Topic 5: Noise in Images

5.1 Introduction

One of the most important consideration in digital processing of images is *noise*, in fact it is usually the factor that determines the success or failure of any of the enhancement or reconstruction scheme, most of which fail in the present of significant noise.

In all processing systems we must consider how much of the detected signal can be regarded as *true* and how much is associated with random background events resulting from either the detection or transmission process. These random events are classified under the general topic of *noise*. This *noise* can result from a vast variety of sources, including the discrete nature of radiation, variation in detector sensitivity, photo-graphic *grain* effects, data transmission errors, properties of imaging systems such as air turbulence or water droplets and image quantsiation errors. In each case the properties of the *noise* are different, as are the image processing operations that can be applied to reduce their effects.

5.2 Fixed Pattern Noise

As image sensor consists of many detectors, the most obvious example being a CCD array which is a two-dimensional array of detectors, one per pixel of the detected image. If individual detector do not have identical response, then this fixed pattern detector response will be combined with the detected image. If this fixed pattern is purely additive, then the detected image is just,

$$f(i,j) = s(i,j) + b(i,j)$$

where $s(i,j)$ is the true image and $b(i,j)$ the fixed pattern noise. More commonly there is an multiplicative factor associated with each pixel, so the detected image is,

$$f(i,j) = a(i,j)\,s(i,j) + b(i,j)$$

where $a(i,j)$ is the multiplicative factor, being the *gain* at each pixel. This effect is most evident when the detected image is constant. This is shown in figure 1 which is image from a CCD camera of a blank, almost evenly illuminated blank card. The structure shown has a variation of ± 3 grey level values and is a typical result for a medium cost CCD camera.

If a range of images are taken at different intensity then each sensor can be calibrated and the $a(i,j)$ and $b(i,j)$ coefficient calculated. The fixed pattern noise can then be easily removed. This is a time consuming and difficult calibration, the most difficult been producing a good even illuminated blank object, and it typically only carried out for critical applications, for example in astronomy where absolute intensity values are required, or when the sensor is very poor, typically infra-red detectors where the fixed pattern noise can dominate the image.

5.2.1 Fixed pattern noise in scanned images

Many satellite and aircraft based imaging systems employ a scanning imaging system where the sensor is a one dimensional array which is swept by a movable mirror to system as shown in figure 2. The second dimension of the scan is produced by the movement of the satellite relative to the ground.

Figure 1: Image of an almost evenly illuminated blank card using a CCD camera showing fixed pattern noise.

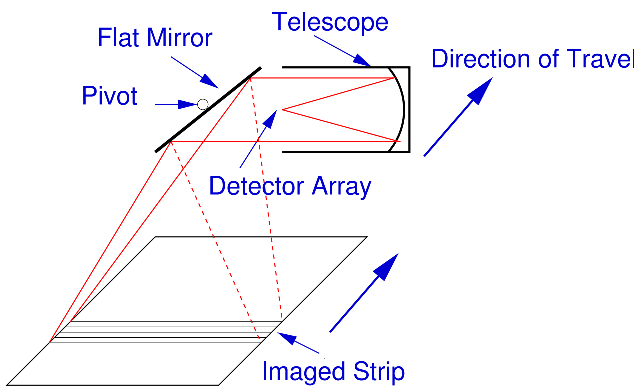

Figure 2: Layout of a typical satellite imaging system.

This scanning scheme results in an image with geometry shown in figure 3 (a), where if the sensor contains K elements, then ever K^{th} line of the image will be detected by the same element of the sensor. In the sensitivity of the individual sensors vary then the image will be striped with horizontal strips, as shown in figure 3 (b), with a repect period of K lines.This is a very common fixed pattern noise effect seen on many satellite or aircraft borne imaging system which can be sucessfully processed.

The first problem to determine the sensor length can be easily found by taking a horizontal projection along the lines so forming the projection

$$p(j) = \sum_{i=0}^{N-1} f(i,j)$$

as shown in figure 4. If we can assume that the horizontal image structure is not periodic[1], then any periodic structure in the projection will give the sensor length. This can be either detected manually of by autocorrelation of $p(j)$.

If we can assume that the projection $p(j)$ is locally smooth we can then fit a smooth function to $p(j)$, usually a low order polynomial, then as shown in figure 5 the local variation of $p(j)$ from this fit is given by the sensor variation. This gives a series of linear equations for the sensor parameters.

[1]True for most natural scenes.

