Data, Noise and Signal (for Astronomers)

S. R. Kulkarni

January 20, 2008–February 27, 2017

Dedicated to the animals on this planet.

Pirvi Press, Pasadena, California

Notes to my Friends

I am writing a book entitled "Signal Processing & Data Analysis For Astronomers". The goal is to equip a first year graduate student in astronomy with the basics of probablity and signal processing that is essential to the pursuit of observational astronomy. Statistical inference and modeling is only discussed on the margins. I do this for the simple reason that inference is not my strength and I feel that I do not possess sufficient depth to justify its inclusion in the book.

Please read the Preface and Notes to Using the Book first. Then please review the present table of contents. Please comment what else you may wish to see.

I attach a sample chapter (Order Statistics) as an illustrative chapter. Is the style readable? Can you suggest more homeoworks?

Thanks
Shri

Contents

Preface

Scientists obtain data with the hope of discovery new phenomena or detecting patterns that may then give an insight into how nature functions. However, the first thing that they encounter is likely noise or a very large background or foreground. Only after careful sifting of the data, in particular after eliminating or throughly understanding the noise, does one usually obtain a reliable sensible. Thus, the path to discovery usually goes through a thicket of noise.

I have written this book with a single purpose: distill the essence of techniques and methodologies relevant to astronomical instrumentation, methodology and observations in one book. The book is aimed at a first year graduate student in astronomy. The book is pedagogical. In this respect this book differs from the many fine books on signal processing written by engineers And books written by mathematicians and statisticians are, well, mathematical.

I deliberately focus on the theoretical underpinnings. It is critical that the reader develop a sense of numbers for both experiment planning as well as analysis. Eventually, numerical simulations or some variant of bootstrap will likely give the most definitive result. However, simulations without a strong understanding is unlikely to be productive nor efficient use of one's time.

It is my experience that real examples provide the best motivation to learn. It is not a common sight to find a scientist tossing a coin or selecting colored balls from a bowl. On the other hand, pedagogy is critical and there is no escape from hard thinking and head banging. For this the book is conceptually divided into two parts: the basic concepts early on, followed by "applied" topics.

The writing of the book was a process of self learning with the hope that I may be able to present a fresh perspective and (very) occasionally a new result. All the figures and tables in this book have been either drawn or programmed by the author.

The best way to read this book is to find a long stretch of time (wintry days or summer holidays) and blaze through it. True enlightenment is only possible via direct understanding, ergo work out every equation and *do the homeworks*.

I would consider it an outstanding success if the book (though written with a strong astronomical bias) would be found useful by other scientists (geologists, biophysicists etc). The author would be delighted to hear of success stories. I continue to seek interesting real problems to be included in the next version of the book.

Notes to Using the Book

The book is conceptually organized into two parts. The first part covers the pedagogical foundation of probability, statistics and signal processing. The second part has applications that are drawn from astronomy and with a view of having more than one area (e.g. probablity and signal processing) as well as rather specialized topics of some interest to astronomy (e.g. Kalman filtering, Allan Variance, Phasor and Bispectrum).

Each chapter has homeworks in up to three different: pedagogy (P), standard homework (H) and interesting open-ended or problems without known solutions ($\star$). The pedagogical problems are expected reinforce the content and the homework problems are designed to be moderately useful and usually require a good knowledge of the material in the first part.

This is a textbook and not a research paper. I have avoided giving references within the main text. Key references are given at the end of each chapter. I would like to acknowledge the following fundamental references: the Wikipedia (which is amazingly comprehensive and adequate pedagogical explanations) and Eric Weinstein's MathWorld. *Numerical Recipes* by Press, Teukolsky, Vetterling & Flannery (1986–2007) is the standard reference that the reader is advised for explanations of a wide variety of algorithms and their implementation. *The Fast Fourier Transform* by E. O. Brigham is the standard textbook for a student wishing to learn about Fourier transforms and their applications.

Each chapter ends with terse references to papers which either present a valuable extension of methodology or an illustrative application of the method under discussion. I have sought the help of colleagues (and still continue to do so) to send me suggestions for interesting examples. However, I must stress that the goal is not a *comprehensive* treatise or compendium of numerical methods or analysis but pedagogy.

Notation

For a textbook to be useful uniformity in style is critical. I have paid some attention to this issue and below you will find the guiding rules for notation.[1]

The probability density function is denoted by lower case and the cumulative function by upper case. For example, $p(x)$ is the probability for the random variate to lie in the interval $[x, x+dx]$ and $P(x) = \int_{-\infty}^{x} p(x)dx$ is the cumulative function. Population mean and variance is indicated by μ and V

[1]THIS IS WORK IN PROGRESS SINCE HARMONIZING NOTATION ACROSS THE PLANNED 20 CHAPTERS WILL REQUIRE SOME TRIAL AND ERROR.

and $\sigma^2 = V$.

The forward Fourier transform is

$$H(f) = \int_{-\infty}^{+\infty} h(t) \exp(-j2\pi ft)dt \tag{1}$$

and the inverse transform is

$$g(t) = \int_{-\infty}^{+\infty} G(f) \exp(-j2\pi ft)df. \tag{2}$$

This is the convention used in Brigham's book (*ibid*) and I find it far preferable compared to the choice in Physics textbooks (with ω instead of f). Both i and j are used for $\sqrt{-1}$.

Chapter 1

Algebra of Probability

1.1 Axioms of Probability

1.2 Bayes Theorem

1.3 Mean, Median & Variance

The probability density function or probability distribution function completely characterizes the signal or noise (or both). However, as in ordinary life, single value descriptors are generally useful ("The weather in Los Angeles is ideal" or "The mean temperature in Los Angeles is 65 °F"). Mean and median are commonly used single value descriptors for probability distributions.

The median is the 50-percentile value. Consider a 101-strong group of fifth graders. Ask them to organize so that the person to the right is taller (and at best equal but not shorter). The median height of this group is the height of the 51st student. Mathematically the median is specified by the equation

$$P(x = m) = \int_{-\infty}^{m} p(x)dx = \frac{1}{2}. \tag{1.1}$$

where $P(x)$ is the cumulative function of x.

A seemingly separate attribute of the median is that it also minimizes the absolute deviation of a sample, i.e.

$$Y = E(|x - c|) = \int |x - c|p(x)dx \tag{1.2}$$

is minimized when $c = m$.

A useful relation to note is that the median and mean cannot deviate by more than the rms (and discussed below):

$$|\mu - m| \leq \sigma. \tag{1.3}$$

The arithmetic mean (usually shortened to "mean") is defined as

$$\mu = \int_{-\infty}^{\infty} p(x) x dx. \tag{1.4}$$

The arithmetic mean has many desirable properties: unbiased and efficient. In addition, the mean minimizes the squared deviation of a given sample. The sum of the squared deviation of the sample can be written as

$$Y = \int (x - c)^2 p(x) dx. \tag{1.5}$$

We seek that value of c which minimizes Y. Differentiating Equation 1.5 with respect to c and setting the result to zero we find $c = \mu$.

An alternative proof which is easier to understand is along the following lines:

$$\begin{aligned} Y &= \int (x - c)^2 p(x) dx \\ &= \int [(x - \mu) + (\mu - c)]^2 p(x) dx \\ &= \sigma^2 + (c - \mu)^2 \end{aligned} \tag{1.6}$$

where

$$\sigma^2 = \int (x - \mu)^2 p(x) dx. \tag{1.7}$$

Clearly, Y is minimized when $c = \mu$ and the minimum value is the "variance" of the distribution (the second central moment). The square root is some times called "sigma", "standard deviation" or "rms". We will revisit these names when discussing samples of finite size.

For a probability density function with a single well defined peak ("unimodal") the mean and median provide a numeric measure of the location of the peak. The third indicator is the "mode", $x = M$, which is the the value at which $p(x)$ is maximum. The primary (and perhaps only) usage of the mode in astronomy is to determine the "sky" level in optical and IR astronomy.

`Other topics: harmonic mean, geometric mean. When are these useful?`

`Exercise: A Light Puzzler.` For a set of five integers, the mean, median and mode are 4, 1 and 5. Infer the values of the integers.

`Exercise: Empirical Relation.` For probability distributions with a single mode ("unimodal") and of no more than "moderate" asymmetry is the following empirical relation appears to be noted:

$$\text{mean} - \text{mode} = 3(\text{mean} - \text{median}). \tag{1.8}$$

As you go along the text, test this claim.

1.3.1 Relation between Mean and Median

The Bienaymé-Tchebycheff inequalities are very powerful relations with application to general relation between mean and median and the central limit theorem.

Consider a random variable, y which can attain only zero or positive values. The expectation value of y is then

$$\begin{aligned} E(y) &= \int_0^\infty y p_y(y) dy \\ &\geq \int_{y_0}^\infty y p_y(y) dy \\ &\geq y_0 \int_{y_0}^\infty p_y dy. \end{aligned} \tag{1.9}$$

Thus

$$P_y(y > y_0) \leq \frac{E(y)}{y_0}, \tag{1.10}$$

the simple but beautiful inequality of Bienaymé.

We now set $y = |x - a|^n$ where x is a random variable with density distribution, $p_x(x)$. Clearly, $y \geq 0$ regardless of the range of x. Consider the specific case, $n = 2$, $a = \mu$ and $y_0 = k\sigma^2$. Then we obtain

$$P_x[|x - \mu)| > k\sigma] \leq \frac{1}{k^2}, \tag{1.11}$$

the Tchebycheff inequality. Neither of these inequalities provide specific useful constraint for a given data set since the constraints are quite weak (k^{-2} should be compared with $\exp(-k^{-2})$ expected of a Gaussian distribution).

Cantelli's inequality is the one-tail version of Equation 1.11 inequality (`needs to be proved.`):

$$P(x - \mu \geq k\sigma) \leq \frac{1}{1 + k^2}. \tag{1.12}$$

Setting $k = 1$ in Equation 1.12 one obtains , $P(x \geq \mu+\sigma) \leq 1/2$. This means that the median $m < \mu + \sigma$. Interchanging the signs within the argument of P or applying Equation 1.12 to $x' = -x$ we find $P(x \leq \mu - \sigma) \leq 1/2$. This means that $m > \mu - \sigma$. Thus the range $[\mu - \sigma, \mu + \sigma]$ *must* include m, the median. In other words, the $|\mu - m| \leq \sigma$, independent of the probability density function.

`Exercise: Tchebycheff Inequality: another proof.` Start with

$$\begin{aligned} \sigma^2 &= \int_{-\infty}^{\infty} (x-\mu)^2 p_x(x) dx \\ &\geq \int_{-k\sigma}^{k\sigma} (x-\mu)^2 p_x(x) dx \end{aligned} \tag{1.13}$$

and derive Equation 1.11 relation.

`Exercise: Demonstrate Equality.` Construct examples when the inequality discussed above becomes equality.

1.4 Correlation and Covariance

1.5 Characteristic and Related Functions

The characteristic function is the Fourier transform of the probability distribution and it has two principal applications: in deriving the probability distribution of sums and differences of independent random variates (§ZZ) and in obtaining the central moments. Letting x be our random variate the characteristic function is defined without any adornments of π and as follows:

$$\phi(t) \equiv E\big[\exp(itx)\big] = \int_{-\infty}^{\infty} p(x) \exp(itx) dx. \tag{1.14}$$

The characteristic function exists provided that $p(x)$ is well behaved and bounded. Now consider the Taylor expansion around $t = 0$.

$$\begin{aligned} \phi(t) &= \int p(x)\big[1 + itx + \frac{(itx)^2}{2!} + \frac{(itx)^3}{3!} + ...\big] dx \\ &= 1 + it\langle x\rangle + \frac{(it)^2}{2!}\langle x^2\rangle + ... \end{aligned} \tag{1.15}$$

Thus the Maclaurin series (Taylor expansion evaluated at $t = 0$) of $\phi(t)$ yield the moments of $p(x)$.

However, many times it is of interest to compute central moments and for this the following subsidiary function is useful:

$$\begin{aligned} \phi_\mu(t) &\equiv \int p(x)\exp[it(x-\mu)]dx \\ &= \exp(-it\mu)\phi(t). \end{aligned} \tag{1.16}$$

One can show straightforwardly,

$$i^k\mu'_k = \left.\frac{\partial^k\phi_\mu(t)}{\partial t^k}\right|_{t=0}. \tag{1.17}$$

The cumulants κ_n are defined by the following equation:

$$\log[\phi(t)] \equiv \sum_{n=1}^{\infty}\kappa_n\frac{(it)^n}{n!}. \tag{1.18}$$

The Maclaurin series of $\log(phi(t))$ yields

$$\log[\phi(t)] = (it)\mu'_1 + \frac{1}{2}(it)^2[\mu'_2 - {\mu'_1}^2] + \frac{1}{3!}(it)^3[2{\mu'_1}^3 - 3\mu'_1\mu'_2 + \mu'_3] + ... \tag{1.19}$$

where μ'_k are the raw moments (see above, Equation XX). Re-arranging:

$$\begin{aligned} \kappa_1 &= \mu'_1 \\ \kappa_2 &= \mu'_2 - {\mu'_1}^2 \\ \kappa_3 &= 2\mu'1^3 - 3\mu'_1\mu'_2 + \mu'_3. \end{aligned} \tag{1.20}$$

The cumulants have a highly desired relation to the central moments:

$$\begin{aligned} \kappa_1 &= \mu \\ \kappa_2 &= \mu_2 \\ \kappa_3 &= \mu_3 \\ \kappa_4 &= \mu_4 - 3\mu_2^2. \end{aligned} \tag{1.21}$$

The first central moment is the mean and the second central moment is the variance. Sometimes the third model is presented as "skewness", $\gamma_1 = \mu_3/\mu_2^{3/2}$ and the fourth moment as "kurtosis", $\gamma_2 = \mu_4/\mu_2^2$.

`The k-statistics are the unbiased estimators of the cumulants. Add to discussion.`

`Gaussian Distribution.` Show that for symmetric Gaussian distribution, $p(x) \propto \exp(-x^2)$ (over the range $[-\infty, \infty]$), all cumulants save the second cumulant are zero.

1.5.1 Probability Generating Function

The probability generating function, $G(z)$, is useful for random variates (r) which take only integer or quantized values:

$$G(z) \equiv E(z^r) = \sum_r p(r) z^r \tag{1.22}$$

where the summation is over all allowed value of r. z is treated as a continuous parameter. Differentiating Equation 1.22 we obtain

$$\begin{aligned} G'(z) &= \sum_r r z^{r-1} p(r) \\ G''(z) &= \sum_r r(r-1) z^{r-2} p(r) = \sum_r r^2 z^{r-2} p(r) - G'(z), \\ G'^k(z) &= \sum_r r(r-1)...(r-k+1) z^{r-k}. \end{aligned} \tag{1.23}$$

Evaluating the functions at $z = 1$ (in many cases, more precisely, as z approaches unity from smaller values) and re-arranging we find

$$\begin{aligned} E(r) &= G'(1), \\ V(r) &= G''(1) + G'(1) - G'(1)^2, \\ E(r(r-1)...(r-k+1)) &= G'^k(1). \end{aligned} \tag{1.24}$$

Let r_1 and r_2 be independent random variables with probability generating functions, $G_1(z) = E(z^{r_1})$ and $G_2(z) = E(z^{r_2})$. Now consider $R = ar_1 + br_2$. Then

$$G_R(z) = E(z^R) = \sum_R z^{ar_1+br_2} p(r_1, r_2) = G_1(az) G_2(bz). \tag{1.25}$$

The probability generating function for the sum, $S = r_1 + r_2$ and the difference, $D = r_1 - r_2$ is

$$\begin{aligned} G_S(z) &= G_1(z) G_2(z) \\ G_D(z) &= G_1(z) G_2(1/z), \end{aligned} \tag{1.26}$$

respectively.

`Exercise: Poisson Distribution`. The Poisson probability distribution is given by $p(r) = (\lambda^r / r!) \exp(-\lambda)$ where $r = 0, 1, 2, ...$. Show that $E[r(r-1)...(r-k+1)] = \lambda^k$.

`Exercise: Infer the Distribution.` You are given that a particular distribution has the following values: $G'(1) = 1$ and all higher derivatives,

evaluated at $z = 1$, are zero. Infer the underlying distribution. (This problem can be clearly solved with knowledge of calculus. However, think what does it mean that the variance is zero. This is a pedagogical problem which demonstrates the power of thinking over purely symbolic manipulations.)

`Exercise: Galton-Watson process.` Let r_1, r_2, ... r_N be independently but identically distributed [with probability generating function, $G(z)$]. N itself is the result of a separate random process [with probability generating function, $G_N(z)$]. Consider the random variable $R = \sum_{j=1}^{N} r_j$. Show that the probability generating function for R is

$$G_R(z) = G_N[G(z)]. \tag{1.27}$$

`Exercise: Vanishing of Surnames.` By tradition the Koreans are restricted to one of about 250 surnames. Assume that in the distant past that there was equipartition between these original names. However, over time there has been a shrinking of surnames even though the population of the Korean peninsula has dramatically increased. Three surnames (Kim, Lee and Park) account for 45% of the population. Can you develop a model to explain the extinction of family lines?

1.6 Transformation of Variables

1.7 Addition and Subtraction Theorem

Consider two random variates, x and y, with probability density function, $p_{xy}(x, y)$. We wish to know the distribution function of their sum, $z = x + y$.

In the x-y plane, lines of constant z slope at 135° with an x or y intercept of z (Figure XX). We compute the cumulative function of z. I Let x be the free variable. Then y must be $z - x$. Then the cumulative function of z is

$$P_z(z) \;=\; \int_{-\infty}^{\infty} dx \int_{-\infty}^{z-x} p_{xy} dy. \tag{1.28}$$

The probability density function of z is dP_z/dz. We Taylor expand $P_z(z)$ around z and can rapidly conclude that

$$p_z(z) = \frac{dP(z)}{dz} \;=\; \int_{-\infty}^{\infty} p_{xy}(x, z - x) dx. \tag{1.29}$$

Let us now assume that x and y are independent. In this case,

$$p_z(z) = \int_{-\infty}^{\infty} p_x(x) p_y(z - x) dx = p_x(z) * p_y(z) \tag{1.30}$$

Thus p_z is the convolution of p_x and p_y.

A well known theorem in theory of Fourier transforms is that convolution in one domain is multiplication in the conjugate domain. Thus

$$\phi_z(z) = \phi_x(z)\phi_y(z). \tag{1.31}$$

Thus the characteristic function of z is simply the product of that of x and y, provided that x and y are independent.

