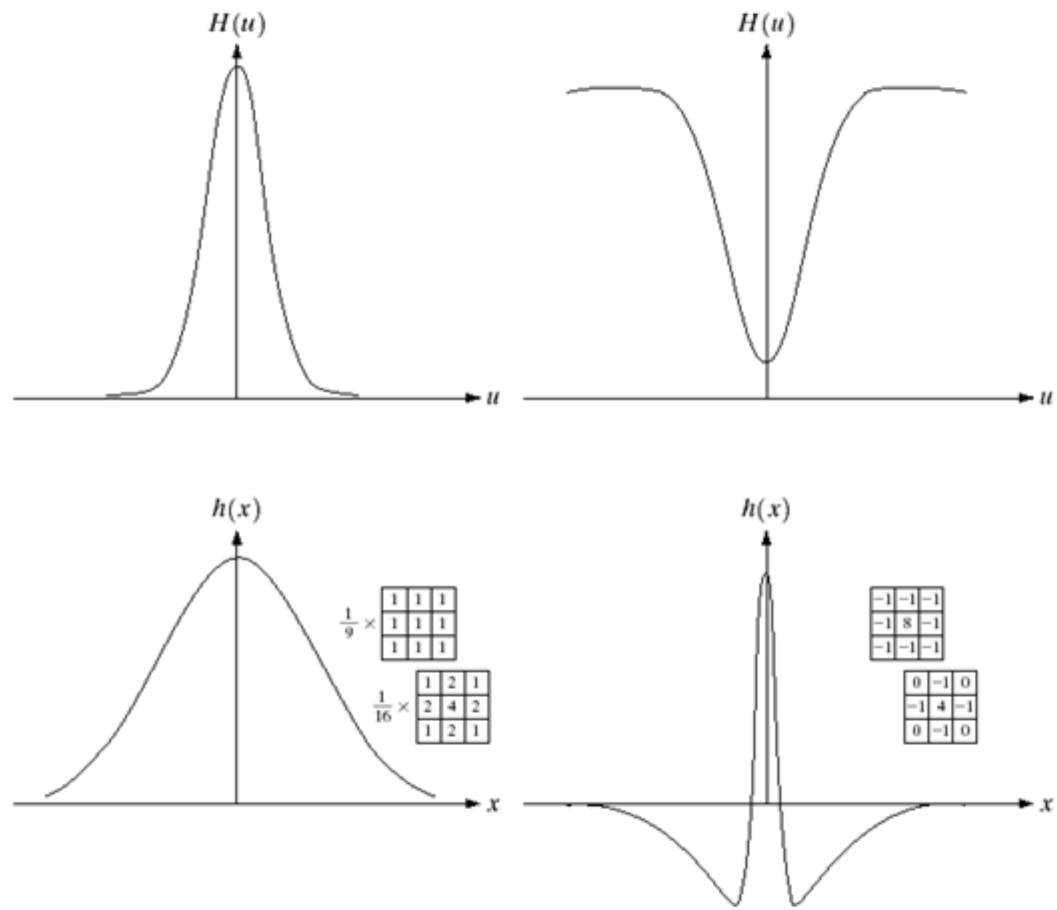

Low Pass Filter (LPF)

Image corrupted by noise

SNR = 20db

LP filtered image

SNR=26db

Low Pass Filter (LPF)

- Suppressing noise, LPF also affects the image
- If an image contains sharp edges of objects, they are blurred