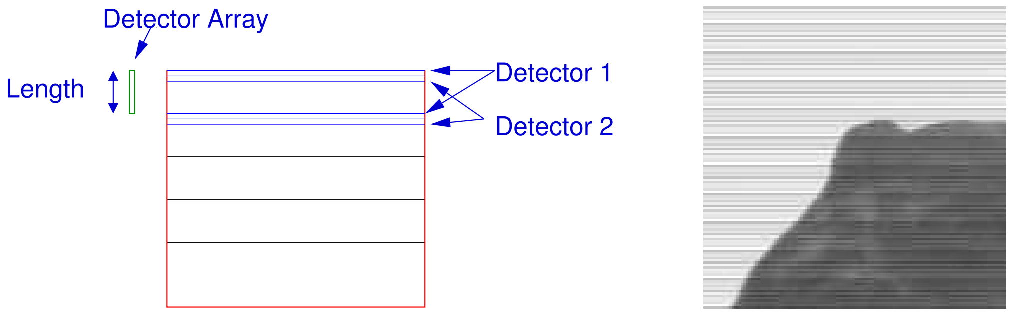

(a) Image Geometry (b) Typical image

Figure 3: Geometry of a image formed by scanning of a short sensor array.

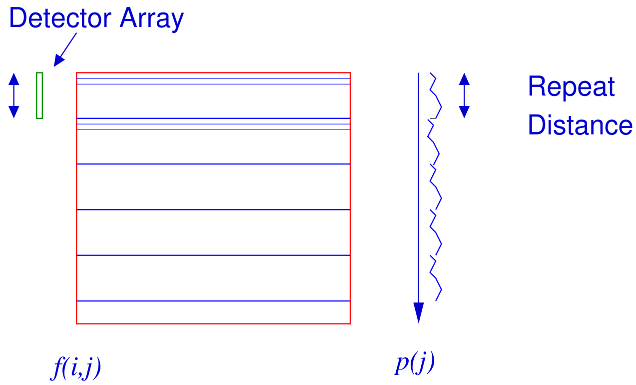

Figure 4: Projection along the lines.

An alternative scheme is to consider the statistics, if an image $f(i,j)$ is scaled by a multipliative factor a and additive term b to give,

$$g(i,j) = a\,f(i,j) + b$$

then the mean of $g(i,j)$ is given by

$$\langle g \rangle = a\,\langle f \rangle + b$$

and it variance by

$$\sigma_g^2 = a^2\,\langle f^2 \rangle - a^2\,\langle f \rangle^2 = a^2\,\sigma_f^2$$

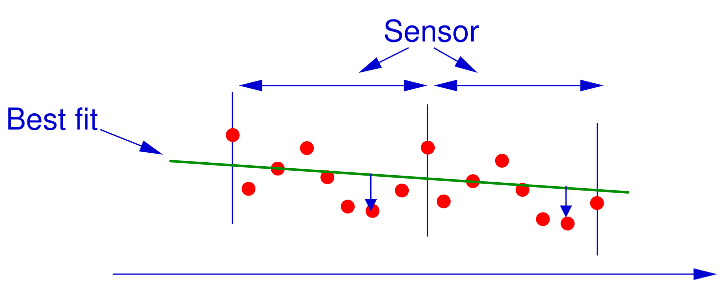

Figure 5: Analysis of the projection, with the sensor variation resulting in deviation from the best- fit line.

where σ_f^2 is the variance of the original image $f(i,j)$.

If we have an images formed by scanning a detector with K elements, as shown in figire 3 then we effectively detect K images of the *same scene*, being,

$$g_k = a_k f + b_k$$

with each detected image being of size $N/K \times N$, being a reduce resolution version of $f(i,j)$, but with different a and b factors which we want to solve for.

The image statistics of mean are variance do *not* depend on the resolution, so and

$$\langle g_k \rangle = a_k \langle f \rangle + b_k \quad \text{and} \quad \sigma_{g_k}^2 = a_k^2 \sigma_f^2$$

so allowing us to solve for a_k and b_k[2]. As with the above fitting scheme, once the a_k and b_k have been obtained the de-striped image $f(i,j)$ can be formed from the subimages.

Both of these schemes are used routinely on satellite images usually by the satellite operator prior to them being released for further analysis. Neither of these scheme give *perfect* destriping and even with the best processed image, close inspection shown some residual horizontal structure from the scanning system.

5.3 Data Drop-out Noise

A common error problem is transmission of digital data is *single bit* transmission errors where a bit it set wrong. In most computer systems where a single bit error, in for example a program code, is totally fatal, complex and computationally expensive data correction and verification is employed to correct such errors, but with the vast amount of data using in digital imaging this level or correction is not usually possible and many digital imaging systems, and especially digital video, suffer from random bit errors.

The effect of random bit errors will depend on which *bit* is corrupted, so for an 8-bit image, corruption of the *least significant* bit will result in an error of $\pm$ in a pixel value, where corruption of the *most significant* will result in a error of ± 128, with the effect of the others given by

Bit Number	Effect
1	± 1
2	± 2
3	± 4
4	± 8
5	± 16
6	± 32
7	± 64
8	± 128

This corruption of the lower significant bits has little visual effect but corruption of the most significant bits turn a *white pixel* to *black* or vice-verse as shown in figure 6 that shown a 128×128 pixel image with (a) 1% of bits corrupted and (b) 5% of bits corrupted. Both images shown random *white* pixels in *black* areas and vice-verse.

[2]The gives the realtive size of of these parameters.