Now let us consider $z = x - y$. In this case,

$$\begin{aligned} p_z(z) &= \int_{-\infty}^{\infty} p_x(z+y)p_y(y)dy \\ &= p_x(z) \odot p_y(z). \end{aligned} \tag{1.32}$$

Thus the p_z is the correlation of p_x and p_y. Correspondingly,

$$\phi_z(z) = \phi_x^*(z)\phi_y(z) \tag{1.33}$$

The convolution theorem for $x + y$ is well known but the correlation theorem for $x-y$ is less known. Since, for two arbitrary functions, convolution of two functions is not expected to be equal to their correlations $x - y$ and $x + y$ do not follow the same distribution. However, if $\phi_x(x)$ and $\phi_y(y)$ are both symmetric or anti-symmetric functions then $x \pm y$ follow similar distributions.

Chapter 2

Probability Distributions

2.1 Uniform Distribution

The simple uniform distribution surprisingly arises in a number of situations. The distribution of the phase of a complex random Gaussian variate is uniformly distributed over the range $[0, 2\pi]$. `Other examples.`

However, the greatest use of uniform variates arises in the generation of random numbers in computers. A number of algorithms exist that can generate a random number uniformly over an integer range, $[a, b]$. The resulting numbers form the basis of generating random numbers with a desired probability density function.

The probability density function of a deviate, distributed uniformly between a and b is $p(a \leq x \leq b) = 1/(b-a)$ and zero otherwise. The mean and variance of a uniform variate is

$$\begin{aligned} \mu &= \int_a^b xp(x)dx = (1/2)(b+a) \\ V &= \int_a^b x^2 p(x)dx - \mu^2 = (1/12)(b-a)^2. \end{aligned} \tag{2.1}$$

Letting $b = -a$ the characteristic function is

$$\phi_0(t) = \int_{-a}^{a} p(x) \exp(itx)dx = \mathrm{sinc}(ta). \tag{2.2}$$

The mean and variance can be obtained by Taylor expansion around $t = 0$ and taking the first and and second derivatives of $\phi_0(t)$ at $it = 0$. For the more general case, utilizing Equation XXX, $\phi_\mu(t) = \phi_0(t) \exp(-it\mu)$ where $\mu = (1/2)(b+a)$.

2.2 Binomial Distribution

Consider an event whose outcome is binary. Let p the probability for success in which case $q = 1 - p$ is the probability for failure. This is called a "Bernoulli" trial. Now consider n Bernoulli trials. The probability for r successes in n trails leads to the "Binomial" distribution:

$$B(r; n, p) = \frac{n!}{(n-r)!r!} p^r q^{n-r} \qquad r = 0, 1, ..., n \tag{2.3}$$

The characteristic function is

$$\begin{aligned} \phi_n(t) &\equiv E(\exp(it)) = \sum_{r=0}^{n} B(r; n, p) \exp(irt) \\ &= [p \exp(it) + q]^n. \end{aligned} \tag{2.4}$$

Here, we have used the well known binomial expansion theorem, $(x+y)^n = \sum_{r=0}^{n} [n!/((n-r)!r!)] x^r y^{n-r}$. Another way to obtain $\phi_n(t)$ is by assigning $y = 1$ to success and zero to failure. With this assignments $y = r$. For a single Bernoulli trial, the characteristic function is $\phi_1(t) = \langle y \rangle = p \exp(it) + q$. Since successive trials are independent we have $\phi_n(t) = \Pi_{j=1}^{n} \phi_1(t) = \phi_1^n(t)$.

The first k moments can be obtained by Taylor expansion of $\phi(t)$ around $t = 0$. evaluating $d^k \phi_n(t)/d(it)^k$ at $it = 0$ (§XX). From this we find

$$\begin{aligned} E(r) &= np \\ V(r) &= npq. \end{aligned} \tag{2.5}$$

The binomial distribution is symmetrical when $p = q$. The skewnessis proportional $p - q$.

The binomial distribution becomes a Gaussian distribution $N(\mu, \sigma^2)$ [§3] with $\mu = np$ and $\sigma^2 = npq$ as n becomes a large number. If, on the other hand, as n increases p decreases so that $np = \lambda$ then the binomial distribution asymptotes to a Poisson distribution, $P(r, \lambda)$ [§2.4].

The multinomial distribution is the generalization in which there is more than two outcomes, say m outcomes. Let $\underline{\mathbf{r}}$

$$\begin{aligned} \underline{\mathbf{r}} &= (r_1, r_2, ..., r_m) \\ \underline{\mathbf{p}} &= (p_1, p_2, ..., p_m) \end{aligned} \tag{2.6}$$

be the number of successes for outcome 1 through m and let the corresponding probabilities be $p_1, p_2, ..., p_m$. Clearly, $\sum_{j=1}^{m} p_j = 1$ and the total number

of tries be $n = \sum_{j=1}^{m} r_j$. The probability distribution and the first two moments are given by

$$\begin{aligned} M(\underline{\mathbf{r}}; n, \mathbf{p}|n) &= \frac{n!}{r_1! r_2! ... r_k!} p_1^{r_1} p_2^{r_2} ... p_k^{r_k} \\ E(r_j) &= np_j \\ V(r_j) &= np_j(1-p_j) \\ \mathrm{cov}(r_i, r_j) &\equiv E(r_i, r_j) - E(r_i)E(r_j) = -np_i p_j \\ \rho_{ij} &\equiv \frac{\mathrm{cov}(r_i, r_j)}{\sqrt{V(r_i)V(r_j)}} \\ &= -\sqrt{\frac{p_i p_j}{q_i q_j}}. \end{aligned} \tag{2.7}$$

The covariance is negative since, given that n is fixed, any increase in one bin has to come at the expense of other bins.

As noted in the above equation, the multinomial distribution has been computed assuming n is known and fixed. However, let us consider the more likely case when n itself is a result of a Poisson process, i.e.

$$p(n) = \frac{\lambda^n}{n!} \exp(-\lambda). \tag{2.8}$$

where $\langle n \rangle = \lambda$. Noting $n = r_1 + r_2 + ...r_k$ and $p_1 + p_2 + ...p_k = 1$. Then

$$\begin{aligned} M(\underline{\mathbf{r}}; n, \mathbf{p}) &= M(\underline{\mathbf{r}}; n, \mathbf{p}|n)p(n) \\ &= \frac{\lambda^n}{n!} \exp(-\lambda) \frac{n!}{r_1! r_2! ... r_k!} p_1^{r_1} p_2^{r_2} ... p_k^{r_k} \\ &= \frac{(p_1\lambda)^{r_1}}{r_1!} \frac{(p_2\lambda)^{r_2}}{r_2!} ... \frac{(p_k\lambda)^{r_k}}{r_k!} \exp(-\lambda) \\ &= \frac{(p_1\lambda)^{r_1}}{r_1!} \exp(-p_1\lambda) \frac{(p_2\lambda)^{r_2}}{r_2!} \exp(-p_2\lambda) ... \frac{(p_k\lambda)^{r_k}}{r_k!} \exp(-p_k\lambda). \end{aligned}$$

Thus, the multinomial distribution in which the total number of events follows a Poisson distribution is exactly the same as independent Poisson distributions with $\lambda_j = p_j\lambda$. In particular, for such an unconstrained distribution the covariance between different bins is zero.

`Need some examples other than histograming data.`

2.3 Exponential Distribution

The exponential distribution lies at the basis of reliability of manufacturing. In particular, the time for a component or system to fail is described by

the exponential distribution. In physics, one may regard the decay of a radioactive nucleus as an example of a component falling apart.

Consider a component (nucleus). Let the probability that the component will fail (undergo fission) in the interval $[t, t+\delta t]$ be $\delta p = \lambda \delta t$. Here the interval is assumed to be sufficiently small $\delta p \ll 1$. For a component to decay in the time interval $[\tau, \tau+\delta t]$ the unit must have survived for all the previous time *and* decayed within $[\tau, \tau+\delta t]$. Thus

$$\begin{aligned} p(\tau)\delta t &= (1-\lambda\delta t)^{\tau/\delta t - 1} \times \lambda \delta t \\ &\approx \lambda \delta t \exp(-\lambda \tau). \end{aligned} \tag{2.9}$$

where we have used the well known approximation $(1-\epsilon)^n \sim \exp(-n\epsilon)$ when $\epsilon \ll 1$ but $n\epsilon$ is a finite number. Thus

$$p(\tau) = \lambda \exp(-\lambda \tau). \tag{2.10}$$

The mean time to decay is

$$\langle \tau \rangle = \int_0^\infty \tau p(\tau) d\tau = \lambda^{-1}. \tag{2.11}$$

and the cumulative function is

$$P(\tau) = 1 - \exp(-\tau/\langle \tau \rangle) \tag{2.12}$$

The fraction of components that survive to time τ is $1-P(\tau)$ and is plotted in Figure ZZ. The survival fractions are 37%, 13%, 5% and 2% by $[1, 2, 3, 4] \times \langle \tau \rangle$.

In the simple exponential model the probability to decay is independent of time. The Weibull distribution allows for increased failure at early times (e.g. infantile death) or later times (wearing out of parts). Consider a cumulative function, $P(\tau) = 1 - \exp(-g(\tau))$. The corresponding probability density function is $p(\tau) = g'(\tau)\exp(-g(\tau))$. The *fractional* rate of units failing at time τ, $\lambda(\tau)$, is given by the rate of failure normalized to the number of units that have survived to time τ. Thus

$$\lambda(t) = \frac{p(\tau)}{1-P(\tau)} = g'(t) \tag{2.13}$$

A standard operating procedure is to invoke power laws when some generality is called for. We set

$$g'(\tau) = \frac{\beta}{\alpha}\Big(\frac{\tau}{\alpha}\Big)^{\beta-1} \tag{2.14}$$

and find

$$p(\tau) = \frac{\beta}{\alpha}\Big(\frac{\tau}{\alpha}\Big)^{\beta-1}\exp\Big[-\Big(\frac{t}{\alpha}\Big)^{\beta}\Big]. \tag{2.15}$$

With this assumption one has the Weibull probability density function where β (the "shape") parameter sets the timescale and α ("scale") provides the normalization.

$\beta = 1$ corresponds to the simple exponential model. $0 < \beta < 1$ correspond to premature failures whereas $\beta > 1$ account for increased failure as components age.

2.4 Poisson Distribution

The Poisson distribution is central to optical and X-ray astronomy or any low-light level experiment. It is a distribution worth knowing since many routine things in life follow Poisson distribution: the number of cars which pass a given point in a road (assuming the traffic rate is constant on average), the number of mutations after being irradiated and so on.

Assume a constant intensity light source (such as a well stabilized single-mode laser) is incident on a photoelectric detector. Our goal is to estimate the variation in the number of photoelectrons detected over say a time T. As in the previous section (§2.3) we divide the interval $[0, T]$ into many small intervals of duration, δt. The intervals are chosen to be small enough that there is negligible chance of detecting more than one photo-electron. Let the probability for photoelectron emission in each such interval be $\lambda\delta t$. We now have a binomial problem with $n = T/\delta t$ and $p = r\delta t$ (see §2.2). Thus while n is a very large number, $pn = \lambda T$ is finite. Thus the probability distribution of the number of photoelectrons, $P(k)$, is

$$\begin{aligned} P(k) &= \frac{n!}{(n-k)!k!}p^r(1-p)^{n-k} \\ &= \frac{1}{r!}[np][(n-1)p]...[(n-k+1)p](1-p)^{n-k} \\ &\approx \frac{1}{k!}(rT)^k\exp(-rT) \end{aligned} \tag{2.16}$$

where two approximations have been made, $n \gg 1$ and $(1-\epsilon)^n \approx \exp(-n\epsilon)$ as n tends to infinity but $n\epsilon$ is a finite number.

Thus the Poisson distribution can be re-expressed as

$$P(k) = \frac{\lambda^k}{k!}\exp(-\lambda). \tag{2.17}$$

where $\lambda = \langle k \rangle$. The mean, variance, mode and median of the Poisson distribution is

$$\begin{aligned} \langle k \rangle &= rT \\ V(k) &= \langle k \rangle \\ \text{med}(k) &= \lambda + 1/3 - 0.02/\lambda \\ \text{mode}(k) &= \text{floor}(\lambda), \quad [\lambda - 1, \lambda] \end{aligned} \tag{2.18}$$

The following relation is particularly useful in the Chapter XX (Semiclassical Theory of Detection):

$$\langle k(k-1)...(k-r+1) \rangle = \lambda^r. \tag{2.19}$$

The characteristic function of this distribution is

$$\phi(t) = \sum_{k=0}^{\infty} P(k) \exp(itk) = \exp[\lambda(e^{it} - 1)]. \tag{2.20}$$

Consider two independently distributed Poisson variates, k_1 and k_2, with mean values λ_1 and λ_2. Let $K = k_1 + k_2$. The characteristic function of K is

$$\begin{aligned} \phi_K(t) &= \exp[\lambda_1(e^{it} - 1)] \exp[\lambda_2(e^{it} - 1)] \\ &= \exp[\Lambda(e^{it} - 1)]. \end{aligned} \tag{2.21}$$

Thus the sum of two independent Poisson variates is also a Poisson variate with $\Lambda = \lambda_1 + \lambda_2$. Clearly this result can be readily generalized to more than two variates. In contrast, the difference of two Poisson variates follows the *Skellam* distribution.

There are several points worth remembering. First, be aware of scaling Poisson data. Specifically, consider the following example. Charge Coupled Detectors (CCDs) are the most widely detectors these days. They are found in your digital camera and the choice detectors for optical telescopes. Each photon (subject to quantum efficiency) produces a photoelectron. Modern CCDs are very well engineered and the noise associated with "reading" the detected photoelectrons is very small. Ideally one would like that one input photoelectron produces one unit of output. However, given more bits for the output costs more money. A compromise is made in which each output unit corresponds to say g input photoelectrons. Letting y be the output and n_p be the number of electrons we have $y = n_p/g$. A typical mistake is to claim that y follows Poisson statistics. This is wrong. We know that variance in n_p, $V(n_p) = n_p$ and thus $V(y) = n_p/g^2 = y/g$. Thus the the rms of y is $\sqrt{y/g}$ and not $\sqrt{y}$.

You can turn around and exploit the fact that the ratio of the measured mean to the measured variance of a Poisson variable is unity to determine the scale factor. APPLICATION ...

Next, as explained in §2.3 the difference in arrival time of successive photons follows an exponential distribution.

Third, at low light levels ($\lambda \ll 1$) the number of integration intervals ($[0, T]$) with zero photo-electrons is $P(0) = \exp(-\lambda)$. Thus the number of intervals with one or more photo-electrons is $P(\geq 1) = 1 - \exp(-\lambda) = \lambda + \lambda^2/2! + \lambda^3/3! + ...$, corresponding to one, two, three and so on photoelectrons. However, for faint levels ($\lambda \ll 1$), $P(\geq 1) = 1 - P(0) \approx \lambda$.

The Poisson distribution is widely used in surveys. Let us say that you have a survey a wide area of the sky looking for an award-winning source (e.g. a black hole-milisecond pulsar binary). Unfortunately no source is seen. You would then like to say something quantative of the area density (the number of sources per square degree) of such sources. Clearly, the lower limit to the area density is zero. To compute an upper limit we need to specificy a confidence level, $\mathcal{P}$. The upper limit is computed by demanding that no source be detected even though true source count was λ no more than $1 - \mathcal{P}$ fraction of the cases, i.e.

$$\begin{aligned} P(0) &= 1 - \mathcal{P} \\ \lambda &< -\log(1 - \mathcal{P}) \end{aligned} \tag{2.22}$$

Thus a true source population with $\lambda = 4.6$ will result in a non-detection about 1% fraction of the cases.

`Literature:`

The median of Poisson distribution is from Adrella & Jodra (2005, Metrika); see Wikipedia also. The discussion of the Weibull distribution is from `www.mathpages.com`.

`Applications`

Problem Set: Probability Distributions

`Exercise PD-2(P): Sum and Difference of Binomial variates.` Let r_1 and r_2 arise from $B(r_1; n, p)$ and $B(r_2; m, p)$). Show that $R = r_1 + r_2$ arises from $B(R; n + m, p)$. Next, show that $\mathcal{R} = r_1 - r_2$ is not a binomial distribution. Compute $E(\mathcal{R})$ and $V(\mathcal{R})$.

`Exercise PD-3(P): Poisson Conditional Distribution.` Let k_1 and k_2 be independent Poisson variates with mean values of λ_1 and λ_2. Let $K = k_1 + k_2$. Show that

$$p(k_1|K) = B(k_1; K, \frac{\lambda_1}{\lambda_1 + \lambda_2}). \tag{2.23}$$

`Exercise PD-4(P): Geometric Distribution.`[1] Let r be the number of trials before a failure takes place. The probability distribution of r is

$$G(r) = p^r q. \qquad r = 0, 1, 2, ... \tag{2.24}$$

Show that the probability generating function is

$$G(z) = \langle z^r \rangle = \frac{p}{1 - qz}. \tag{2.25}$$

From this show that $\langle r \rangle = p/q$ and $V(r) = pq^{-2}$.

[1]There are two different definitions of geometric distribution. One definition is the number of trials before a failure takes place. The steps can range from 0 (the first event is itself a failure) and upwards. The other definition is the number of trials at which the first failure takes place. In this case, the steps can range from 1 and upwards. Here, we use the first definition.

Solutions to Problem Set of Probability Distributions

Chapter 3

Gaussian Distribution

Some history.

The Gaussian distribution for a single variate is given by

$$N(\mu, \sigma^2) = \frac{1}{\sqrt{2\pi}} \exp\Big[- \frac{(x-\mu)^2}{2\sigma^2} \Big]. \tag{3.1}$$

The characteristic function is

$$\begin{aligned} \phi_N(t) &= \langle \exp(it) \rangle \\ &= \exp[-t^2\sigma^2] \end{aligned} \tag{3.2}$$

3.0.1 χ^2 Distribution

3.0.2 The Central Limit Theorem

Chapter 4

Generation of Random Numbers

Our goal in this section is to generate random numbers with a specified distribution function, say $p_y(y)$. All computer generated number schemes assume the availability of high quality random numbers with a uniform distribution[1].

A distinct advantage of a random number which is uniform over (say) the interval $[0, 1]$ is that the cumulative distribution also spans this interval and by definition the range $[0, 1]$ is equally probable. Deviates whose inverse cumulative function is readily calculable can be easily obtained from uniform random deviates (this is called as the "inverse" approach).

There are three methods by which variates with a desired distribution can be generated: transformation, inverse of the cumulative function and rejection methods. We discuss each method with examples.

Let x be a random number uniformly distributed between $[0, 1]$. Then $p_x(x) = 1$ for $0 \leq x \leq 1$ and zero otherwise. The simplest approach is to find a transformation, $y(x)$ so that the resulting distribution of the y's follows the desired distribution. From Chapter XX, we know that

$$p_y(y) = p_x(x)\left|\frac{dx}{dy}\right|. \tag{4.1}$$

[1]This is the modern approach. In the past, random numbers were obtained from real experiments. A fair mechanical or electrcal roulette wheel can be expected to generate uniform random variates. A radioactive source can serve as a proxy for exponential distributions. Finally, teams of unemployed astronomers armed with a fair coin each can be usefully employed to generate binomial distributions. One of the best known examples of "genuine" random numbers is a book published by RAND corporation in 1948 titled "A Million Random Digits". This books si regarded as one of RAND's greatest publication.