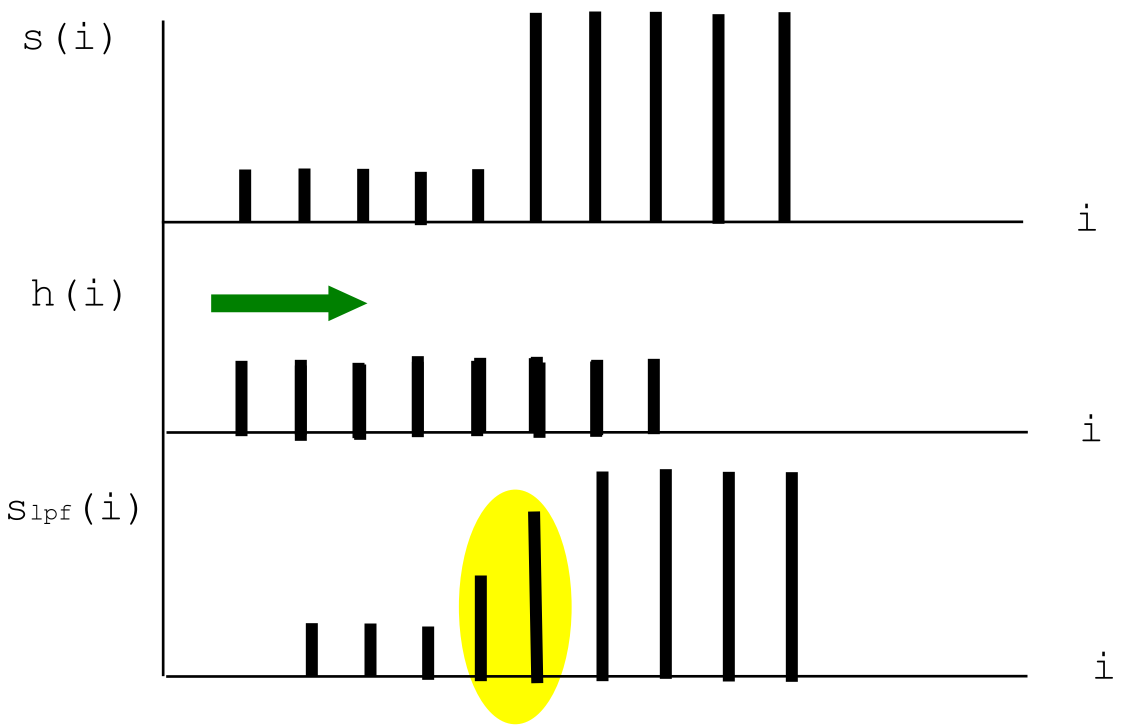

Anisotropic Difusion

- An ideal noise suppression filter must
  1. Efficiently reduce noise level especially in uniform areas
  2. Do not blur object edges

- Anisotropic Diffusion is an adaptive filter that avoids LP filtering across object boundaries

	h_u	
h_l	1.0	h_r
	h_d	

where weights h are calculated based upon the difference between the adjacent image pixels in corresponding directions

Quiz: *the greater*

$|s(i,j) - s(i, j-1)|$ *the smaller/larger* h_l?

- There are several Anisotropic Defusion filters that employ different equations to calculate h

Image Pyramids

- An image may look quite different at different scales. A powerful and simple structure often used in Computer Vision Image Pyramid which represents the image at more than one resolution

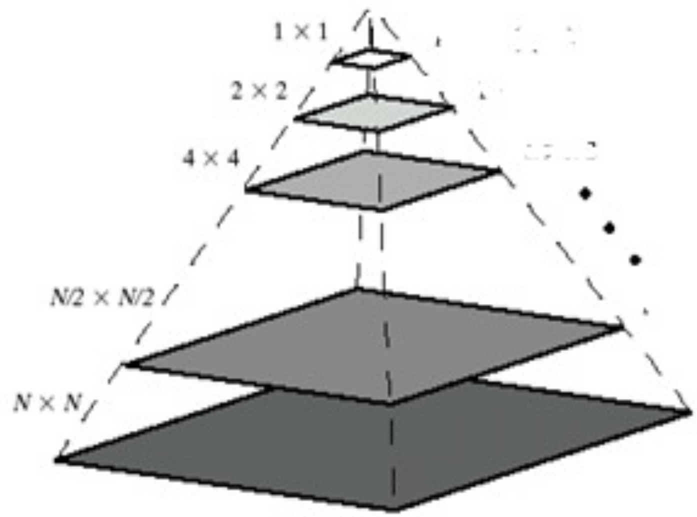

Gaussian Image Pyramids

- Apply Gaussian Lowpass filtering to an image I and down sampling the filtered image by a factor of 2,

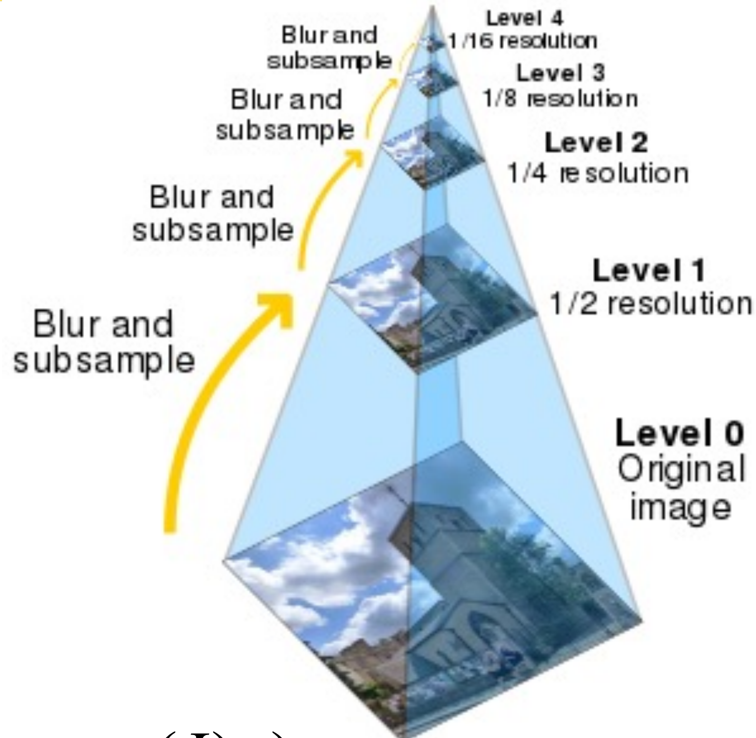

$$P_{Gaussian}(I)_{n+1} = S^{\downarrow}(GLPF_{\sigma} \otimes P_{Gaussian}(I)_{n})$$

$$P_{Gaussian}(I)_{1} = I$$

Suggested Reading

- E. R Davies, Computer Vision: Principles, Algorithms, Applications, Learning, Academic Press; 5th edition; 2017
  - 2.2, 3.1-3.3
- D Forsyth, Computer Vision. A Modern Approach
  - Chapter 4

OpenCV 4.6.0

- Tutorials – Python
  - https://docs.opencv.org/4.6.0/d2/d96/tutorial_py_table_of_contents_imgproc.html
  - Image Thresholding
  - Smoothing Images
  - Image Pyramids
  - Histograms in OpenCV
  - Image Transforms in OpenCV