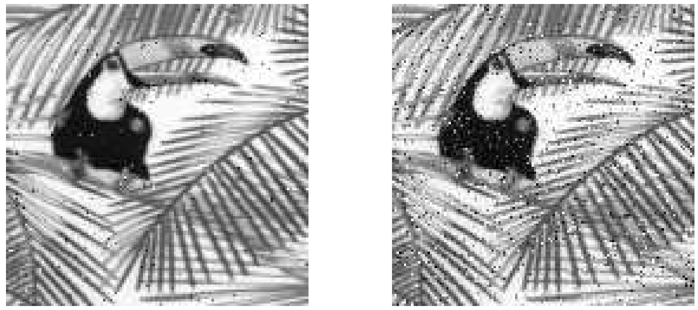
(a) (b)

Figure 6: Images corrupted by single bit data dropout noise, (a) with 1% corruption and (b) with 5% corruption.

This noise is characterised by producing isolated *black* pixels in *white* regions and vice versa, so isolated pixels that vary significantly from their immediate neighbours. This characteristic allows us to determine if a point is corrupted, and so correct it. We will discuss the actual correction techniques in the next section. Such noise it thus statistically different from the image, and so can be recognised and removed, such noise is known as deterministic.

5.4 Detector or Shot Noise

All imaging systems *measure* some type of radiation, either optical intensity, or electron current, or radio power, or X-ray counts etc.; however, due to the discrete nature of all radiation the final system is actually counting particles, typically photons or electrons. It should be noted that many imaging system do not *count* the incident radiation particles directly, but for example with a video camera the incident optical photons generate electrons from a phosphor layer, and the electrons are subsequently *counted* as a current. This is in effect *two* counting processes, one associated with the photon interaction with the phosphor and one with the detection of the resultant electrons, but since the statistical properties of these processes are identical, they can not be separated.

For an imaging system that is assumed to be linear and space invariant, we can treat each image point (or pixel) as independent. Thus each pixel point can be modelled by a *source* of particles and a *detector* where the value of the image pixel is given by the number of detected particles in a given time interval. We can define the *source* brightness as the average, or expected, number of particles incident on the detector in unit time, denoted by $\langle u \rangle$. So if to form the pixel value the detector is observed for a time Δt, then the expected pixel value will be

$$\langle f \rangle = \Delta t \, \langle u \rangle$$

For a particular interval, the actual number of particles detected will be an integer, drawn at random from a probability distribution that characterises the source emission process. For a source of constant brightness $\langle u \rangle$ and a time interval Δt this probability function will be *Poissonian* with mean $\langle f \rangle$. However if we perform this experiment a number of times we will not always count the same number of particles, in fact the actual number of particles detected will be an integer drawn at random from a probability distribution that chararcterises the source emission process. For a source of constant brightness μ and a constant observation interval of Δt, this probability distribution will be *Poissonian* with mean $\langle f \rangle$. So that the probability of counting f particles in one given interval is given by,

$$p(f) = \frac{\langle f \rangle^f \exp(-\langle f \rangle)}{f!} \qquad \text{for } f = 0, 1, 2, \ldots \infty$$

A plot of $p(f)$ for $\langle f \rangle = 4$ is shown in figure 7.

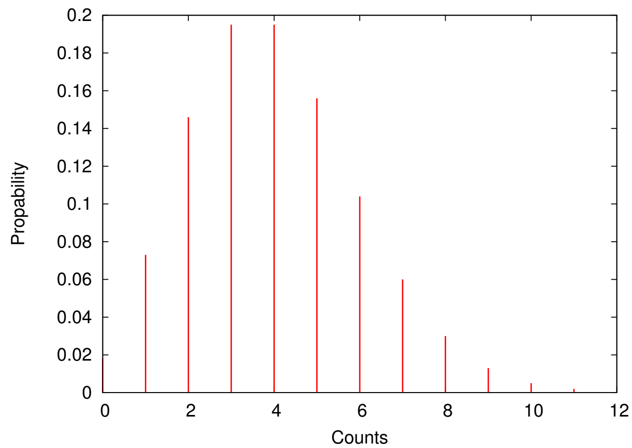

Figure 7: Possionian distribution with $\langle f \rangle = 4$.

The imaging system can now be considered as two dimensional array of independent sources and detectors, as shown in figure 8, so we can define a *mean* or expected values for each pixel, given by

$$\langle f \rangle (i, j)$$

and thus a probability distribution function for each pixel, given by,

$$p(f(i,j)) = \frac{\langle f \rangle (i,j)^{f(i,j)} \exp(-\langle f \rangle (i,j))}{f(i,j)!}$$