4.0.1 Exponential and Related Deviates

The transformation $y = -\log(x)$ is an elegant way to generate exponential distribution. The proof is quite simple.

$$p_y(y) = p_x \Big| \frac{dx}{dy} \Big| = x = \exp(-y). \tag{4.2}$$

χ^2_2 also follows the exponential distribution. Given the χ^2 addition theorem (REF) we find

$$u = -\log(x_1) - \log(x_2) - ... \log(x_n) \tag{4.3}$$

follows χ^2_{2n} distribution, provided of course x_j are independently distributed.

The exponential distribution is also χ^2_2

4.0.2 Poisson and Related Deviates

4.0.3 Gaussian Deviates

Let x_1 and x_2 be independent uniform variates over the range $[0, 1]$. Then

$$\begin{aligned} p(x_1, x_2) &= p(x_1)p(x_2) = 1 \qquad 0 \le x_j \le 1 \\ &= 0 \qquad\qquad\qquad\quad \text{otherwise} \end{aligned} \tag{4.4}$$

The Box-Muller transformation

$$\begin{aligned} y_1 &= \sqrt{-2\log(x_1)}\cos(2\pi x_2) \\ y_2 &= \sqrt{-2\log(x_1)}\sin(2\pi x_2) \end{aligned} \tag{4.5}$$

results in y_1 and y_2 being independent Gaussian variates with zero mean and unit variance.

The proof is as follows. From Equation 4.5 we find

$$\begin{aligned} x_1 &= \exp\big[-\frac{1}{2}(y_1^2 + y_2^2)\big] \\ x_2 &= \frac{1}{2\pi}\arctan\Big(\frac{y_2}{y_1}\Big). \end{aligned} \tag{4.6}$$

The probability distribution of y_1 and y_2 is given by (REF)

$$\begin{aligned} p(y_1, y_2) &= p(x_1, x_2)\frac{\partial(x_1, x_2)}{\partial(y_1, y_2)} \\ &= \begin{Vmatrix} \frac{\partial x_1}{\partial y_1} & \frac{\partial x_1}{\partial y_2} \\ \frac{\partial x_2}{\partial y_1} & \frac{\partial x_2}{\partial y_2} \end{Vmatrix} \end{aligned}$$

$$\begin{aligned} &= \frac{1}{2\pi}\exp\Big[-\frac{1}{2}y_1^2-\frac{1}{2}y_2^2\Big] \\ &= \frac{1}{\sqrt{2\pi}}\exp\Big[-\frac{1}{2}y_1^2\Big]\frac{1}{\sqrt{2\pi}}\exp\Big[-\frac{1}{2}y_2^2\Big]. \end{aligned} \tag{4.7}$$

Thus y_1 and y_2 are independent Gaussian variates with zero mean and unit variance. This transformation is ideally suited to generating circular complex Gaussian numbers (see Chapter XX).

A variation known as the Marsaglia polar method speeds up the Box-Muller transformation by eliminating trignometric calculations. Consider two independent random variates, u, v, over the range $[-1, 1]$. Accept the pair only if the pair lies within the unit circle, $R^2 = u^2 + v^2 < 1$ (see Figure 4.1). Thus the joint distribution of u and v is

$$p(u,v)dudv = \frac{1}{\pi}dudv. \tag{4.8}$$

Noting that

$$\begin{aligned} u &= R\cos(\theta) \\ v &= R\sin(\theta) \end{aligned} \tag{4.9}$$

we find

$$p(R,\theta) = p(u,v)\left\| \begin{array}{cc} \frac{\partial x_1}{\partial R} & \frac{\partial x_1}{\partial \theta} \\ \frac{\partial x_2}{\partial R} & \frac{\partial x_2}{\partial \theta} \end{array} \right\| \tag{4.10}$$

Thus

$$\begin{aligned} P(R,\theta)dRd\theta &= \frac{d\theta}{\pi}dR \\ &= \frac{d\theta}{2\pi}d\rho \end{aligned} \tag{4.11}$$

where $\rho = R^2$. Thus, θ and ρ follow a uniform distribution over the range of $[0, 2\pi]$ and $[0, 1]$. Noting that $\cos(\theta) = u/R$ and $\sin(\theta) = v/R$ and the uniform distribution of ρ and θ we re-express y_1 and y_2 (Equation 4.5) and find

$$\begin{aligned} y_1 = \sqrt{-2\log(x_1)}\cos(2\pi x_2) &\rightarrow u\sqrt{-2\log(\rho)/\rho} \\ y_2 = \sqrt{-2\log(x_1)}\sin(2\pi x_2) &\rightarrow v\sqrt{-2\log(\rho)/\rho}. \end{aligned} \tag{4.12}$$

The method is inefficient by $4/\pi = 1.27$ but is compensated by not undertaking computationally expensive trignometric computations.

The recent method is the Ziggurat method.

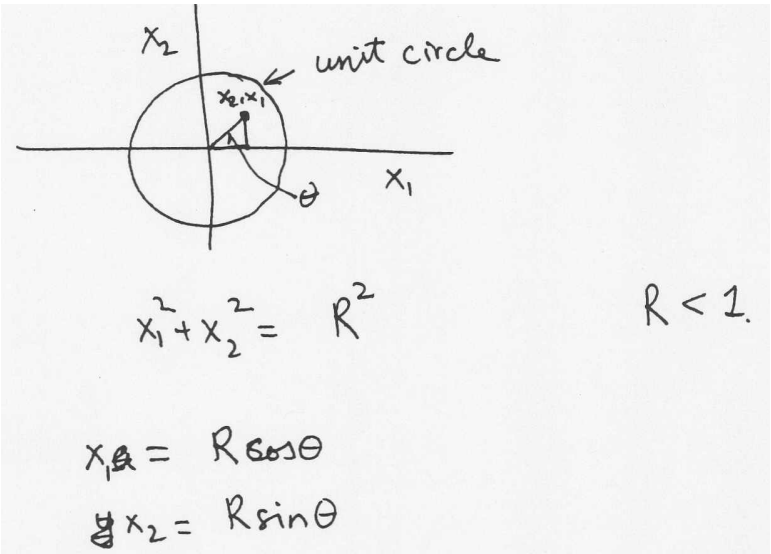

Figure 4.1: Marsaglia polar form.

4.0.4 Other Deviates

The lifetime of a product whose probability to fail is expected to be constant with time follows an exponential law. The Weibull distribution allows for the probability to fail to change with time:

$$W(z) = \alpha\beta z^{\beta-1}\exp(-\alpha z^{\beta}) \qquad z > 0. \tag{4.13}$$

and zero otherwise. The mean and variance are respectively

$$\begin{aligned} \mu &= \alpha^{-1/\beta}\Gamma(1+1/\beta) \\ \sigma^2 &= \alpha^{-2/\beta}\big[\Gamma(1+2/\beta) - \Gamma^2(1+1\beta)\big]. \end{aligned} \tag{4.14}$$

If $\beta < 1$ then the failure rate decreases with time (as with people, "infantile" mortality) whereas $\beta > 1$ corresponds to an increase with time (as with your car, wearing out of parts). A second application of the Weibull probability function is that it can mimic an exponential ($\beta = 1$) to an appoximate Gaussian ($\beta \sim 3.4$).

The cumulative function of this distribution is

$$F(y) = \int_0^y W(z)dz = 1 - \exp(-\alpha y^{\beta}) \tag{4.15}$$

We know that $x = F(y)$ is uniformly distributed between 0 and 1. Thus,

$$y = \Big[-\frac{1}{\alpha}\log(1-x)\Big]^{1/\beta} \tag{4.16}$$

follows the Weibull distribution.

Other examples include $y = \tan(x)$ and $y = \sin(x)$.

NOTES. To read up.

`http://www.mathworks.com/company/newsletters/news_notes/clevescorner/spring01_cleve.`

4.1 Pitfalls

It is not uncommon to find someone using sums of uniform random variates as a proxy for Gaussian distribution (appealing to the central limit theorem; §XX) or the sum of binomial variates as a proxy for Poisson variate. These two (and any other such approach) is risky because the main purpose of simulation is not to derive the mean value (this usually can be calculated quite easily) but to determine the probability distribution of the outcome of the experiment. It is the distribution of extreme values that should be motivating the simulation analysis. The sum of a few tens or even hundreds of

`Exercise: Random numbers of questionable significance.` Let r be random number distributed between $[-1, 1]$. Consider the sum of n such numbers, $R = \sum_{j=1}^{n} r_j$. For definitiness, set $n = 100$. What is the expected and mean and variance of R? Compare the distribution of R with that of Gaussian distribution (whose mean and variance match the expected value). Howe deep should the simulation (i.e. the total number of R's you generate) be to adequately characterize the difference in R and an ideal Gaussian distribution?

Chapter 5

Order Statistics

Say you have undertaken a Fourier transform of an X-ray time series of a neutron star and you are looking for periodic signal but with an unknown period. [The profile is assumed, for simplicity, to be a pure sinusoid]. The usual approach is to Fourier transform the X-ray time series and then search the power spectrum. The frequency bin corresponding to that of the pulsar period would be expected to show a much larger value relative to other bins. In the same spirit, let us say that you are looking for a short burst from a neutron star or a spiky emission from a flare star. We will assume that you you have done the usual processing of data (including rejecting interfering signals) and that the resulting time series follows a single distribution, $p(h)$. Again the approach is to look for a spike in your time series. Likewise, the search for a Higgs boson, consists of looking for a spike in an energy spectrum of events seen at the LHC.

In each of the above cases, on inspection, you may find a maximum that appears to be quite large and enticing. Should you get excited? You instinctively realize that the more bins you have searched the larger will be the maximum value that you will find. Thus, your excitement has to be contained until you understand the probability distribution of the maxima.

One can generalize and think about the second highest bin, the third highest bin and so on. We note that the median is the $(n+1)/2$ highest bin[1]. Continuing, the minimum is the one with smallest value of input series. The statistics of extrema (maximum or minimum), the median and the range (maximum$-$minimum) is called as "order statistics"[2].

[1]The median is well defined when n is odd. However, this is nit-picking since the main application of order statistics is when n is large.

[2]In books of statistics the minimum is the first order statistic and the maximum is the nth maximum. However, most of our applications involve maxima and so we focus on the nth statistic and .

5.1 Extrema

Let us denote the measurements by h_i where $i = 1, ..., n$ and $\mathcal{H} = \max(h)$. We assume that h_i are independent of each other and drawn from a probability density distribution of $\rho(\mathcal{H})$. Say that h_1 happens to be the maximum value. Then all other values must be less than h_1. The joint probability density function then yields:

$$\rho(\mathcal{H} = h_1) = p(\mathcal{H})\Big[1 - P(\mathcal{H})\Big]^{n-1} \tag{5.1}$$

where $P(h) = \int_{-\infty}^{h} p(h)dh$ is the cumulative probability function. However, the maximum could be any one of h_i and so we find

$$\rho(\mathcal{H}) = np(\mathcal{H})\Big[1 - P(\mathcal{H})\Big]^{n-1}. \tag{5.2}$$

The cumulative distribution, $\varrho(\mathcal{H})$, is obtained by integrating $\rho(\mathcal{H})$ from $-\infty$ to $\mathcal{H}$. This yields

$$\varrho(\mathcal{H}) = \Big[1 - P(\mathcal{H})\Big]^{n}. \tag{5.3}$$

Let us now look at the probability distribution of $\mathcal{H}$ when $p(h)$ is a Gaussian distribution, $N(\mu, \sigma^2)$ and for $\mu = 0$ and $\sigma^2 = 1$. Then,

$$p(h) = \frac{1}{\sqrt{2\pi}} \exp(-h^2/2), \qquad P(h) = \frac{1}{2}\Big[1 + \mathrm{erf}(h/\sqrt{2})\Big] \tag{5.4}$$

where

$$\mathrm{erf}(t) = \frac{2}{\sqrt{\pi}} \int_0^t \exp(-t^2)dt, \qquad \mathrm{erfc(t)} = 1 - \mathrm{erf}(t). \tag{5.5}$$

Substituting Equation 5.4 into Equation 5.2 and 5.3 and using the well known approximation $(1 - \epsilon)^n \approx \exp(-n\epsilon)$ when ϵ is vanishingly small but $n\epsilon$ is finite we find

$$\begin{aligned} \rho(\mathcal{H}) &\approx np(\mathcal{H}) \exp\Big[-\frac{n-1}{2}\mathrm{erfc}(\mathcal{H}/\sqrt{2})\Big] \\ \varrho(\mathcal{H}) &\approx \exp\Big[-\frac{n}{2}\mathrm{erfc}(\mathcal{H}/\sqrt{2})\Big]. \end{aligned} \tag{5.6}$$

As can be seen from Figure 5.1 the probability density function of $\mathcal{H}$ is concentrated – a result of multiplying a Gaussian tail $\propto \exp(-h^2/2)$ with a step function $\exp[-(n-1)/2\,\mathrm{erf}(\mathcal{H}/\sqrt{2})]$.

For experimental planning it helps to have a typical value for the maximum (for a given n). A simple way to determine a typical maximum value

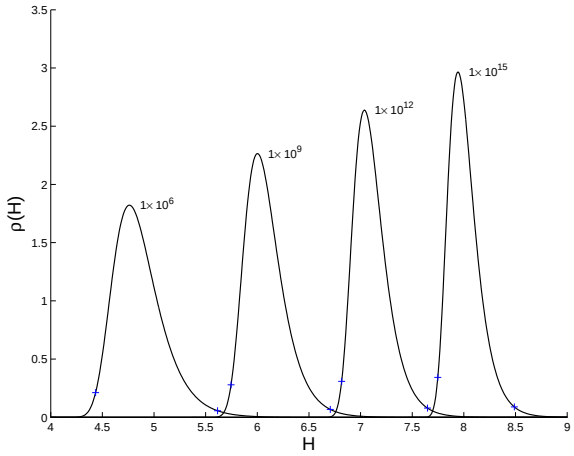

Figure 5.1: Plot of $\rho(H)$, the probability density function for the maximum, $\rho(H)$ for sample size of $n = [10^6, 10^9, 10^{12}, 10^{15}]$ and assuming that the input series follows a Gaussian density function with zero mean, namely, $p(h) = N(0, 1)$. The 1% and 99% confidence points are marked by a "+" and the intervals are [4.4, 5.6], [5.7, 6.7], [6.8 7.6] and [7.7, 8.5]; Notice the width is approximately one σ, as one would expect for any one single quantity.

(for a given n) is given by the mode of $\rho(\mathcal{H})$ since for a single experiment or run the mode is the most likely value. The peak of $\rho(\mathcal{H})$ yields the mode, H_{mode} and can be obtained the usual way. We find

$$H_{\text{mode}} = \sqrt{2 \log \Big(\frac{n-1}{H_{\text{mode}} \sqrt{2\pi}} \Big)}. \tag{5.7}$$

Note that H_{mode} increases very slowly with n being $\propto \sqrt{\ln(n)}$. Next, the above equation is not a closed form solution for H_{mode}. However, since H_{mode} on the right hand side appears as $\sqrt{\ln H_{\text{mode}}}$ we can set $H_{\text{mode}} = 6$ (typical for $n = 10^9$) and find approximate value for H_{mode}. With this simplification for $n = [10^6, 10^9, 10^{12}, 10^{15}]$ we find $H_{\text{mode}} = [4.7, 6.0, 7.1, 8.0]$ which can be compared to the exact value of $H_{\text{mode}} = [4.76, 6.01, 7.04, 7.95]$.

In order to interpret a particular experiment the cumulative distribution (Equation 5.6) is useful. We would like to know the chance (probability) of obtaining a value as high as the measured value, $\mathcal{H} = H$. If this probability is reasonable (say between 0.01 and 0.99) then you have no reason to suspect that you have a detection.

It is common practice that many searches are designed to identify singular peaks e.g. giant spikes from pulsars, strong bursts from other galaxies and so on. However, sometimes a minimum is also very interesting. Possibilities

range from occultation (e.g. by a rock in the solar system, a bird in the dome) to an extreme scattering event caused by structures in the interstellar medium or a defocussing lens in the interstellar medium. The above discussion also applies to the statistics of the minimum.

The statistics of extrema are very sensitive to imprecision in our knowledge of σ or the measurement itself or more precisely the signal-to-noise ratio (SNR) of the bin in question, $\mathcal{H}/\sigma$. A small imprecision, say 10%, will change $\mathcal{H}$ from say 7.5 to 8.5. (for an assumed $n = 10^{12}$). Given the narrowness of $\rho(\mathcal{H})$ we can say that such a shift has a dramatic impact on the inference. The cumulative probability goes from an acceptable value of 97% to 99.9921%. Likewise a difference of a single σ (owing say due to some additive error) would also have a dramatic effect. The way to recognize the presence of such biased scale factor or additive constants is to look at the distribution of the top m maxima.

The arguments which led us to derive Equation 5.2 lead us to the following probability density function for the mth highest point in sample of n elements:

$$\rho(\mathcal{H}_m) = np(\mathcal{H}_m) \times {}^{n-1}C_{m-1} P(\mathcal{H}_m)^{n-m} [1 - P(\mathcal{H}_m)]^{m-1} \qquad (5.8)$$

where ${}^nC_r = (n)!/(r!(n-r)!)$. In Figure 5.2 I plot the maxima and higher order maxima for zero mean, unit variance Gaussian variates. Notice the decreasing variance of higher order maxima. Roughly one expects the variance to decrease as $m^{-1/2}$ (with the median enjoying the smallest variance). The decreasing width of the variance means that spikes due to noise occur at rather specific values (and thereby the pattern of the top m maxima provide a distinction between noise and signal).

5.2 The Median

Perhaps the most important percentile is the median (IntREF). While the mean minimizes the sum of the square of the differences between the mean and the sample values (residuals) the median minimizes the sum of the absolute value. As a result the median is naturally more robust against outliers.

Consider a sequence of measurements, h_i as in the previous section (§5.1). Let us re-arrange this sequence to produce an ordered series, g_i where $g_i < g_j$ for $i < j$. Let $r_i = i/n$. Clearly, $0 < r_i \leq 1$. The next measurement, h_{n+1} can be below the current minimum or above the current maximum or lie between a sequential paper, say, g_i and g_{i+1} or a total of $n+1$ possibilities. Thus the probability for any of these possibilities is $1/(n+1)$. This is a useful and

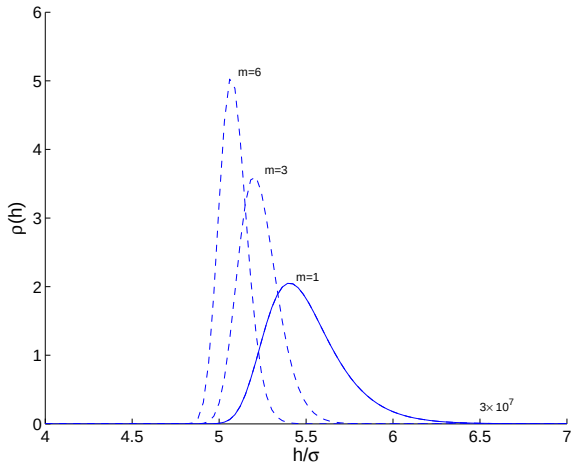

Figure 5.2: The distribution of mth maxima for a set ($n = 3 \times 10^7$) of zero mean, unit variance Gaussian variates. The m values are marked next to each of the curves. The $m = 1$ curve corresponds to $\rho(\mathcal{H})$.

general result to keep in mind. We say general because the result applies to any valid probability distribution.