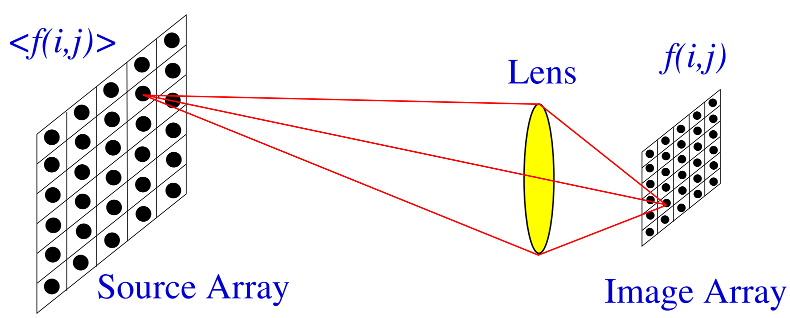

Figure 8: Source and detector arrays

A simulated set of images is shown in figure 9 where the average number of particles, or photons, per pixel is specified. This shown that for low numbers of photons noise dominates but as the number of photons increases image structure become more visible. The analysis of photon limited images is rather complex and beyond a course this level; we will have to make assumption to get any further.

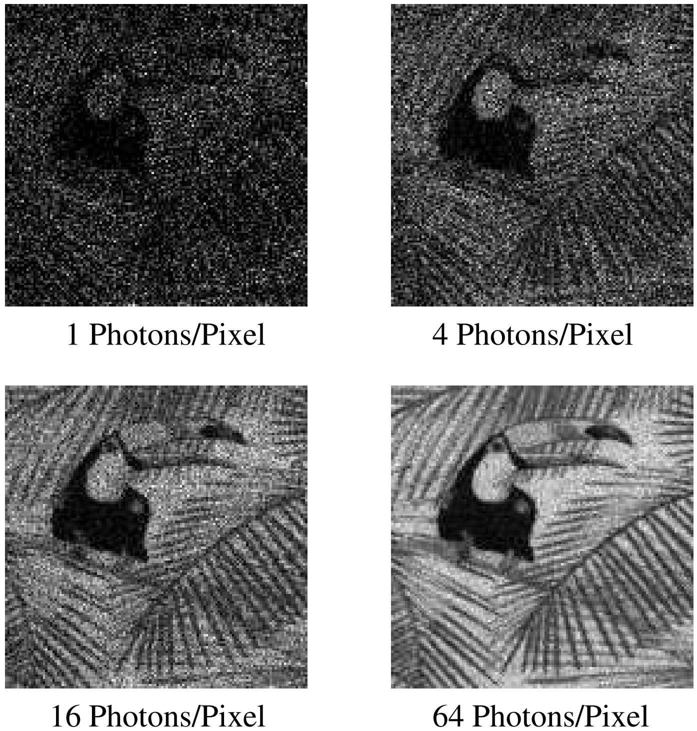

Figure 9: Simulated image image of the toucan with varying numbers of photons per pixel.

5.5 Gaussian Approximation

Since it is discrete valued, the Poissonian probability distribution is mathematically difficult to work with, however by *The Central Limit Theorem* for large expectation values a Poisson distribution is well approximated by a Gaussian distribution of mean and variance equal to the expectation value.

$$p(n) = \frac{u^n \exp(-u)}{n!} \rightarrow \frac{1}{(2\pi u)^{1/2}} \exp\left(\frac{-(n-u)^2}{2u}\right)$$

for large u. The comparison between the Gaussian approximation and the Poisson distribution for $u = 20$ is shown in figure 10 and shows this approximation is valid to within about 1% for values of $u > 20$

In most imaging systems, with the important exception of low light level astronomical images, the expected number of *particles* per pixel can be regarded as large, typically many 100s to 1000s. We can therefore approximate the probability distribution for each image pixel by

$$p(f) = \frac{1}{(2\pi\langle f\rangle)^{1/2}} \exp\left(\frac{-(f-\langle f\rangle)^2}{2\langle f\rangle}\right)$$

The detected value f, can now be regarded as a *true* signal $\langle f\rangle$, which is characteristic of the source and a deviation of *noise* n, characteristic of the detection process, so that for a particular pixel,

$$f = \langle f\rangle + n$$

where n is a random variable with probability distribution

$$p(n) = \frac{1}{(2\pi\langle f\rangle)^{1/2}} \exp\left(\frac{-n^2}{2\langle f\rangle}\right)$$

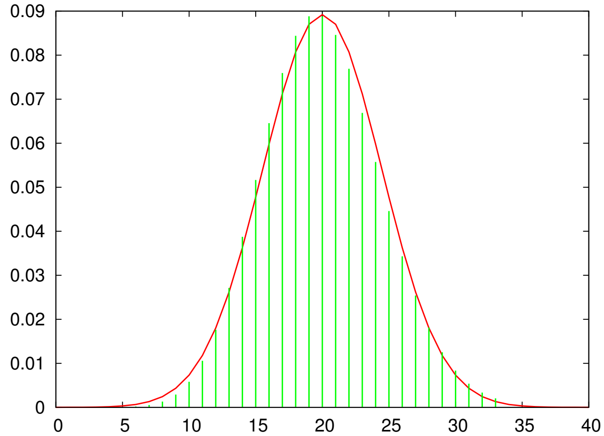

Figure 10: Comparison of Poisson and Gaussian distribution for mean of 20

Therefore for a two dimensional image we have the model that the detected image $f(i,j)$ is given by

$$f(i,j) = \langle f\rangle(i,j) + n(i,j)$$

where *each* additive noise term is characterised by a probability distribution function with mean and variance given by the expectation value of the signal at that pixel, ie.

$$p(n(i,j)) = \frac{1}{(2\pi\langle f\rangle(i,j))^{1/2}} \exp\left(\frac{-n(i,j)^2}{2\langle f\rangle(i,j)}\right)$$

So that the *noise* term $n(i,j)$ is *dependent* on the signal.

5.5.1 Low contrast approximation

If however we assume that the image is of *low-contrast* so that the pixel expectations $\langle f\rangle(i,j)$ can be regarded approximately constant, then the probability distribution of the noise can be taken as

$$p(n(i,j)) = \frac{1}{(2\pi\mu)^{1/2}} \exp\left(\frac{-n(i,j)^2}{2\mu}\right)$$

where μ is the average of the expectation values across the image, given by

$$\mu = \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \langle f\rangle(i,j) = \langle(\langle f\rangle(i,j))\rangle$$

In this case the probability distribution of the *noise* is a Gaussian which only depends on the *average* of the “true” signal and *not* on the details of the “true” signal, so it can be regarded as *independent* of the signal. Under these assumptions, we have a *noise* model where the noise is additive, independent of the image and is characterised by a zero mean Gaussian probability distribution with variance equal to the mean of the true signal.

This assumption can be considered from physical grounds as follows. If a pixel has an expectation value of $\langle f\rangle$, then the measured value of the pixel f, will be drawn from a Gaussian probability distribution centred about $\langle f\rangle$ with standard deviation $\sqrt{\langle f\rangle}$, therefore there is a 98% probability that the value of f will be in the range $\langle f\rangle \pm 2\sqrt{\langle f\rangle}$. Now if $\langle f\rangle$ is large and does not vary significantly across the image then $\sqrt{\langle f \rangle}$ is almost constant, so it can be taken as being independent of the image.

In many image collection systems these approximations are valid, particularly in thermal infra-red, electron microscope, X-ray microscope, medical X-rays, NMR images, PET images and almost all video systems. In these systems the *particles* being counted are different, ranging from low energy infra-red photons to gamma photons and thermal electrons to high energy elastic electrons in high voltage microscope. In many systems there are more than one *particle* counting process, but assuming linearity, the overall effect is still modelled by the particle counting model.

This additive Gaussian noise model is *not* valid where images either contain very few counted particles per pixel or the images are very high contrast. Low light level systems, such as image intensifiers and optical telescope systems do not obey these assumptions. The optical telescope operates in two possible modes, the *long exposure* and the *short exposure*. A long exposure image is formed over several minutes, being recorded on photographic plate or CCD array system. In this case there are a large number of photons incident per pixel but the image is high contrast, typically 16 bits, so the Gaussian signal dependent noise model can be used. Short exposure images use ultra fast photographic material or *photon-cameras* which measure individual light photons. Photographic material can be used with an expectation of about $12 \rightarrow 20$ photons per pixel while the "photon-camera" records the time and coordinate of *each* photon, and operates at a expectation values of 0.1 photons per pixel. In these cases the full Poisson noise model must be used, which complicates the processing significantly.

5.6 Properties of Gaussian Additive Noise

In the case considered above where we assume an *additive* noise model, the detected image $f(i,j)$ is given by,

$$f(i,j) = s(i,j) + n(i,j)$$

where the term $s(i,j) = \langle f \rangle(i,j)$ is the *true signal* and $n(i,j)$ is the *noise* term. The noise term is a Gaussian random distribution with zero mean, so that,

$$\langle n(i,j) \rangle = 0 \quad \text{and} \quad \langle |n(i,j)|^2 \rangle = \sigma_n^2$$

where we have $\sigma_n^2 = \langle s(i,j) \rangle$. We also have that the noise distribution is completely *independent* of the image, and is thus *uncorrelated* with the image. This can be expressed mathematically as,

$$\langle s(i,j)\, n(i,j) \rangle = 0$$

which is just the zero shifted correlation. The variance of $f(i,j)$ is

$$\sigma_f^2 = \left\langle |f(i,j) - \langle f(i,j) \rangle|^2 \right\rangle$$

so by expansion we get that the variance of the detected image $f(i,j)$ is

$$\sigma_f^2 = \sigma_s^2 + \sigma_n^2$$

where σ_s^2 is the variance of the signal; so the addition of noise alters the variance and not the mean.