Next, we compute the confidence range for the median. For this calculation it is helpful if n is even. In this case the inferred median is bracketed by g_q and g_{q+1} where $q = n/2$. The probability that the true median lies in this range is the probability of having $n/2$ points below and above the population median. This probability that such is the case is $P = {}^nC_q 2^{-n}$. Thus, for instance, if $n = 10$ the probability that the true median is in the range $[g_5, g_6]$ is 0.246. The confidence level (P) can be increased by increasing the range. Thus for instance the probability for the range $[q-1, q+2]$ is larger and given by

$$P = [{}^nC_{q-1} + {}^nC_q + {}^nC_{q+1}]2^{-n}. \tag{5.9}$$

Note that the confidence level derived thus applies to any $p(h)$.

The density of r on the axis $[0, 1]$ is uniform. With increasing n the probability density function of r becomes continuous. Thus $p(r) = 1$ and $P(r) = r$. A given percentile, say t, is the point closest to $m = tn$. The probability density that there are m points above interval $[t, d + dt]$ and $n - m - 1$ below this interval is given by (cf. Equation 5.8) is

$$\rho(n; t) = {}^nC_{m-1}\, t^{n-m}[1-t]^{m-1}. \tag{5.10}$$

In the asymptotic limit the distribution becomes normal, $\rho(n; t) \propto N(t, t(1-t)/n)$. Thus, for the median, $t = 1/2$ and we find $\rho(n; 1/2) \approx N(1/2, 1/(4n))$.

For the specific case of noise which follows Gaussian distribution, $N(\mu, \sigma^2)$, the variance of the median estimator can computed exactly:

$$V(H_{\rm med}) = \frac{4n}{\pi(2n+1)}\sigma^2 \tag{5.11}$$

where n is the sample size. In contrast, the variance of the arithmetic mean is $\sigma^2/(n-1)$ (Internal REF). Thus the median, for Gaussian distribution, is $\sqrt{\pi/2}$ worse estimator than the arithmetic mean. The median, on the other hand, is *robust* against the presence of strong fluctuations due to say Radio Frequency Interference (RFI) or cosmic rays hitting the CCD. These events with high amplitude are rare but owing to their large magnitude they will bias the arithmetic mean and the variance as well. The median estimator is essentially immune to such rare and extreme interfering events.

Incidentally Equation 5.11 should not lull the reader into thinking that the median is always a worse estimator relative to the mean. For example, the variance of the median is the same as that of the mean for exponential distribution and uniform distribution (PS-1, PS-6) and the median is a better of estimator of μ for a Laplace distribution (PS-1, PS-5). In Chapter(IntREF) we learnt that the Cauchy distribution is so pathological as to not have any central moments. In particular, for this distribution neither the mean nor variance are defined. `OS-4(P)` nicely demonstrates the median is well defined for this distribution.

The median is also useful in analysis of data, not merely in data reduction. Indeed median analysis is essential when two different underlying distributions contribute to the measurements (as discussed above) or when the underlying distribution is not Gaussian. It is safe to assume that any astronomical class of objects is *not* described by a single Gaussian. Examples include supernovae of all sorts (even if classified into Ia and II) and luminosity functions of stars of a given spectral type. The reason is simple. In the first case, it is most likely that there are multiple channels of producing type Ia supernovae (indeed the systematics of the difference between supernovae produced in elliptical and star-forming galaxies) and in the second case there may be external perturbations (e.g. stars classified as A type from low resolution spectroscopy may have faint companions and thus the luminosity will be biased towards higher values).

Another use of medians is when the variance of the measurements are not well known or if a small fraction of measurements are suspected to be of dubious value. The χ^2 approach is powerful but only if the assumptions of normality of the data is true and the variances are known. These are very strong assumptions and application of χ^2 may yield results which a level of confidence that is not well deserved.

`Literature:` The `Wikipedia` entry on "Order Statistics" is comprehensive and understandable. The classical result of Equation 5.11 is quoted from Eric Weinstein's Mathworld. Gott et al. (Astrophysical Journal, 549, 1–17, 2001) provide a sound pedagogical discussion of when the use of median is preferred (§5.2). OS-8 is from S. K. Mishra (Economics Bulletin, 3, pp. 1–6, 2004).

`Applications:` Gott et al. (*ibid*) for application of median in determining a robust value of H_0 given measurements with poorly known variance. Frail et al. (2011, arXiv:1111.0007) for real examples of scaling or additive errors in measurements and a subsequent application of top m maxima. See Israel & Stella (ApJ 468, pp. pp. 369–379, 1996) for pre-whitening power spectrum of X-ray sources prior to searches for periodic signals.

Problem Set: Ordered Statistics

`Exercise OS-1(P):` Reason out Equation 5.8. Write a program to evaluate Equation 5.8 and thence reproduce Figure 5.2.

`Exercise OS-2(P):` Generate (say) $n = 101$ Gaussian random numbers with zero mean and unit variance. Determine the median and median. Repeat the exercise say $N = 10^5$ times. Plot the histogram of the mean and median estimators. Notice that the histogram of the median is more spread out than that of the mean. Redo the same exercise for an exponential distribution, $p(h) = \exp(-h)$ for $h \geq 0$ and 0 otherwise, and the Laplace distribution, $p(h) = 1/2 \exp(-|h|)$.

`Exercise OS-3(P):` Continuing with the case of $p(h) = \exp(-h)$ derive the mode of $\mathcal{H} = \max(h)$ as a function of the number of independent bins, n. Plot the distribution of $\rho(\mathcal{H})$ for $n = [10^2 10^6, 10^9, 10^{12}]$.

`Exercise OS-4(P):` The Cauchy distribution,

$$p(h) = \frac{1}{\pi(1 + h^2)} \tag{5.12}$$

with $-\infty < h < \infty$ is a distribution for which the central moments (mean and variance, in particular) are formally indefinite (IntREF). However, the median is well defined. Generate a sequence, say h_k for $k = 1, 2, .., n$ with n=101. Generate $m = 10^4$ such runs and for each run obtain the mean and the median. Comment on the histograms of the means and the medians.

`Exercise OS-5(H):` The Laplace distribution (with $\mu = 0$) is the distribution of two identically distributed exponential variates (hence sometimes this distribution is called the double exponential). Show that the maximum likelihood estimator of μ is given by the median and the median is the better estimator compared to the mean.

`Exercise OS-6(H):` The probability density function of the median can be obtained from Equation 5.8 with $m = (n + 1)/2$. Show (that is, derive and not merely restate Equation 5.11) that for large n and assuming Gaussian distribution, $N(\mu, \sigma^2)$ that the variance of the estimator of the median is $(\pi/2)\sigma^2/n$. Next, provide a general expression relating the the variance of the median to n and σ^2 (for large values of n). Extend the result for any percentile.

`Exercise OS-7(H):` The statistics of the Fourier transform of a series of measurements asymptotes to a uncorrelated Gaussian variates (Internal REF) as

the number of measurements increases. As a result, the statistics of the power spectrum (s) asymptote to χ^2_2 or exponential distribution, $p(s) = \exp(-s/A)/s_0$ where the A is the mean value and for white noise A is a constant, independent of frequency. However, many X-ray binaries exhibit a strong frequency dependence in their power spectrum with strong red noise (increasing power at low frequencies), quasi-periodic oscillations (resulting in broad humps at moderate frequencies) and flicker noise (contributing to the spectrum at high frequencies). These may arise from hot gas in the disk and the corona. Regular pulsations may arise from a specific region (e.g. the polar cap of the neutron star). Thus you expect to see a strong single-bin peak (assuming that the pulsation can be represented by a pure sinusoid) but against a varying background, the power spectrum $\mathcal{P}(f) = A(f)$ where $A(f)$ is given to you. Now work out the statistics of $\mathcal{H}$ (in the absence of any signal).

`Exercise OS-8(⋆):` In general the computation of the median involves $n\log(n)$ steps. In contrast the mean requires n steps and allows you to access the data sequentially. It has been suggested a weighted arithmetic mean yields an approximate value for the median:

$$w_i = \frac{1}{|x_i - \nu_1|} \qquad \nu_1 = \sum_{i=1}^{n} x_i w_i \Big/ \sum_{i=1}^{n} w_i. \tag{5.13}$$

Here, ν_1 is initially set to the mean in determining the weights, w_i. Then ν_1 is computed. The two steps repeated until necessary precision in convergence is obtained. I wonder what weights would be appropriate to similarly compute a certain percentile.

`Exercise OS-9(⋆):` For reasons discussed in the previous exercise it is sometimes advantageous to use median instead of the mean. This is especially the case when combining CCD images of the same field. After the images are registered to the same reference grid a median image is obtained by taking, for each position on the sky, a median of the pixels which map to that piece of sky. Assume that for a given position, the Poisson parameter is λ and that we have n images. Numerically compute the median and variance of the median as a function of λ (say from $\lambda = 0.1, 1, 3, 10, 20, 50, 100$) and n (say 5, 10, 30, 100).

Compare your numerical median estimator to an analytical estimate (REF: Wiki/Poisson Distribution).

$$\text{Median} = \text{floor}(\lambda + 1/3 - 0.02/\lambda) \tag{5.14}$$

To my knowledge there is no known simple expression for the variance of the

median. Make a sensible guess and obtain a fitting formula as a function of λ and n.

Solutions to Problem Set of Ordered Statistics

`Solution to OS-1(P):` For the mth maximum we need to have one point in interval, $[\mathcal{H}_m, \mathcal{H}_m + d\mathcal{H}_m]$ (for which there are n choices), $m-1$ points above this interval and $n-m$ below this interval. The number of choices for the latter is ${}^{n-1}C_{m-1}$ and the probability for a point being above the interval is $1-P(\mathcal{H}_m)$ but $P(\mathcal{H}_m)$ for being below the interval. This then leads to

$$\rho(\mathcal{H}_m) = np(\mathcal{H}_m) \times \frac{(n-1)!}{(m-1)!(n-m)!} P(\mathcal{H}_m)^{n-m}[1-P(\mathcal{H}_m)]^{m-1}. \quad (5.15)$$

For large n direct evaluation of the above Equation is not possible with even double precision arithmetic. We recognize that the primary interval of interest is the vicinity where $\rho(\mathcal{H}_m$ is not exceedingly close to zero. This happens when $P(\mathcal{H}_m)$ approaches unity. We then use the approximation $(1-\epsilon)^n \approx \exp(-n\epsilon)$ when $\epsilon \to 0$ but $n\epsilon$ tends to a non-zero value. Furthermore, we make use of the approximation that $n-m \approx n$ when $n \gg 1$ and $m \ll n$.

$$\rho(\mathcal{H}_m) \approx \frac{1}{(m-1)!} p(\mathcal{H}_m)[1-P(\mathcal{H}_m)]^{m-1} \exp\Big[-(n-m)[1-P(\mathcal{H}_m)]\Big] \quad (5.16)$$

`Solution to OS-2(P):` See Figures 5.3, 5.4 and 5.5. The Gaussian variates were generated with the Box-Muller transformation; for exponential $h = -\log(x)$ where x is a uniformly distributed random number, $[0,1]$; the Laplace variates were generated as the difference between two unit variance exponential variates or alternatively as $\mathrm{sgn}(y)\log(1-2y)$ where y is uniformly distributed over the range $[-1/2, 1/2]$.

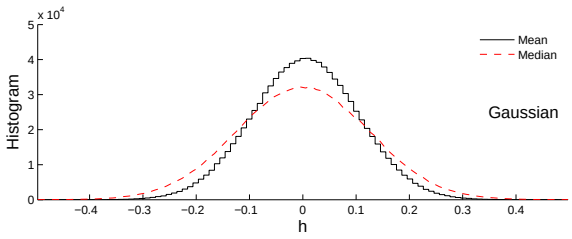

Figure 5.3: A sequence of $n = 1001$ Gaussian variates, h, with zero mean and unit variance were generated, $p(h) = 1/\sqrt{2\pi}\exp(-h^2/2)$. The mean and median for each run was obtained. The run is repeated $N = 10^6$ times and the histogram of the resulting means and medians is shown above.

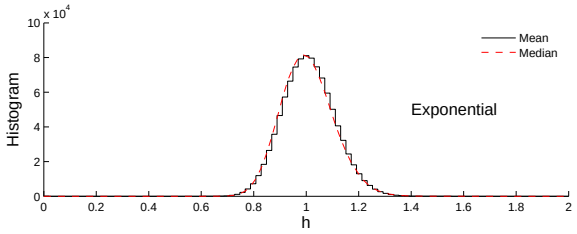

Figure 5.4: Same as Figure 5.3 with $n = 101$, $N = 10^6$ and $p(h) = \exp(-h)$ for $h \geq 0$ and 0 otherwise.

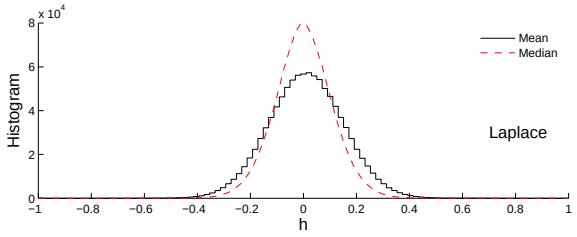

Figure 5.5: Same as Figure 5.3 with $n = 101$, $N = 10^6$ and $p(h) = \exp(-|h|)$.

`Solution to OS-3(P):` The probability density function and the cumulative function are $p(h) = \exp(-h)$ and $P(h) = 1 - \exp(-h)$. The probability density function and the cumulative function of $\mathcal{H}$, the maximum of h, is given by (cf. Equation 5.6)

$$\rho(\mathcal{H}) \approx n \exp(-\mathcal{H}) \exp\Big[-(n-1)\exp(-\mathcal{H})\Big], \qquad \varrho(\mathcal{H}) \approx \exp\Big[-n\exp(-\mathcal{H})\Big] \tag{5.17}$$

The mode of $\rho(\mathcal{H})$ is the value of $\mathcal{H}$ at the peak of $\rho(\mathcal{H})$ and can be shown to be

$$H_{\text{mode}} = \ln(n-1). \tag{5.18}$$

Thus, $H_{\text{mode}} = [13.8, 20.7, 27.6]$ for $n = [10^6, 10^9, 10^{12}]$. The plots of $\rho(\mathcal{H})$ can be found in Figure 5.6. Notice the constant width of $\rho(\mathcal{H})$ with respect to the sample size, n. Compare this to Figure 5.1.

`Solution to OS-4(P):` The transformation $h = \tan(\pi(x - 1/2))$ generates zero mean Cauchy distribution with pdf given by Equation 5.12; here x is a random variate uniformly distributed over the range $[0, 1)$. A simulation run will show that the median is nicely distributed over the approximate interval between -1 and $+1$. In contrast, the mean is wildly and randomly distributed over a much larger range.

`Solution to OS-5(P):` Not solved yet.

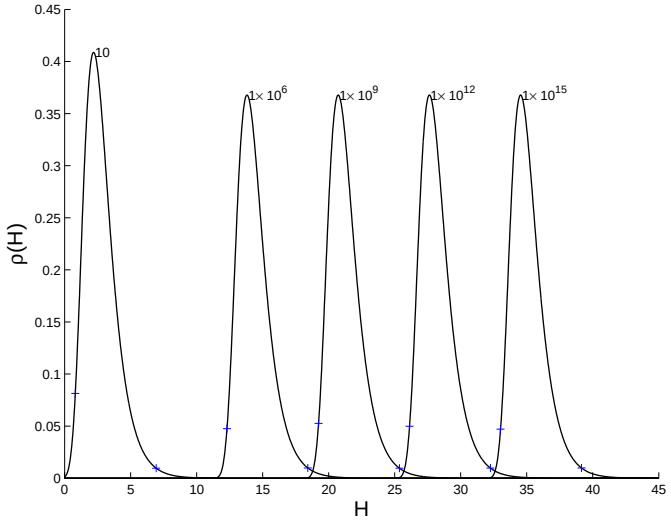

page

Figure 5.6: Plot of $\rho(H)$, the probability density function for the maximum, $\rho(H)$ for sample size of $n = [10, 10^6, 10^9, 10^{12}, 10^{15}]$ and assuming that the input series follows a unit mean exponential distribution. The 1% and 99% confidence points are marked by a "+" and the intervals are [0.8 6.9], [12.3, 18.5], [19.2 25.3], [26.1 32.3], [33.0, 39.2] – essentially a constant width with respect to n.

`Solution to OS-6(H)`: From Figures 5.3 and 5.4 we see that the distribution of the median looks like a Gaussian distribution, despite differing $p(h)$. We will make the assumption that for large values of n the distribution of the median, at least in the vicinity of the median, is also Gaussian, $\rho(h_m) \propto \exp(-(h - h_m)^2/(2\sigma_h^2)$; here $\rho(h_m)$ is the probability density function of the sample medium. If so, it is easy to show that

$$\frac{d^2\rho}{dh_m^2} = -\frac{1}{\sigma_{h_m}^2}. \tag{5.19}$$

While this equation is true for a Gaussian distribution, at any value of h, we will make use of this result only in the vicinity of the peak $\rho(h_m)$ for which our simulations, as discussed above show is a plausible approximation.

The probability density function for the sample median, h_m, is obtained from Equation 5.8 with $r = (n-1)/2$ and the result is

$$\rho(h_m) = \frac{n!}{r!r!} p(h_m) P^r [1 - P(h_m)]^r. \tag{5.20}$$

Taking the $\log(\rho)$ and differentiating it twice yields

$$\frac{d^2\log(\rho)}{dh_m^2} = \frac{d^2\log(p)}{dh_m^2} - r\Big(\frac{dP}{dh_m}\Big)^2\Big[\frac{1}{P^2} + \frac{1}{(1-P)^2}\Big] + r\frac{d^2P}{dh_m^2}\Big[\frac{1}{P} - \frac{1}{1-P}\Big], \tag{5.21}$$

The inverse of the variance of the median is obtained by evaluating the above equation at the true value of the median, $H_m = h_m$. The median is defined as the value of h_m which satisfies the condition, $P(h_m = H_m) = 1/2$ and thus the last term in Equation 5.21 is identically equal to zero. Substituting this value in Equation 5.21 and noting $dP/dh = p(h)$ we obtain

$$\begin{aligned} \frac{1}{\sigma^2_{H_m}} &= \Big[- \frac{d^2 \log(p)}{dh_m^2} + 8r\, p(h_m)^2 \Big]_{h_m = H_m} \\ &\approx 4(n-1)p(H_m)^2 \end{aligned} \tag{5.22}$$

In arriving at Equation 5.22 we ignored the first term in Equation 5.21 (which is approximately the variance of a single sample) relative to the second term which is $\propto n\sigma^2$.