We can Fourier transform the input image $f(i,j)$ to get,

$$F(k,l) = S(k,l) + N(k,l)$$

where $N(k,l)$ is the Fourier transform of the noise distribution $n(i,j)$. Noting that the noise distribution is zero meaned, then from Parseval's theorem we have that

$$\langle |N(k,l)|^2 \rangle = \langle |n(i,j)|^2 \rangle = \sigma_n^2$$

This noise is assumed to be a totally random process, characterised only by its variance, so that

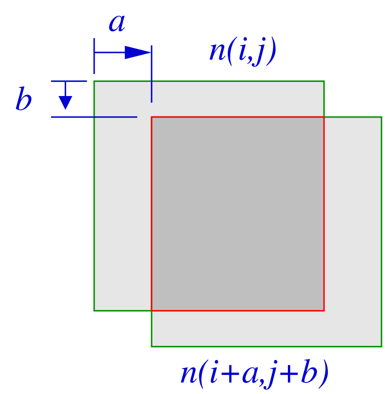

Figure 11: Autocorrelation of the noise with itself.

there is no correlation between separated noise points, and so the noise is not correlated with itself so that the auto-correlation,shown in figure 11, of the noise is given by,

$$\langle n(i,j)\, n(i+a,j+b) \rangle = \sigma_n^2 \delta_{a,b}$$

The auto-correlation is therefore a delta function located at zero. The power spectrum of the noise, $|N(k,l)|^2$, is now the inverse Fourier transform of the auto-correlation function, which is given by[3]

$$|N(k,l)|^2 = \text{constant}$$

In practice, the noise distribution is a random process, and for any realisation of the $n(i,j)$ the power spectrum will only be approximately constant as shown in figure 12.

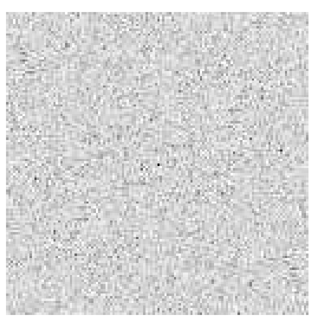

Figure 12: Fourier transform of one realisation of Gaussian noise.

However if we average over a range of realisations, $n_p(i,j)$ each with Fourier transform $N_p(k,l)$ then this will be constant. Also noting which shows that mean of the power spectrum is given by σ_n^2 we get that,

$$\frac{1}{P}\sum_{p=1}^{P} |N_p(k,l)|^2 = \sigma_n^2$$

[3]See *Fourier Transform* booklet.

This shows that the the Fourier transform of the noise $N(k,l)$ will have approximately constant amplitude, of σ_n at all spatial frequencies, which is typically referred to as *white-noise*.

In most imaging systems, the Fourier transform of the true signal, $S(k,l)$ is very sharply peaked about the low spatial frequencies while the Fourier transform of the noise is approximately constant at all spatial frequencies. Therefore the noise has little effect at low spatial frequencies, but may become dominant at high spatial frequencies where $S(k,l)$ is small as shown in figure 13.

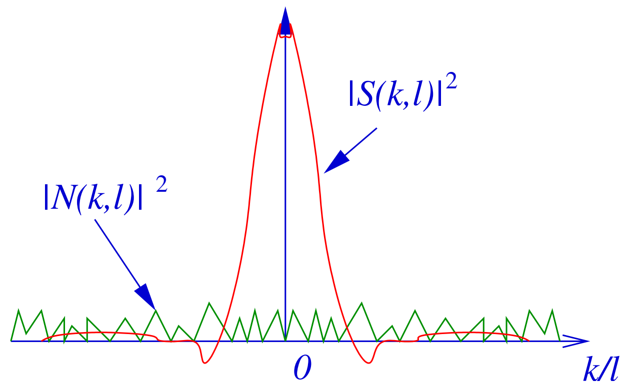

Figure 13: The effect of additive noise in Fourier space.

This effect is shown for images in figure 14 where additive noise has been added to the `toucan` image. This results in a *granular* texture added to the image. In Fourier space in figure 14 (b), there is extra high frequency power, which is reduced with a *low-pass*[4], then the low-pass filtered image has significantly reduced granular noise. The image is however significantly *smoothed* and the edges, where are also associated with high spatial frequency, significantly blurred. This is a fundamental problem with additive noise since it randomly corrupts the image it *cannot* be removed without degrading the underlying image.

6 Signal to Noise Ratio

When we have a linear, space invariant imaging system corrupted by additive Gaussian noise, it is frequently useful to characterise the *effect* of the noise by a single value, referred to as *Signal to Noise Ratio*, or SNR. In this case the detected can be written as

$$f(i,j) = s(i,j) + n(i,j)$$

then if we define the variances of the signal, $s(i,j)$, and the noise, $n(i,j)$, by,

$$\begin{aligned} \sigma_s^2 &= \langle |s(i,j) - \langle s(i,j) \rangle|^2 \rangle \\ \sigma_n^2 &= \langle |n(i,j)|^2 \rangle \end{aligned}$$

and then the SNR can be defined[5] by;

$$\text{SNR} = \frac{\sigma_s}{\sigma_n}$$

[4] This is covered in detail ion the next section.

[5] There are many definitions of SNR in the Image and Signal Processing literature, which are either the square, or the log of this definition.