We now apply Equation 5.22 for the specific case of a Gaussian distribution, $N(\mu, \sigma^2)$. With $H_m = \mu$ we find

$$\sigma^2_{H_m} \approx \frac{\pi}{2} \frac{\sigma^2}{n-1} \tag{5.23}$$

leading us to conclude that the variance of the median is $\pi/2$ worse than that of the mean.

Next, we apply the general formula (Equation 5.22) to an exponential distribution, $p(h) = \exp(-h)$ for $h \geq 0$ and zero otherwise. The cumulative distribution is $P(h) = 1 - \exp(-h)$. The mean, $\mu = 1$. The median is obtained by noting that $P(H_m) = 1/2$ or $H_m = \log(2)$. We note that $p(H_m) = 1/2$. Thus $\sigma^2_{H_m} \approx 1/(n-1)$. The variance in the sample mean is also $\approx 1/(n-1)$.

`Solution to OS-7(H):` The probability density function for each channel is an exponential but with a mean which is channel dependent. In particular, let $A(f)$ be the expected mean value for channel f. Then, to the extent $A(f)$ is characterized we can convert the value in each channel, $P(f)$ to a significance level, $h(f) \propto \exp(-P(f)/A(f))$. We see that $h(f)$, being a probability, is uniformly distributed on $[0, 1]$. Thus, the probability density of $\mathcal{H}$ is

$$\rho(\mathcal{P}_m) = n\mathcal{H}^{n-1} \qquad \varrho(\mathcal{H}) = \mathcal{H}^n. \tag{5.24}$$

The problem with this approach is that if small errors in determining $A(f)$ can lead to significant variation in $h(f)$. For this reason, determining $A(f)$ with high precision is important. One approach is to "pre-whiten" the spectrum.

`Solution to OS-8(⋆):` Not thought about it.

`Solution to OS-9(⋆):` Not thought about it.

Chapter 6

Bias & Confusion

Chapter 7

Digitizing Intensity

7.1 1-bit Quantization

Let $\mathcal{N}$ be a zero mean, unit variance Gaussian variate $[N(0,1)]$. Let $\mathcal{S}$ be the signal (assumed, for the purpose of the discussion, to be steady). The input and the output to the 1-bit digitizer is thus

$$\begin{aligned} \mathcal{Y} &= \mathcal{S} + \mathcal{N} && (7.1) \\ y &= \mathrm{sgn}(\mathcal{Y}) && (7.2) \end{aligned}$$

Our goal is to evaluate the signal-to-noise ratio of y, defined as $\langle y \rangle / \sqrt{V(y)}$ where $V(y)$ is the variance and $\langle y \rangle$ is the mean value. We wish to compare this with the SNR of the input signal,

$$\mathrm{SNR}(\mathcal{Y}) = \mathcal{S} \tag{7.3}$$

(since the variance of the input signal is unity).

Given that $\mathcal{N}$ follows Gaussian statistics with zero mean and unit variance we deduce that the probability distribution of $\mathcal{Y}$ is

$$p(\mathcal{Y}) = \frac{1}{\sqrt{2\pi}} \exp\Big(- \frac{(\mathcal{Y} - \mathcal{S})^2}{2} \Big). \tag{7.4}$$

The mean and the variance of $\mathcal{Y}$ is simply $\mathcal{S}$ and 1, respectively. The mean of the digitized signal is

$$\begin{aligned} \langle y \rangle &= \frac{1}{\sqrt{2\pi}} \Bigg[\int_0^\infty \exp\Big(- \frac{(\mathcal{Y} - \mathcal{S})^2}{2} \Big) d\mathcal{Y} - \int_{-\infty}^0 \exp\Big(- \frac{(\mathcal{Y} - \mathcal{S})^2}{2} \Big) d\mathcal{Y} \Bigg] \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\mathcal{S}}^{\mathcal{S}} \exp\Big(- \frac{\mathcal{Y}^2}{2} \Big) d\mathcal{Y} \end{aligned}$$

$$= \operatorname{erf}(\mathcal{S}/\sqrt{2}) \tag{7.5}$$

The second moment of y is simply unity[1]

$$\begin{aligned} \langle y^2 \rangle &= \frac{1}{\sqrt{2\pi}} \left[\int_0^\infty \exp\Big(-\frac{(\mathcal{Y}-\mathcal{S})^2}{2} \Big) d\mathcal{Y} + \int_{-\infty}^0 \exp\Big(-\frac{(\mathcal{Y}-\mathcal{S})^2}{2} \Big) d\mathcal{Y} \right] \\ &= 1 \end{aligned} \tag{7.6}$$

and thus

$$V(y) = \langle y^2 \rangle - \langle y \rangle^2 = 1 - \operatorname{erf}(\mathcal{S}/\sqrt{2})^2. \tag{7.7}$$

Thus the SNR of a *single* sample of the digitized signal is

$$\mathrm{SNR}(y) = \frac{\operatorname{erf}(\mathcal{S}/\sqrt{2})}{\sqrt{1 - \operatorname{erf}(\mathcal{S}/\sqrt{2})^2}} \tag{7.8}$$

With Equations 7.3 and 7.8 we can now plot the SNR of a 1-bit digitizer (Figure 7.1).

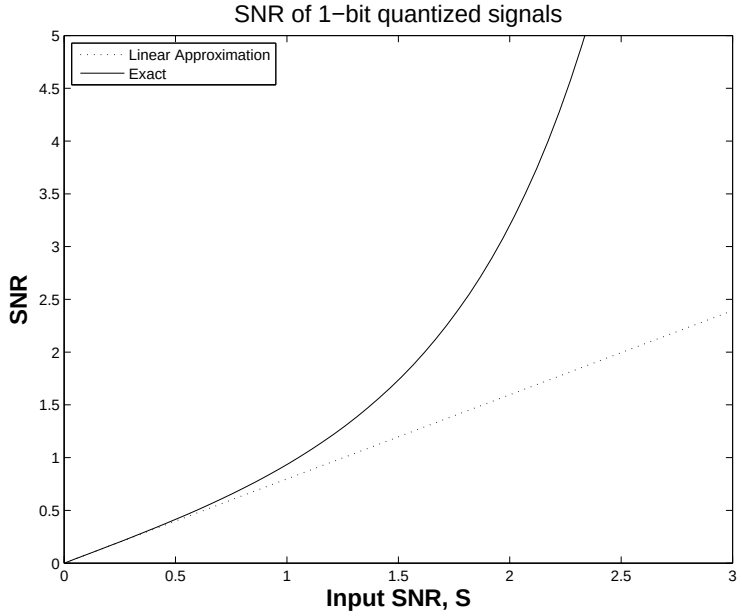

Figure 7.1: Signal-to-noise ratio of 1-bit digitized intensity signal.

[1] At least in the framework we specified and that is $\mathcal{S}$ is noiseless. Any real signal will induce its own additional noise and a correct treatment would take effect into account. However, 1-bit digitizers are or should be employed for weak signals and so this discussion is academic even in the context of an academic discussion.

At low input SNR (by which we mean $\mathcal{S}$ since the input noise has a variance of unity) the output SNR of the 1-bit digitizer bears a linear relation with the input SNR. However, as the input SNR increases the output SNR increases exponentially – a very curious state of affairs. Let us try to understand both these features.

First we consider the case of $\mathcal{S} \ll 1$. Taylor expansion of the RHS of Equation 7.8 leads to

$$\mathrm{SNR}(y) = \sqrt{\frac{2}{\pi}}\mathcal{S} \tag{7.9}$$

Thus the SNR of the 1-bit signal is $\sqrt{\pi/2}$ worse than that of the input signal by about 20%.

Now consider the opposite case, $\mathcal{S} \gtrsim 1$. In this limit, $\mathrm{erf}(\mathcal{S}/\sqrt{2}$ approaches unity and. Thus, as can be seen from Equation 7.8, $\mathrm{SNR}(y)$ then increases rapidly. In fact, it can be shown to increase exponentially! It is easy to understand this result since in this limit the variance decreases (most of the samples are +1) and the signal saturates to unity. Thus we are in the paradoxical situation in which the SNR of the digitized signal exceeds that of the analog signal. In fact, this phenomenon presents a very good opportunity to understand the limitations of the concept of SNR.

The fundamental quantity that matters in measurement theory is inferring the range of the underlying physical quantity. In our case it is $\mathcal{S}$. From Equation 7.5 we find that variation in $\langle y \rangle$ is

$$\Delta\langle y \rangle = \sqrt{\frac{2}{\pi}} \exp(-\mathcal{S}^2/2)\Delta\mathcal{S}. \tag{7.10}$$

We invert this equation and obtain

$$\Delta\mathcal{S} = \sqrt{\frac{\pi}{2}} \exp(\mathcal{S}^2/2)\Delta\langle y \rangle. \tag{7.11}$$

For $\mathcal{S} \gtrsim 1$ you can see that a small variation in the measured value of y results in a larger variation in the inferred value of $\mathcal{S}$. This simple example shows the limitation of using SNR as a proxy for the confidence interval of $\mathcal{S}$.

The lesson learnt is simple: normally increasing SNR also means better knowledge of the inferred parameter. However, here, the larger SNR is a mirage. Our knowledge of $\mathcal{S}$ decreases with increasing value of $\mathcal{S}$. On the other hand our knowledge that the input signal has a strength comparabele to the noise becomes more certain with increasing $\mathcal{S}$.

This situation is similar to the photon statistics in the low light and zero background case: a certainty of detection but near absence of a quantitative

knowledge of the amplitude of the signal. On some reflection you can see that the probability distribution of the negative state (-1) indeed approach Poisson distribution as $\mathcal{S}$ becomes larger. Here we have an example of a Gaussian process which approaches Poisson limit (rather than the other way around).

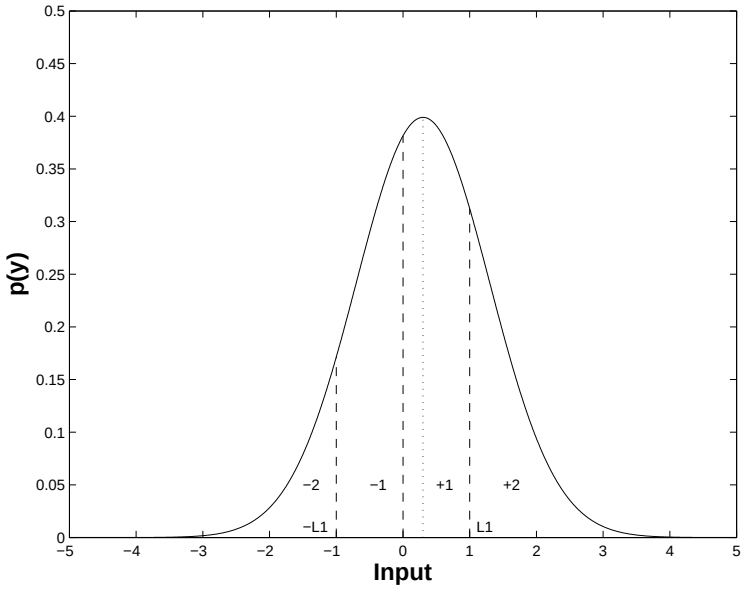

Figure 7.2: Two bit digitizer.

7.2 2-bit Digitizer

The next level of sophistication is a 2-bit digitizer with transitions at 0, $\pm L_1$ (Figure 7.2). The four possible states $(j = 1, 2, 3, 4)$ for the digital signal y are $S(j) = -Q, -1, +1, +Q$ where for generality we let the state 1 and 4 have a value ≥ 2. Given the discussion in the previous section we assume that $\mathcal{S} \ll 1$.

Adopting the same model as before (Equation 7.2) we find the probability of each state for a signal to be as follows:

$$\begin{aligned} P(1) &= 1/2\Big[1 - \mathrm{erf}((L_1 + \mathcal{S})/\sqrt{2})\Big] \\ P(2) &= 1/2\Big[\mathrm{erf}((L_1 + \mathcal{S})/\sqrt{2}) - \mathrm{erf}(\mathcal{S}/\sqrt{2})\Big] \\ P(3) &= 1/2\Big[\mathrm{erf}((L_1 - \mathcal{S})/\sqrt{2}) + \mathrm{erf}(\mathcal{S}/\sqrt{2})\Big] \end{aligned}$$

$$P(3) \;=\; 1/2\Big[1-\mathrm{erf}((L_1-\mathcal{S})/\sqrt{2})\Big]. \tag{7.12}$$

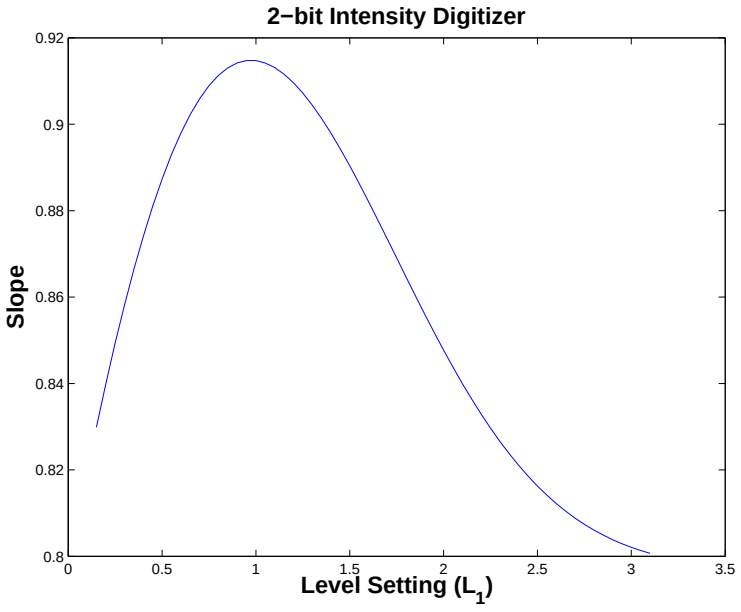

Figure 7.3: Slope of the SNR-$\mathcal{S}$ for a 2-bit digitizer as a function of L_1.

The mean and second moment can be computed as before: $\langle y\rangle = \sum P(j)S(j)$ and $\langle y^2\rangle = \sum P(j)S(j)^2$.

$$\begin{aligned}\langle y\rangle &= \frac{Q-1}{2}\left[\mathrm{erf}((L_1+\mathcal{S})/\sqrt{2})-\mathrm{erf}((L_1-\mathcal{S})/\sqrt{2})\right]+\mathrm{erf}(\mathcal{S}/\sqrt{2})\\ &\approx (Q-1)\exp(-L_1^2/2)\sqrt{\frac{2}{\pi}}\mathcal{S}+\sqrt{\frac{2}{\pi}}\mathcal{S}\\ &\approx \mathcal{S}\sqrt{\frac{2}{\pi}}\left[(Q-1)\exp(-L_1^2/2)+1\right]\end{aligned} \tag{7.13}$$

where the approximation is valid for $\mathcal{S}\ll 1$.

$$\begin{aligned}\langle y^2\rangle &= Q^2P(1)+P(2)+P(3)+Q^2P(4)\\ &= 1+(Q^2-1)[1-1/2\mathrm{erf}((L_1+S)/\sqrt{2})-1/2\mathrm{erf}((L_1-S)/\sqrt{2})]\\ &\approx 1+(Q^2-1)\Big[1-\mathrm{erf}(L_1/\sqrt{2})\Big].\end{aligned} \tag{7.14}$$

We note that the output signal increases linearly with the strength of the input signal. In the limit we are considering $\mathcal{S}\sim 0$ the variance is approximately independent of the input signal. For $Q=2$ we find the output SNR

is related to the input SNR as

$$\mathrm{SNR}(y) \approx \mathcal{S}\left(\sqrt{\frac{2}{\pi}}\right) \frac{(Q-1)\exp(-L_1^2/2)+1}{\sqrt{1+(Q^2-1)[1-\mathrm{erf}(L_1/\sqrt{2})]}} \tag{7.15}$$

The simplest choice is $Q = 2$ for which $L_1 \sim 1$ optimizes the output SNR (see Figure 7.3). The efficiency is about 92%. The efficiency monotonically decreases to that of the 1-bit digitizer for smaller values of L_1. For larger values of L_1 the efficiency decreases since states 1 and 4 become increasingly sparsely populated.

Exercise: Find out which combination of Q and L_1 maximizes the efficiency of a 2-bit digitizer.

Chapter 8

Combining Experiments

Consider two experiments aimed at measuring the same quantity. Each experiment has a different variance: σ_i^2 for $i = 1, 2$. What is the best way to combine the two measurements, x_j?

Let w_i be the weights assigned to each measurement:

$$X = w_1 x_1 + w_2 x_2 \text{ with} \tag{8.1}$$

The two weights should add to unity in order to ensure that the mean value of X is indeed the sample mean. Thus,

$$1 = w_1 + w_2. \tag{8.2}$$

We use the signal-to-noise ratio (SNR) as the metric to maximize. In the mean $\langle X \rangle = w_1\mu_1 + w_2\mu_2$ and the variance of X, $V = w_1^2\sigma_1^2 + w_2^2\sigma_2^2$. Thus the SNR is

$$\mathcal{S} = \frac{w_1\mu + w_2\mu}{\sqrt{w_1^2\sigma_1^2 + w_2^2\sigma_2^2}}. \tag{8.3}$$

The goal is to determine the weights which maximize SNR(X). Given Equation 8.2 there is only one free parameter and we arbitrarily choose the free parameter to be w_1. The condition for the maximum of SNR(X) is obtained by setting the partial derivative of SNR(X) with respect to w_1 to zero.

$$\frac{\partial \mathcal{S}}{\partial w_1} = (w_1^2\sigma_1^2 + w_2^2\sigma_2^2) - (w_1 + w_2)w_1\sigma_1^2 = 0 \tag{8.4}$$

which when combined with 8.2 yields

$$w_1 = \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} = \frac{\sigma_1^{-2}}{\sigma_1^{-2} + \sigma_2^{-2}}. \tag{8.5}$$

By symmetry the weighting can be generalized to $w_j \propto \sigma_j^{-2}$. This is the origin of the so-called weighted mean. The SNR of the weighted mean is obtained by subsitutiong the above weights into Equation 8.3 and is

$$\mathcal{S} = \mu\sqrt{\sum_i \frac{1}{\sigma_i^2}} = \Big(\sum_i \frac{\mu^2}{\sigma_i^2}\Big)^{1/2}. \tag{8.6}$$

For the case of $\sigma_i = \sigma$ we recover the usual result of the SNR improving by $\sqrt{n}$ relative to a single measurement.

Next, consider an experiment which measures $j = (1, 2, ..., n)$ separate channels. Say, the measurement noise is the same in each channel. However, the expected signal strength is different in each channel. The expected strength in each channel can be known in advance of the measurement (except for the overall strength of the signal), e.g. a resolved spectral line with a known center wavelength but unknown flux.

As before consider

$$X = \sum_{j=1}^{n} w_j x_j \quad \text{with} \tag{8.7}$$

with the requirement that the sum of the weights add to unity, $\sum_j w_j = 1$. The SNR of X is

$$\mathcal{S} = \frac{\sum_{j=1}^n w_j \mu_j}{\sigma\sqrt{\sum_{j=1}^n w_j^2}}. \tag{8.8}$$

Setting the partial derivatives to zero yields

$$\frac{\partial \mathcal{S}}{\partial w_i} = \mu_i \sum_{j=1}^{n} w_j^2 - w_i \sum_{j=1}^{n} w_j \mu_j = 0. \tag{8.9}$$

Thus $w_i \propto \mu_i$ and with the requirement for unit sum for the weights we find

$$w_i = \frac{\mu_i}{\sum_{j=1}^n \mu_j}. \tag{8.10}$$

This is the so-called "matched filter" solution. Substituting w_i into Equation 8.8 yields the SNR of the Xf:

$$\mathcal{S} = \frac{1}{\sigma}\sqrt{\sum_i \mu_i^2} = \Big(\sum_i \frac{\mu_i^2}{\sigma^2}\Big)^{1/2}. \tag{8.11}$$

For the special case, $\mu_i = \mu$ we recover the well known result that $\mathcal{S} = \sqrt{n}(\mu/\sigma)$.

Finally, we consider the most general case: allow each channel to have its own variance, σ_i^2 and its own mean value μ_i. By following the steps as above we find

$$w_i = \frac{\mu_i/\sigma_i^2}{\sum_j \mu_j/\sigma_j^2} \tag{8.12}$$

maximzes $\mathcal{S}$. We could call this as the weighted matched filter. The SNR is

$$\mathcal{S} = \Big(\sum_i \frac{\mu_i^2}{\sigma_i^2} \Big)^{1/2} \tag{8.13}$$

In effect, the square of the SNR of n measurements, when combined optimally, is the the sum of the squares of the SNRs of each measurement (or channel).

8.1 Application: Optical Imaging

Optical imaging is done with a camera followed by a CCD detector. We will assume that the pixel size is chosen to be a fraction of width of the point spread function (PSF). For space based imagers the PSF width is approximately λ/D whereas for ground-based imagers it is set by atmospheric seeing (angular diameter, $\theta_{\rm FWHM}$). The primary source of noise is due to the emission from the background (zodiacal dust emission, reflection by interstellar grains and emission from the atmosphere). CCDs are getting so good that one, in most cases, ignore read noise.

For simplicity we will consider circularly symmetric extended sources. The surface brightness of the source is described by

$$I(\vec{\theta}) = \mathcal{I} g(\theta) = \mathcal{I} \exp\Big[-(\theta_g/\theta)^2 \Big] \tag{8.14}$$

where $\vec{\theta} = (\theta_x, \theta_y)$, $\theta^2 = \theta_x^2 + \theta_y^2$ and $\mathcal{I}$ is the central surface brightness (unit: photons per square centimeter per second per pixel). Let $\mathcal{B}$ be the background surface density and assumed to be constant across the image. The unit for both $\mathcal{I}$ and $\mathcal{B}$ is photons per square centimeter per second per pixel.

Say your goal is to estimate $\mathcal{I}$. Going forward we will assume that the mean background level has been subtracted out from the images (with negligible noise penalty). Following the discussion centered around Equation 8.12

the pixel data should be combined with the following weight

$$w(\vec{\theta}) \propto \frac{g(\vec{\theta})}{I(\vec{\theta}) + \mathcal{B}} \tag{8.15}$$

where it is understood that the integral of $w(\vec{\theta})$ is unity. The weighted sum then yields an optimal value for the central surface brightness of the source:

$$\mathcal{I}_0 = \int w(\vec{\theta}) I(\vec{\theta}) d\vec{\theta}. \tag{8.16}$$

The SNR of $\mathcal{I}_0$ is

$$\mathcal{S} = \left(\mathcal{I}_0 \int \left[\frac{g(\theta) I(\theta)}{I(\theta) + \mathcal{B}} \right] 2\pi\theta d\theta \right)^{1/2} \tag{8.17}$$

Above we assumed that the collecting area (A) and the integration time (t) was unity. The SNR should be increased by $\sqrt{At}$ for the general case.

Problem Set: Combining Experiments

`Exercise CE-1(P):` Qualitatively argue that the in the bright limit ($\mathcal{I} \gg \mathcal{B}$) the best SNR is $\sqrt{F_S}$ where F_S is the total number of photons detected. Verify that Equation 8.17 yields this value. Show that in the opposite limit ($\mathcal{I} \lesssim \mathcal{B}$)

$$\mathcal{S} \approx \frac{\mathcal{I}_0}{\sqrt{2\mathcal{B}}}\sqrt{\pi\theta_g^2} = \frac{F_G}{\sqrt{2\pi\theta_g^2\mathcal{B}}}. \tag{8.18}$$

Thus at any given SNR level the limiting surface brightness $\propto \theta_g^{-1}$. the SNR of the integrated flux density scales as θ_g^{-1}.

`Exercise CE-2(P):`

Need some ideas

Solutions to Problem Set: Combining Experiments

`Solution to CE-1(P):` The flux density from the object is $F_G = I(\theta)2\pi\theta d\theta = \mathcal{I}_0\pi\theta_g^2$. In the limit of for $\mathcal{B} \gg \mathcal{I}_0$ Equation 8.17 simplifies to

$$\mathcal{S} = \Big(\mathcal{I}_0 \int g(\theta)2\pi\theta d\theta\Big)^{1/2} = F_G \tag{8.19}$$

In the opposite regime

$$\mathcal{S} = \mathcal{I}_0\Big(\int \frac{g(\theta)I(\theta)}{\mathcal{B}}2\pi\theta d\theta\Big)^{1/2}. \tag{8.20}$$

In the mean $I(\theta) = \mathcal{I}_0 g(\theta)$ and thus Equation 8.20 simplifies to $\mathcal{S} = \mathcal{I}_0\sqrt{\pi\theta_g^2/(2\mathcal{B})}$.

Chapter 9

Gaussian Distribution: Famous Examples

9.1 Astrometry

9.2 Fourier Transforms

9.3 Ruze's Formula (Strehl Ratio)

One of the basic formula familiar to radio astronomers is Ruze's formula for radio telescopes. Assume that each segment is well laid down so as to hug a parabola. However, the segment surface is rough (to some degree). Let δ_k be the path length error of each "patch" (which is either a region of a segment or the segment itself). The corresponding phase error is

$$\phi_k = 2 \times 2\pi\delta_k/\lambda \tag{9.1}$$

where the factor of two accounts for reflection (in and out). We will assume that each panel or region contributes equal amount of flux. Thus the primary error ("piston") is the phase error shown above. Let $p(\phi)$ describe the statistics of the phase error. We assume $\langle\phi\rangle = 0$ and note the variance is

$$V(\phi) = \left(\frac{4\pi\sigma_\delta}{\lambda}\right)^2 \tag{9.2}$$

where σ_δ^2 is the variance in the piston errors.

Each region/panel contributes the following electric field at the focus:

$$E_R(k) \;\; = \;\; E_0 \cos(\phi_k)$$

$$E_I(k) \;=\; E_0 \sin(\phi_k). \tag{9.3}$$

The total electric field is thus

$$\begin{aligned} E_{\rm tot} &= \sum_k E_0 \cos(\phi_k) + j \sum_k E_0 \sin(\phi_k) \\ &= E_R({\rm tot}) + jE_I({\rm tot}) \\ &= nE_0 \Big\langle \exp(j\phi) \Big\rangle \end{aligned} \tag{9.4}$$

where

$$\exp(\phi) = \int p(\phi) \exp(j\phi) d\phi. \tag{9.5}$$

Here, $p(\phi)$ is the probability distribution of ϕ and the integration limits range from $-\infty$ to $+\infty$.

We now consider the specific case of a Gaussian distribution

$$p(\phi) = \frac{1}{\sigma_\phi \sqrt{2\pi}} \exp\Big(- \frac{\phi^2}{2\sigma_\phi^2} \Big). \tag{9.6}$$

As will become apparent it is better to compute a slightly more general expression, namely the Fourier transform of $p(\phi)$,

$$\begin{aligned} G(t) &\equiv \frac{1}{\sigma_\phi \sqrt{2\pi}} \int \exp\Big(- \frac{\phi^2}{2\sigma_\phi^2} \Big) \exp(jt\phi) d\phi \\ &= \exp\Big(- \frac{\sigma_\phi^2 t^2}{2} \Big); \end{aligned} \tag{9.7}$$

here we utilize the well known expression for the characteristic function (Fourier transform) of the normal distribution. As expected the sine component is zero leaving only the cosine transform. Thus

$$E_{\rm tot} \;=\; nE_0 \exp\Big(- \sigma_\phi^2/2 \Big). \tag{9.8}$$

The specific quantity we seek is simply $I_{\rm tot} = \langle (E_R + jE_I)(E_R - jE_I) \rangle$. At first blush one may think that $I_{\rm tot} = E_{\rm tot} E_{\rm tot}^*$. However, this is incorrect, since, in general, $\langle xx^* \rangle \neq \langle x \rangle \langle x^* \rangle$.

The correct approach is to derive the mean value of $I_{\rm tot}$ directly:

$$I_{\rm tot} \;=\; E_0^2 \sum_{k,l} \exp(j\phi_k - j\phi_l) \tag{9.9}$$

which requires the evaluation of

$$\Big\langle \exp(j\phi - j\phi') \Big\rangle \;=\; \int\int d\phi d\phi' p(\phi) p(\phi') \exp(j\phi - j\phi')$$

$$= \delta(\phi,\phi') + (1-\delta(\phi,\phi'))\Big\langle \exp(j\phi)\Big\rangle. \qquad (9.10)$$

Thus

$$\langle I_{\rm tot}\rangle = E_0^2\Big[n + n(n-1)G(1)\Big]. \qquad (9.11)$$

The maximum value of is obtained when all the rays are combined in phase, $I_{\rm tot} = n^2E_0^2$. Normalizing with this value we obtain the celebrated Ruze formula for the aperture efficiency of radio telescopes:

$$\eta = \frac{1}{n} + \frac{n-1}{n}\exp\left[-\left(\frac{4\pi\sigma_\delta}{\lambda}\right)^2\right]. \qquad (9.12)$$

The first term $(1/n)$ sets a minimum value to the aperture efficiency and is obtained even when all the contributing rays have random phases. In this case, $E_{\rm tot}$ is zero but there is still some power resulting from incoherent combination of all the rays. The second term arises from coherent sum of all the rays and is related to the quality of the surface.

The Ruze formula includes only piston errors (no tilts, no focus). Assuming that $n \gg 1$ we see that $\eta \geq 0.5$ requires $\sigma_\delta \leq \lambda/15$, a rough rule of thumb known to radio astronomers.

Chapter 10

The Fourier Transform

10.1 The Ubiquitous Transform

The Fourier transform appears in many fields of physics and signal processing:

- Quantum Mechanics. The wave function in spatial coordinates and that in momentum space bear a Fourier relation to each other.
- A lens is a Fourier transforming machine. NEED MORE.
- The basis of interferometry is the van Cittert-Zernike theorem. The visibility function and the far field intensity distribution bear a Fourier relation to each other.

The definition of Fourier transforms depends on the sub-fields. We will use the following definition for a forward Fourier transform:

$$H(f) = \int_{-\infty}^{\infty} h(t) \exp(-j2\pi ft)dt. \tag{10.1}$$

We will use lower case symbols for functions of time domain and the upper case for their Fourier transforms. The inverse transform is defined to be

$$h'(t) = \int_{-\infty}^{+\infty} H(f) \exp(j2\pi ft)df \tag{10.2}$$

Since this is a book on signals and noise that you encounter in real systems we will not worry whether all functions have Fourier transform. We will assume that the above integrals exist and that the $h't(t) = h(t)$ because (*what result?*)

$$\int_{-\infty}^{+\infty} \exp(j2\pi f(t-t')df = \delta(t-t'). \tag{10.3}$$

For moderately discontinuous signals we note that

$$h't(t) = \frac{h(t^+) + h(t^-)}{2}; \tag{10.4}$$

otherwise know as the *Dirichlet* result?? We are now ready to explore the essentials of Fourier transforms.

If you are serious about understanding this book you have to become familiar with the following ?? Fourier transforms.

Box Car. Consider the very simple function

$$\begin{aligned} b(t,T) &= 1 \qquad |t| < T/2 \\ &= 0 \qquad \text{otherwise} \end{aligned} \tag{10.5}$$

The Fourier transform is

$$\begin{aligned} B(f) &= \frac{\exp(j\pi fT) - \exp(-j\pi fT)}{j2\pi f} \\ &= T\frac{\sin(\pi fT)}{\pi fT} \\ &= T\text{sinc}(\pi fT). \end{aligned} \tag{10.6}$$

The box-car function occurs so frequently in Fourier transforms that a new function "$\text{sinc}(x) \equiv \sin(x)/x$" has been invented. We note that $\text{limit}_{x\to 0}\text{sinc}(x) = 1$. This can be seen by considering Taylor expansion of $\sin(x)$ or by the application of L'Hopital's rule.

Note the following: the first zero of the function is at $f = T^{-1}$ i.e the first zero occurs at the inverse of the full width of the box car function. Next, $B(0) = T$ which is the integral of $b(t,T)$.

Cosine and Sine. The Fourier transform of a pure tone, $l(t) = cos(2\pi f_0 t)$ is

$$L(f) = \frac{1}{2}\delta(f + f_0) + \frac{1}{2}\delta(f - f_0). \tag{10.7}$$

A δ function at $f = f_0$ is expected. However, we also have a δ function at the negative frequency $-f_0$. Indeed, all real signals have power at negative frequencies and the negative functions are as real as positive frequencies. We will discuss this slightly disconcerting fact in some details below.

The Fourier transform of a pure sinusoidal tone, $l(t) = sin(2\pi f_0 t)$ is

$$L(f) = -\frac{1}{2j}\delta(f + f_0) + \frac{1}{2j}\delta(f - f_0). \tag{10.8}$$

Note that the components are purely imaginary and have the opposite phase. Compare and contrast with the Fourier transform of a cosine pure tone.

Delta Function. Consider $h(t) = \delta(t - t_0)$. The Fourier transform is

$$H(f) = \exp(j2\pi f t_0). \tag{10.9}$$

The amplitude $H(f)$ is unity for all values of f. The phase of the $H(f)$ increases linearly with frequency, $\phi(f) = 2\pi f t_0$. The phase gradient is proportional to t_0:

$$\frac{1}{2\pi}\frac{d\phi}{df} = t_0. \tag{10.10}$$

DC or Constant Value. Consider a dc signal, $h(t) = k$ for all t. The Fourier transform is

$$H(f) = k\delta(f). \tag{10.11}$$

As expected such a "dc" signal has no power at any frequency other than zero.

Exponential.

Gaussian.

10.2 Properties of Fourier Transforms

10.2.1 Linear Transformations.

All said and done, a Fourier transform is a linear transform of the input (say time series) data. Maintaining this perspective will make Fourier transforms less obscure. What are the consequences of linear transforms? Simply scaling up the input data will result in the output being scaled up by the same factor. The Fourier transform of the sum of two data streams is the sum of the Fourier transform of each data stream.

10.2.2 Symmetric in f and t.

Next, notice that the forward (Equation 10.1) and the inverse transform (Equation 10.2) is symmetric in f and t. The choice of the sign of the exponent ($-$ for the forward transform and $+$ for the inverse transform) is arbitrary. Thus if you know the forward transform for (say) the boxcar function then you also know the inverse transform for the boxcar function.

10.2.3 Real Signals.

Most signals that you encounter are real functions. For example, the time series of a song recorded in audio files (e.g. a Compact Disk) is a series of integers. Let $h(t)$ be a real function and as before let $H(f)$ be its Fourier transform. Let us inspect the $H(f)$ for $f < 0$. We find

$$\begin{aligned} H(-f) &= \int_{-\infty}^{\infty} h(t)\exp(j2\pi ft)dt \\ &= \int_{-\infty}^{\infty} \Big[h(t)\exp(-j2\pi ft)\Big]dt \\ &= H^*(f). \end{aligned} \tag{10.12}$$

Thus the Fourier transform of real function exhibits a certain symmetry: the negative frequency component has the same amplitude and the opposite phase as that of the positive frequency component. The official name of this symmetry is "Hermetian" symmetry.

10.2.4 Sine and Cosine Transforms.

The Fourier Transform of $h(t)$ can be written as the sum of the sine and cosine transform:

$$H(f) = H_c(f) - jH_s(f) \tag{10.13}$$

where

$$\begin{aligned} H_c(f) &\equiv \int_{-\infty}^{\infty} h(t)\cos(2\pi ft)dt \\ H_s(f) &\equiv \int_{-\infty}^{\infty} h(t)\sin(2\pi ft). \end{aligned} \tag{10.14}$$

are the subscripts refer to "cosine" and "sine" transforms.

Now, let $h(t)$ be an even function about $t = 0$ $\big[$i.e. $h(-t) = h(t)\big]$ and $g(t)$ be an odd function $\big[$i.e. $g(-t) = -g(t)\big]$. We find $H(f) = H_c(f)$ and $G(f) = -jG_s(f)$. Any arbitrary function, say $y(t)$ can always be expressed as the sum of an even and odd function:

$$\begin{aligned} y(t) &= y_e(t) + y_o(t) \\ y_e(t) &= \frac{1}{2}[y(t) + y(-t)] \\ y_o(t) &= \frac{1}{2}[y(t) - y(-t)]. \end{aligned} \tag{10.15}$$

With this decomposition we see that $Y_c(f)$ maps to y_e and $Y_s(f)$ to y_o.

10.2.5 Time Scaling.

Say $h(t)$, a function of time, is replayed back at a faster pace[1] $g(t) = h(\alpha t)$ where $\alpha < 1$. Then a change of variable in Equation 10.1 demonstrates

$$G(f) = \alpha^{-1} H(f/\alpha). \tag{10.16}$$

Thus the Fourier transform expands with contraction in time domain.

10.2.6 Modulation and Demodulation.

Modulation (and Demodulation) consist of multiplying a signal by a pure tone to change the center frequency of the signal.