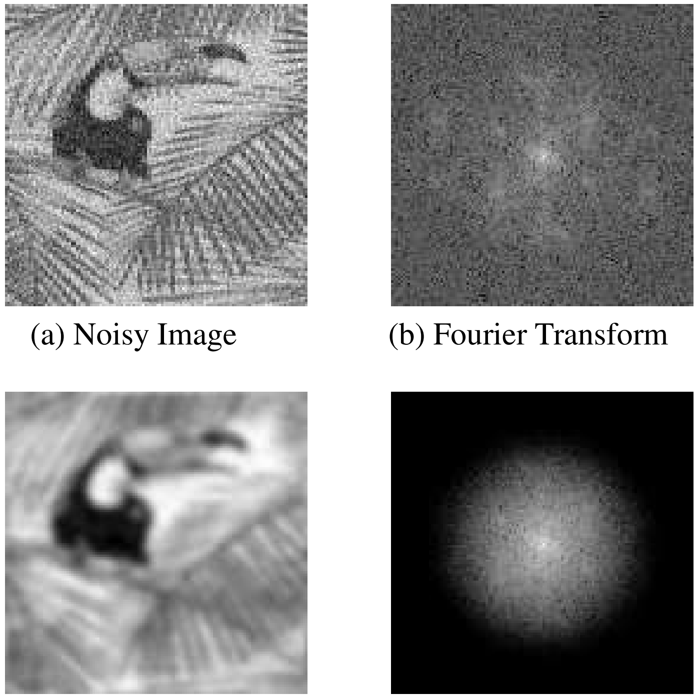

(a) Noisy Image (b) Fourier Transform

(c) Low-pass filtered (d) Fourier Transform

Figure 14: Example of low-pass filtering a noise image.

being the ratio of the standard deviations of the true signal and the additive noise.

Since the measured image is $f(i,j)$ it is useful to specify the SNR in terms of the detected σ_f, the standard deviation of $f(i,j)$ as follows; $\langle n(i,j)\rangle = 0$, so that

$$\langle f(i,j)\rangle = \langle s(i,j)\rangle$$

and thus by expansion of the definition of σ_f^2, we get that,

$$\sigma_f^2 = \sigma_s^2 + \sigma_n^2$$

so that the SNR can also be written as

$$\mathrm{SNR} = \sqrt{\frac{\sigma_f^2}{\sigma_n^2} - 1}$$

6.1 Calculation of SNR from a Detected Image

To calculate the SNR of an image, from the above expression, we are required to find, or estimate, two of the three sigma values $\sigma_f, \sigma_s, \sigma_n$. From a detected image $f(i,j)$, the standard deviation of the image σ_f, is easily obtained, but neither the standard deviation of the *true* signal nor the noise is available directly. For whole image in figure 15, $\sigma_f^2 = 5287$.

If we have an image where we know aprior, that there is *no* signal, or that the signal is constant, ie. an area of sky or water as shown in figure 15. Then we can assume that all variances within this region are caused by noise and the standard deviation within this region is σ_n, in this case $\sigma_n^2 = 1.85$, hence allowing an estimate for the $SNR \approx 53.4$ This method is image dependent and

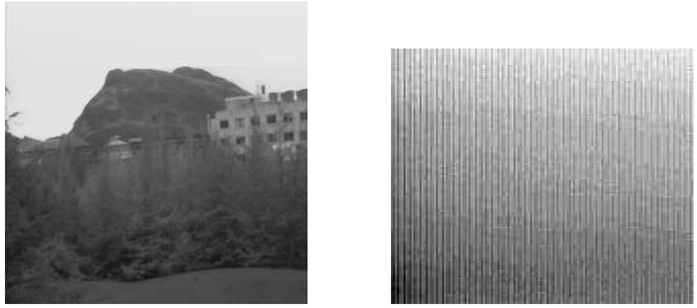

Whole Image Piece of Sky

Figure 15: Image containing a blank area, sky in this case, used to determine Signal to Noise Ratio of SNR ≈ 53.4.

assumes that a constant part of the image can be found, either manually, or by an automated process.

Consider the case where we have *two* realisations of the same image, ie. two images that contain the identical signal but since they were detected completely separately will have totally independent realisations of the same noise process. Two such images can be obtained from adjacent frames of a video signal of a stationary scene. We have now detected two images, $f(i,j)$ and $g(i,j)$ given by,

$$\begin{aligned} f(i,j) &= s(i,j)+n(i,j) \\ g(i,j) &= s(i,j)+m(i,j) \end{aligned}$$

where $n(i,j)$ and $m(i,j)$ are different realisations of the same noise process, so that we have that,

$$\begin{aligned} \langle n(i,j)\rangle &= \langle m(i,j)\rangle = 0 \\ \langle |n(i,j)|^2\rangle &= \langle |m(i,j)|^2\rangle = \sigma_n^2 \\ \langle m(i,j)\,n(i,j)\rangle &= 0 \end{aligned}$$