Amplitude modulation (AM) radio works by multiplying the voice signal ($h(t)$) with a pure tone (carrier) in the range 0.3 to 3 MHz. For simplicity assume that the voice signal is also a pure tune in the audio band, $h(t) = A\cos(2\pi ft + \phi)$ and the carrier tone is $l(t) = \cos(2\pi ft)$. Note that $300\,Hz < f < 10,000\,\text{Hz}$ and $0.3 < f_0 < 3\,\text{MHz}$. The wide separation of f_0 and f simplifies further discussion. Let the product be $y(t) = h(t)l(t)$ which we expand to yield more elementary terms:

$$\begin{aligned} y(t) &= A\cos(2\pi ft + \phi)\cos(2\pi f_0 t) \\ &= A/2\cos[2\pi(f_0 - f) - \phi)] + A/2\cos[2\pi(f_0 + f)t + \phi)] \end{aligned} \tag{10.17}$$

Thus the voice signal is now at transformed to radio frequencies in the AM band (see Figure XX). There are two bands: the upper side band ($f < f_0$; USB) and the lower side band ($f > f_0$; LSB). The upper side band has the positive frequencies and the lower side band has the negative frequencies. Either one band has all the information one needs. Demodulation consists of reversing the process. This is what the radio does so that you can listen to $h(t)$ but shifted to the audio range.

10.3 Correlation

Correlation is something that you use all the time. A more common name is "pattern recognition". Consider a radar (or sonar) signal that is launched

[1] In the previous century, music was recorded on grooves etched on flat circular plates – the so-called gramophone. Playback consisted of spinning the the plate and a needle which sensed the groove was the transducer. Smaller plates were spun at 45 rotations per minute (rpm) and larger ones at 33 1/3 rpm. Time scaling could be easily arranged by playing a 33 rpm recorded at the higher 45 rpm speed, yielding hilarious squeaky sounds. Equation 10.16 lay at the heart of the very successful musical group called the *The Chipmunks*.

from a transmitter (or a bat). The signal bounces back and is received by a receiving antenna. The returned signal is a paler version of the transmitted signal. One can show that the most optimal way of determining the radar echo is by cross-correlating the reflected signal with the transmitted signal. LOOK UP BAT STUFF.

Letting $g(t)$ be the return signal and $h(t)$ the transmitted signal we can write the correlation function as

$$c(\tau) \equiv h(\tau) \cdot g(\tau) = \int_{-\infty}^{+\infty} dt\, h(t+\tau) g(t) \tag{10.18}$$

here τ is called the "lag" and in the specific case of a radar echo one computes $c(\tau)$ for positive values of τ. In other applications, it may be desirable to compute negative lags (or both).

The Fourier transform of $c(\tau)$ is given by

$$C(f) = \int_{-\infty}^{\infty} d\tau\, \exp(-j2\pi f\tau) d\tau \int_{-\infty}^{\infty} dt\, h(t+\tau) g(t). \tag{10.19}$$

As before, we assume that for any signals that we experimentally measure (as opposed to signals that exist for infinite time etc) the order of integration can be freely exchanged. Justified thus we proceed and find

$$\begin{aligned} C(f) &= \int_{-\infty}^{\infty} dt\, g(t) \int_{-\infty}^{+\infty} h(t+\tau) \exp(-j2\pi f\tau) d\tau \\ &= \int_{-\infty}^{\infty} dt\, g(t) \int_{-\infty}^{\infty} dt'\, h(t') \exp[-j2\pi f(t'-t)] \\ &= \int_{-\infty}^{\infty} dt\, g(t) \exp(+j2\pi ft) \int_{-\infty}^{\infty} dt'\, h(t') \exp(-j2\pi ft') \\ &= G^*(f) H(f). \end{aligned} \tag{10.20}$$

The order of the correlation does matter (as can be gathered by contemplation or tracing the steps in Equation 10.20):

$$\begin{aligned} c'(\tau) &= g(\tau) \cdot h(\tau) = \int_{-\infty}^{\infty} dt\, g(t+\tau) h(t) \\ C'(f) &= G(f) H^*(f) = C^*(f). \end{aligned} \tag{10.21}$$

Thus the order matters. Furthermore, $C'(f)$ is Hermetian conjugate of $C(f)$. Thus the famous correlation theorem: the Fourier transform of the correlation of two functions is the product of the Fourier transform of the lagged function and the conjugate of the Fourier transform of the unlagged function.

For the specific case where $h(t) = g(t)$ we have

$$a(\tau) = \int_{-\infty}^{\infty} dt\, h(t)h(t+\tau). \tag{10.22}$$

This is called as the auto-correlation function. Note that $a(-\tau) = a(\tau)$. The Fourier transform of the auto-correlation function is

$$\begin{aligned} A(f) &= \int_{-\infty}^{\infty} d\tau \exp(-2j\pi f\tau) \int_{-\infty}^{\infty} dt\, h(t)h(t+\tau) \\ &= H(f)H^*(f). \end{aligned} \tag{10.23}$$

is the Power Spectrum of $h(t)$. By analogy the quantity $C(f)$ is called as the "Cross-power" spectrum.

`Exercise: Correlation Function.`

1. Show by explicit integration that $a(-\tau) = a(\tau)$.
2. What special relations exist if one of the functions [$g(t)$ or $h(t)$] is even; if both are even; and if both are real.
3. Prove

$$\begin{aligned} h(t)\cdot[g(t)+x(t)] &= h(t)\cdot g(t) + h(t)\cdot x(t) \\ [h(t)\cdot g(t)]\cdot x(t) &= h(t)\cdot[g(t)\cdot x(t)] \end{aligned} \tag{10.24}$$

10.4 Convolution

As we will learn later (§XX) the output of a filter can be understood as the convolution of the input signal ($x(t)$) and the impulse response function of the filter ($h(t)$). The convolution function is defined to be

$$v(t) \equiv h(t) * x(t) = \int_{-\infty}^{\infty} d\tau\, h(t-\tau)x(\tau). \tag{10.25}$$

The Fourier transform is given by

$$\begin{aligned} V(f) &= \int_{-\infty}^{\infty} dt\, \exp(-j2\pi ft) \int_{-\infty}^{\infty} dt\, h(t-\tau)x(\tau) \\ &= \int_{-\infty}^{\infty} d\tau\, x(\tau) \int_{-\infty}^{\infty} dt\, h(t-\tau)\exp(-j2\pi ft) \\ &= X(f)H(f). \end{aligned} \tag{10.26}$$

It appears that convolution is independent of the order in which the functions are convolved.

A convolution is obtained by

1. Flip one of the two functions, either the impulse response function or the input time series around $\tau = 0$ (the latter makes sense to me).

2. Delay function by t response function by τ.

Apart from their use in their own right, convolution and correlation are of great use in Fourier "algebra". Convolution in time domain results in multiplication in Fourier domain (Equation 10.26). The auto-correlation provides a convenient way to determine the power spectrum (Equation 10.23). Thanks to symmetry between time and frequency, one could equally state that multiplication in time domain corresponds to convolution in frequency domain and that the Fourier transform of the power spectrum yields the auto-correlation function.

`Exercise: Convolution is distributive and associative.` Prove that

$$\begin{aligned} h(t) * [g(t) + x(t)] &= h(t) * g(t) + h(t) * x(t) \\ [g(t) * h(t)] * p(t) &= g(t) * [h(t) * p(t)] \end{aligned} \tag{10.27}$$

What about commutative?

10.5 Parseval's Theorem

Let us set $g(t) = x(t)^2$Then we know from the convolution theorem that

$$\begin{aligned} \mathcal{F}[x(t) \times x(t)] &= \mathcal{F}(h) * \mathcal{F}(h) \\ \int_{-\infty}^{\infty} dt\, x^2(t) \exp(-2j\pi\sigma t) &= \int_{-\infty}^{\infty} df X(f) X(\sigma - f) \end{aligned} \tag{10.28}$$

This equality is true for any σ and in particular for $\sigma = 0$. Thus we find

$$\begin{aligned} \int_{-\infty}^{\infty} dt\, x^2(t) &= \int_{-\infty}^{\infty} df\, X(f) X(-f) \\ &= \int_{-\infty}^{\infty} df X(f) X^*(f) \end{aligned} \tag{10.29}$$

where assume that $h(t)$ is a real signal (cf Equation XX). This is the celebrated Parseval's theorem. Imagine that $x(t)$ is a the signal and specifically its current. Then $x^2(t)$ is proportional to the power dissipation. Thus the left side hand represents the work done by the signal. We can now understand why $X(f)X^*(f)$ is the "power" spectrum (the power unit frequency at frequency f).

`An insight: In Equation 10.28 the RHS is the Fourier component of the power at frequency` σ. `It is equal to the sum of the cross-power spectrum whose difference in frequency is` σ. `Kind of obvious after the fact. Not seen this mentioned before.`

Chapter 11

Sampling and Interpolation

11.1 Baseband Sampling

Your CD player has the audio signal (in left and right channels) recorded at a standard rate of 44.1×10^3 samples per second with 16-bits sampling (for each of Left and Right channels). The simplest playback consists of obtaining a stable stream of samples (despite you moving around and despite your grimy fingerprints; whence the tremendous redundancy in recording and a big output buffer). The sampling theorem assures us that these numbers have all the information to faithfully reconstruct the audio signal. In this section the most basic explanation for this "sampling".

First we have to understand a very important function: the sampling function or the sha function.

The Fourier transform of a run of δ-functions regularly spaced by Δt is also another run of δ-functions regularly spaced but by $\Delta f = 1/\Delta t$. This function is so important that it has its own name:

$$s(t, \Delta t) = \mathrm{sha}(t, \Delta t) = \sum_{k=-\infty}^{\infty} \delta(t - k\Delta t) \tag{11.1}$$

The Fourier transform is (see Papoulis 1962; see below for a plausible graphical development)

$$S(f) = \frac{1}{\Delta t} \sum_{k=-\infty}^{\infty} \delta(f - k\Delta f). \tag{11.2}$$

`Exercise: Graphical demonstration of the sha Function.` Consider the following function

$$s_n(t, \Delta t) = 1 + 2 \sum_{k=1}^{n} \cos\left(2\pi \frac{kt}{\Delta t}\right) \tag{11.3}$$

Sketch both $s_n(t, \Delta t)$ and $S_n(t)$. This graphical exercise is expected to show why the Fourier transform of a sha function is another sha function but with inverse spacing.

11.2 Fourier Interpolation Theorem

11.3 Bandpass Sampling

11.4 Fourier Series

The traditional approach in physics is to teach Fourier series first and then Fourier transform. This is mathematically a sound approach since issues of the existence of a Fourier transform for a function of infinite duration is not all that obvious. However, here (as in engineering books) we are dealing with signals of finite duration and so we can happily dispense these sorts of subtle mathematical concerns. As will become clear from the discussion below it is much easier to discuss Fourier Series after discussion Fourier Transforms.

11.5 References

The Compact Disc Story by Kees A. Schouhamer Immink, J. Audio Eng. Soc., 46, 458–463 (1998)

Chapter 12

Power Spectrum

12.1 Window Functions

12.2 Filterbanks

12.3 Autocorrelators

Chapter 13

Basic Signal Processing

Signal processing requires that you digitize the signal. Here, we discuss the steps that are undertaken prior to digitization. The most general input signal is a band-pass signal: the power spectrum of a signal is bounded by f_l (lower edge) and f_u (upper edge). A signal with $f_l = 0\,\mathrm{Hz}$ is called as "base-band" signal.

The simplest case is when the analog signal is already a baseband signal. In this case the signal processing consists of using a low-pass filter prior to sampling (Figure 13.1). The sampling rate should be sufficiently high to satisfy the sampling theorem: $f_s \geq 2B$ where B is the bandwidth of the filter. It is difficult to make analog filters with "brick-wall" response (a sharp cutoff) and so there will be some aliasing due to the response of the filter for frequencies greater than B.

Figure 13.1: Digital sampling of a baseband signal.

Simple examples include telephony and CD signals (Table 13.1). Professional (studio) signals, like DAT, are sampled at 48 kHz (many audiophiles claim that a 20-kHz low pass filter gives is inadequate for sampling rate of 44.1 kHz).

However, radio telescopes produce bandpass signals, example "L-band" is a band in the 1.4-GHz (relevant to the study of cosmic H I), "P-band" is a band centered around 0.3 GHz (relevant to observations of the hyper-fine line of deuterium) and so on. For reasons of performance a receiver usually responds to a band (usually fractional bandwidth less than unity).

Signal	f_s (kHz)	Bits	Channels	Band (kHz)
Telephone	8	8	Mono	0.2–3.4
AM Radio	11.025	8	Mono	??
FM Radio	22.050	16	Stereo	??
CD	44.1	16	Stereo	0.020-20
DAT	48	16	Stereo	0.02-20

Table 13.1: From a course by David Marshall, Cardiff School of Computer Science.

13.1 Simple Mixing to Baseband

The simplest bandpass case is when the center frequency (the center of the band) is fixed. In this case, we have two choices: convert the bandpass signal to baseband or make use of the bandpass sampling theorem (§ REF). We will discuss the first possibility.

The approach consists of using is a local oscillator tuned to f_l (Figure 13.2). Multiplying (or "mixing" in the jargon)[1] the radio frequency (RF) signal by the local oscillator (LO) followed by a low-pass filter of the output (the "Intermediate Frequency", IF) does the job (see Figure 13.2).

The usual way to understand the signal processing shown in Figure 13.2 is to consider the RF to be a pure tone, $b(t) = A(f)\cos[2\pi ft + \phi(f)]$. The output of the mixer is then

$$\begin{aligned} o(t) &= A(f)\cos(2\pi ft + \phi(f)) \times \cos(2\pi f_0 t) \\ &= \frac{A(f)}{2}\cos[2\pi(f - f_0)t + \phi(f)] + \frac{A(f)}{2}\cos[2\pi(f + f_0)t + \phi(f)]. \\ &\rightarrow \frac{A(f)}{2}\cos[2\pi(f - f_0)t + \phi(f)] \end{aligned} \tag{13.1}$$

The low pass filter is designed to eliminate signals at the sum frequency leaving only the difference (first term) as the output.

By setting f_{LO} to f_l and adopting the circuitry shown in Figure 13.2 results in a baseband signal with bandwidth $B = f_u - f_l$ (see Figure 13.3). This baseband signal can then be sampled at a rate $f_s > B$.

[1] A mixer is an analog device that multiplies the input signals, a and b. A diode operating at the knee of the i-V curve is used to square a signal. A pair of "balanced" diodes are used to obtain $(a + b)^2$ and $(a - b)^2$. The difference between these two signals is $\propto ab$. Another approach is a *Gilbert* cell. INVESTIGATE.

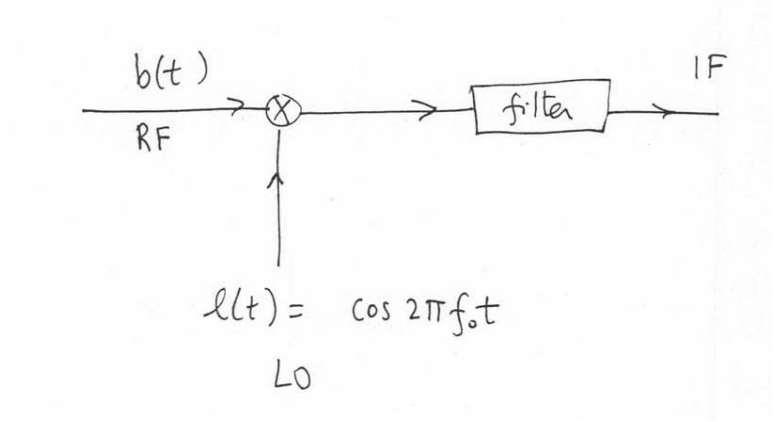

Figure 13.2: Conversion of Radio Frequency to Intermediate Frequency.

Figure 13.3: Simple mixing to baseband.

13.2 Phasor Algebra

In the previous section, we used the cumbersome trignometric relations in arriving Equation 13.1. The "phasor" approach allows one to swiftly compute the filtered output signal. The rules for phasor algebra are as follows:

1. Replace the real signal (always represented by the cosine with appropriate phase added for a sine) by the complex exponential and an amplitude of 1/2. This is the phasor representation.

$$\cos[2\pi ft + \phi(f)] \rightarrow 1/2 \exp[j2\pi ft + j\phi(f)]. \tag{13.2}$$

2. Prior to multiplication of two signals conjugate the phasor of the signal whose frequency is lower of the two.

3. The real part of the product is the low pass frequency output of the mixer.

Applying these rules we obtain the following output:

$$\begin{aligned} o(t) &= \frac{A(f)}{2} \cos[2\pi(f - f_0)t + \phi(f)], f > f_0 \\ &= \frac{A(f)}{2} \cos[2\pi(f_0 - f)t - \phi(f)], f < f_0. \end{aligned} \tag{13.3}$$

However, usually radio telescopes have a range of operating bands (RF central frequencies and bandwidths). In the next two sections we two more general purpose solutions.

13.3 Complex Sampling

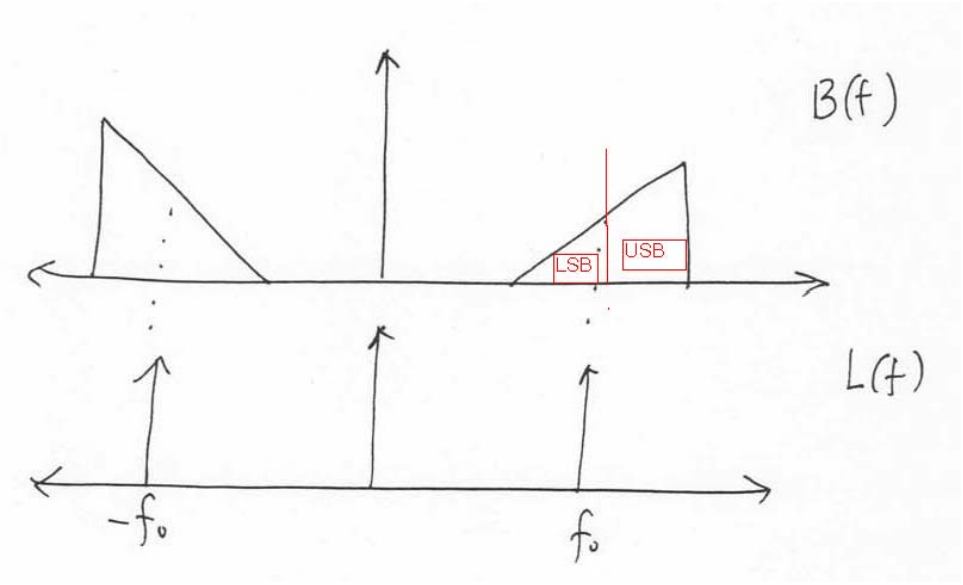

Figure 13.4: Mixing.