If we now the *normalised correlation* $\hat{C}$ between two images as defined by

$$\hat{C} = \frac{\langle (fg - \langle f\rangle\langle g\rangle)\rangle}{\left[\langle |f-\langle f\rangle|^2\rangle\langle |g-\langle g\rangle|^2\rangle\right]^{1/2}}$$

expressions for $f(i,j)$ and $g(i,j)$ are substituted into this relation, noting that $\langle f\rangle = \langle g\rangle = \langle s\rangle$ and all term containing $\langle n\rangle$ and $\langle m\rangle$ are zero, then we get, after a bit of manipulation, that

$$\hat{C} = \frac{\sigma_s^2}{\sigma_s^2+\sigma_n^2}$$

so then the SNR can be found directly from

$$\mathrm{SNR} = \sqrt{\frac{\hat{C}}{1-\hat{C}}}$$

Therefore for two realisations of the same image that differ only by random additive noise, the SNR can be found from the image correlation, as defined above. Clearly if we have more than two images of the same same scene, then we can form a SNR between each pair of images, and obtain an improved estimate for the SNR of the imaging system by averaging.

This measure of SNR is useful in giving an indication of the noise in an image, but the exact visual effect of such noise is highly image dependent. Figure 4.5 shown same simple image with various levels of additive noise to produce SNR of 1, 2, 4 and 8 respectively, showing a marked deterioration in visual quality for SNR of less than 4.

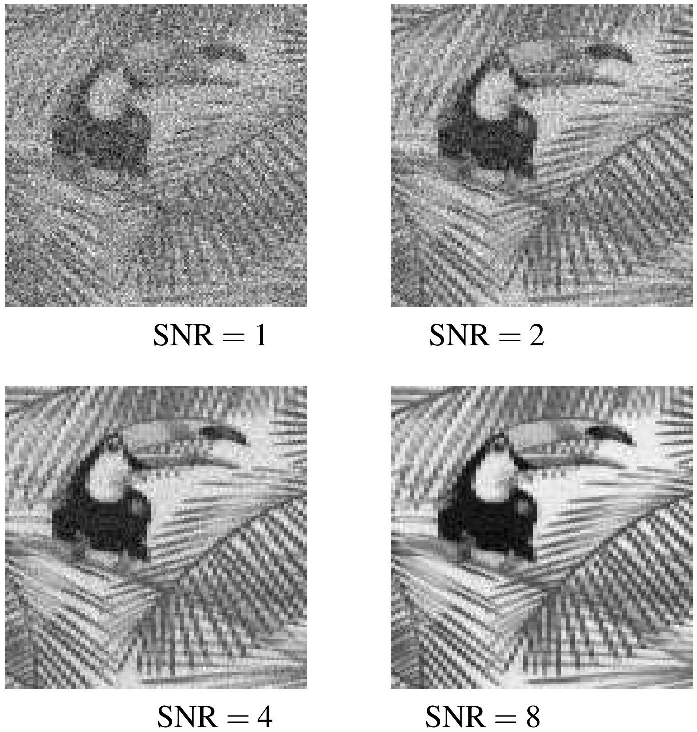

Figure 16: Images with additive Gaussian noise with a) SNR = 1, b) SNR = 2, c) SNR = 4, and d) SNR = 8.

Workshop Questions

6.1 Shape of Poisson Distribution

Use *Maple* or *gnuplot* to plot the Poisson distribution for means of 1,2,4 and 8, and comment on the shapes.

6.2 Poisson to Gaussian

Given that the Poisson distribution can be approximated by a Gaussian of of the form

$$\frac{u^n \exp(-u)}{n!} \rightarrow \frac{1}{\sqrt{2\pi u}} \exp\left(-\frac{(n-u)^2}{2u}\right)$$

when u is large. Plot the Poisson and corresponding Gaussian distribution for a range of u in the range $8 \rightarrow 20$, and comment on the result.

6.3 Variance with Noise

If an image is corrupted with signal independent additive noise being given by,

$$f(i,j) = s(i,j) + n(i,j)$$

where $s(i,j)$ is the *true image* or signal, and $n(i,j)$ is the noise then show that:

$$\sigma_f^2 = \sigma_s^2 + \sigma_n^2$$

where σ_f^2, σ_s^2 and σ_n^2 are the variances of the image, signal and noise respectively.

6.4 Calculating SNR by Correlation

If you have two images of the same scene taken at different times, show that the SNR can be estimated by

$$\text{SNR} = \sqrt{\frac{c}{1-c}}$$

where c is the *Normalised Correlation* between the two images given by

$$c = \frac{\langle f - \langle f \rangle \rangle \langle g - \langle g \rangle \rangle}{\left(\langle |f - \langle f \rangle|^2 \rangle\right)^{\frac{1}{2}} \left(\langle |g - \langle g \rangle|^2 \rangle\right)^{\frac{1}{2}}}$$

where $f(i,j)$ and $g(i,j)$ are the two images.

6.5 Variance of Noise

Extend the techniques above for calculating SNR in question 6.4 to give a scheme for calculating the variance of the noise in an image.