In contrast to the previous approach the LO is place at the center of the RF band. For simplicity assume that the input RF signal consisting of two unit cosines frequencies $f = f_0 \pm f'$. Following mixing, these two tones will appear at the same IF frequency. Special processing is needed to separate the lower- and upper-side band tones. The techniques are complex sampling (a software approach) and single-side band separation (a hardware approach). Both techniques exploit the difference in sign that can be seen in Equation **??**.

A bandpass signal, $b(t)$, is split and one signal mixed with a local oscillator with phase $\theta_0 = 0$ ("cosine" LO) and phase $\theta_0 = \pi/2$ ("sine" LO); see Figure 13.5. The corresponding two outputs, $h_c(t)$ and $h_s(t)$, are filtered (to eliminate the sum frequencies). The power spectrum of h_c and h_s would look similar with USB and LSB overlapping. The first channel is referred to as "X", "real", "in-phase" (I), or "cosine". The second channel is "Y", "imaginary", "quadrature" (Q) or "sine". The channels are sampled by analog-to-digital (A/D) units. Each pair of sample is treated as a complex number:

$$z(k) = x(k) + jy(k). \tag{13.4}$$

We now show that the LSBs and USBs of $Z(f)$ are properly separated.

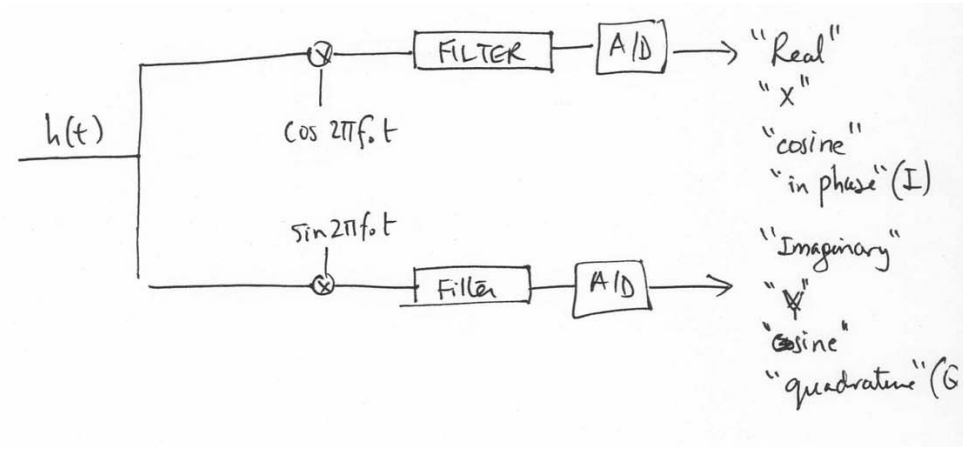

Figure 13.5: Complex Sampler.

Consider a pair of RF unit amplitude phasors with frequency, $f = f_0 \pm f'$:

$$\cos[2\pi(f_0 \pm f')t + \phi(\pm f')]. \tag{13.5}$$

Using Equation **??** we find

$$\begin{aligned} X(f') &= \cos[2\pi f' t + \phi(+f')] \\ Y(f') &= \cos[2\pi f' t + \phi(+f') - \pi/2] = +\sin[2\pi f' t + \phi(+f')] \\ X(-f') &= \cos[2\pi f' t - \phi(-f')] = \cos[-2\pi f' t + \phi(-f')] \\ Y(-f') &= \cos[2\pi f' t - \phi(-f') + \pi/2] = -\sin[-2\pi f' t + \phi(-f')] \end{aligned} \tag{13.6}$$

Thus

$$\begin{aligned} Z(+f') &= X(+f') + jY(-f') = \exp[+j2\pi f' t + j\phi(+f')] \\ Z(-f') &= X(-f') + jY(-f') = \exp[-j2\pi f' t + j\phi(-f')]. \end{aligned} \tag{13.7}$$

Thus the USB signal is now an analytical signal with $f' > 0$ and LSB is an analytical signal with $f' < 0$ (see Figure ZZ).

13.4 Single Side-Band Mixer

The band-pass signal is, as before, mixed with local oscillators and converted to two baseband signals, $h_s(t)$ and $h_c(t)$. A $\pi/2$ phase shift is applied to the former signal to produce $h'_s(t)$. The sum and difference of h'_s and h_c result in USB and LSB respectively (see Figure 13.6).

The "$\pi/2$" phase shifter is somewhat akin to anti-reflection coatings. The purpose of the shifter is to introduce $\pi/2$ shift over the entire bandwidth.

A graphic proof of how the SSB mixer works is presented in Figure 13.7.

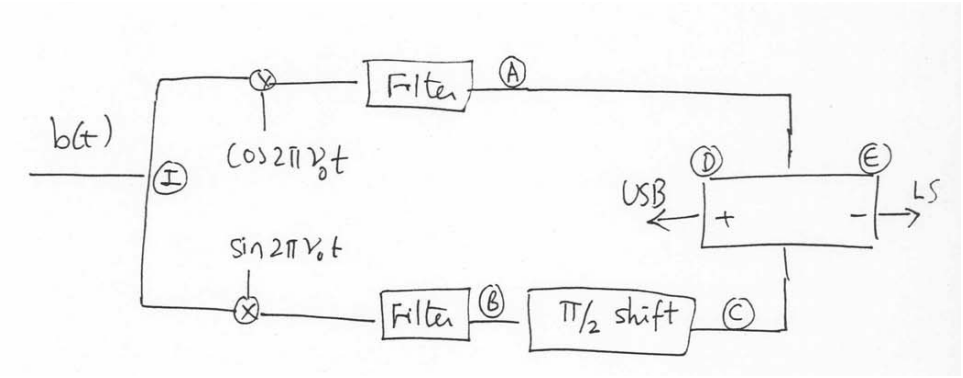

Figure 13.6: Single Side-Band Mixer.

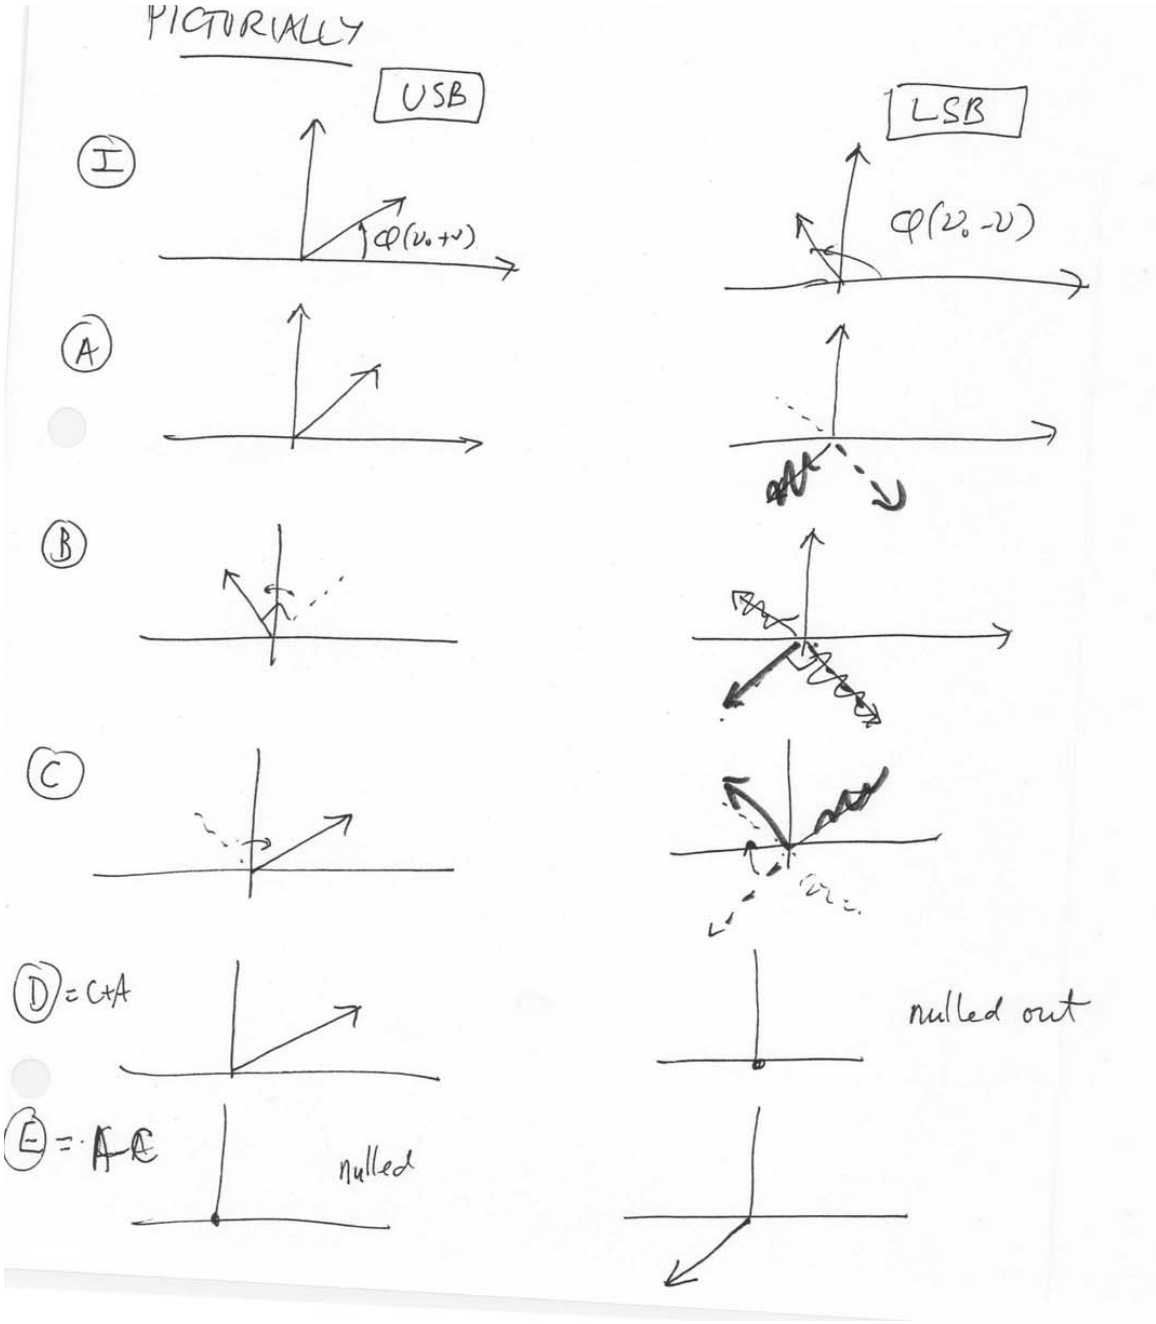

Figure 13.7: SSB Mixer - Graphic proof.

Chapter 14

Discrete Fourier Transform (Applications)

The practical realization of a Fourier Transform is the Discrete Fourier Transform (DFT). The simplest case for a DFT is a band limited (Bandwidth, B) signal sampled uniformly, $h(k\Delta t)$ where $k = [0, 1, ..., n-1]$ and $\Delta t = 1/f_s$ where $f_s = 2B$ is the sampling frequency. The duration of the sampled signal is thus $n = 2B\Delta t$. The symmetry between time and frequency (see previous Chapter) requires that the Fourier transform be also computed on an even grid, say spacing, Δf. Thus the DFT is given by

$$H(l\Delta f) = \sum_{k=0}^{n-1} h(k) \exp(-j2\pi kl\Delta t\Delta f). \tag{14.1}$$

What is the minimum frequency resolution at which you should compute the Fourier Transform? There are several ways to answer this question. The first answer is based on the time-frequency symmetry of Fourier transforms. Thus if the critical sampling interval is the inverse of the total bandwidth (positive and negative frequencies), $\Delta t = (2B)^{-1}$, then Δf is the inverse of the total duration, $T = n\Delta t$ or $\Delta f = f_s/n$.

Another approach is to consider a band limited signal, band width B and sampled at the critical rate, $f_s = 2B$. Since the input time series is sampled at a critical rate each successive element is independent of other elements. Thus over a period of T we have $n = 2B\Delta t$ independent elements. The Fourier Transform is simply a linear combination of these independent elements and can at most also be made up of N independent elements. Since these elements include both positive and negative frequencies the frequency spacing must be $\Delta f = f_s/n$.

Thus with the above choice for Δf the DFT equation is

$$H(l) = \sum_{k=0}^{n-1} h(k) \exp(-2\pi j \frac{lk}{n}). \tag{14.2}$$

Certain channels are worthy of note. The "dc" bin (or "l=0") and the "Nyquist" bin[1] (or $l = n/2$) are both purely:

$$H(0) = \sum_{k=0}^{n-1} h(k), \qquad H(n/2) = \sum_{k=0}^{n-1} (-1)^k h(k) \tag{14.3}$$

with the former being the sum of all elements and the latter being the sum of the differences of successive elements. A signal which is not zero mean will produce a strong spike at $l = 0$. Suppression of this peak is done by subtracting the mean value of the input time series from each element of the time series. Next, note that $H(n - L) = H(-L)$. The negative frequencies are thus found between $k = n/2 + 1$ and $k = n - 1$. Next, the sequence, H is periodic with period P. For instance, $H(n) = H(0)$.

The inverse DFT is essentially the same as the forward DFT except that the $-j$ in the exponential (Equation 14.2) is replaced by $+j$ and a scaling factor of $1/n$ is applied[2]. The scaling factor ensures that a DFT followed by an inverse DFT yields the initial function. As an example consider the the function $h(k) = 0$ except for $h(1) = 1$. In this case the DFT is $H(l) = 1$ for $l = 1, ..., n$. The inverse DFT, if it incorporates the scaling factor of $1/n$, will reproduce the input function, h.

14.1 The Fast Fourier Transform

The calculation of a specific frequency component requires (for each of cosine and sine transforms) n trignometric computations, n multiplications and n additions. Now if one wants to compute, say m frequeny components, then the number of computations is $4mn$; henceforth, we assume that the trignometric computations are "pre computed".

As noted above, a full exploration of the frequency requires that $m \approx n$ and thus the DFT would require $4mn$ or $O(n^2)$ multiply or add operaitons. Fourier Transforms are used for signal processing, convolution of images or

[1]The critical sampling rate is sometimes referred to as the Nyquist rate; the term Nyquist as used here is common, at least amongst astronomers.

[2]There are other choices including for example a scaling factor of $\sqrt{1/n}$ for both the forward and inverse transforms.

time series (via the convolution theorem), searching for periodicities and lossy compression of data. Many of these involve sequences with large value of n and $O(n^2)$ can become quite prohibitive.

A Fast Fourier Transform (FFT) evaluates equation 14.2 in only $O(n \log(n)$ multiply and add operations. The algorithm was apparently known to Gauss and re-discovered in the sixties by Cooley & Tukey in the sixties. The approximate number of computation steps needed for an FFT is $4n \log_2(n)$. `[Explain how it works]`.

As an algorithm the FFT has had a transformative impact on signal processing and certainly without it the development of most of modern radio astronomy (aperture synthesis, spectrometers, pulsar research, low frequency astronomy) would have been held back. The FFT is a clear instance in which development of a methodology has had a profound effect, comparable to the discovery of a major phenomenon.

In this chapter we consider two specific uses of the FFT. The first is the generation of a time series which follow Gaussian statistics and have a specified power spectrum. The second is the detection of a sine wave buried in noise.

14.2 Statistics of DFT

We re-express the DFT of a uniformly sampled series as the sum of the cosine and sine components, $H(l) = C(l) + jS(l)$ where

$$C(l) = \sum_{k=0}^{n-1} h(l) \cos(2\pi lk/n), \qquad S(l) = \sum_{k=0}^{n-1} h(l) \sin(2\pi lk/n). \quad (14.4)$$

We will only specify the first two moments of h:

$$\langle h \rangle = 0 \qquad \mathrm{cov}(h(l), h(l')) = \sigma^2 \delta_{ll'} \quad (14.5)$$

where the covariance function is defined as $\mathrm{cov}(x, y) = \langle xy \rangle - \langle x \rangle \langle y \rangle$. The covariance of $C(l)$ and $S(l')$:

$$\begin{aligned} \mathrm{cov}[C(l), S(l')] &= \Big\langle \sum_{k=0}^{n-1} \sum_{k'=0}^{n-1} h(l) h(l') \cos(2\pi lk/n) \sin(2\pi l'k'/n) \Big\rangle \\ &= \frac{1}{2} \sum_{k=0}^{n-1} \sum_{k'=0}^{n-1} \langle h(l) h(l') \rangle \big[\sin(2\pi(l'k' + lk)/n) + \sin(2\pi(l'k' - lk)/n) \big] \\ &= 0. \end{aligned} \quad (14.6)$$

It is pretty straightforward to show that

$$\mathrm{cov}[C(l), C(l')] = (\sigma_H^2/2)\delta_{lk}, \qquad \mathrm{cov}[S(l), S(l')] = (\sigma_H^2/2)\delta_{lk} \quad (14.7)$$

where $\sigma_H^2 = n\sigma^2$. Note that these covariances (Equations 14.6, 14.7) go to zero only for the specific case when $\Delta f = 1/T$ (more on this later). The complex representation, $Z(l) = C(l) + jS(l)$ are circular Gaussian variates. As will be seen later (Chapter InternalRef) the computation of higher order moments is relatively easy for such variates.

So far we have only specified the first two moments of h. For large values of n we can argue that $C(l)$ and $S(l)$ will follow Gaussian statistics with zero mean and variance, $n/2\sigma^2$ – provided that h follows the conditions specified by Equation 14.5 . Next, the virtue of the lack of covariance between $C(l)$ and $S(l')$ assures that the sine and cosine components are themselves independent variates. Thus, at this point we have fully specified the probability distributions of each element of the power spectrum.

Many analyses involve the power spectrum, $P(l) = C(l)^2 + S(l)^2$. Given Equation 14.6 we can easily show that $P(l)$ follows an exponential or χ^2 distribution (see InternalRef) with mean value of $\sigma_H^2 = n\sigma^2$. Given that the cosine and sine components for each frequency channel are independent Gaussian variates it also follows that $P(l)$ are also independent from each other. This means in particular that the statistics of a sum of any subset of power spectra bins, say $P(l_1) + P(l_2) + ... + P(l_p)$ follows χ^2 distribution with $2p$ degrees of freedom. This result is useful in pulsar astronomy ("harmonic summing"; see InternalREF).

14.3 Generating Gaussian Noise with a Given Power Spectrum

Frequently one needs to generate a time series or a spatial series (map, image) with a given power spectrum. For instance, you may wish to simulate the detection of a faint feature in a radio spectrum which already has bright lines and a sloping continuum; or you may wish to simulate the formation of speckles arising from a Kolmogorov spectrum (either in the decimeter band or at optical wavelengths); or the intensity distribution of the CMB. Another application is to generate band limited spectrum so that you can study the effect of digitizing the data.

A time series generated by $h(k)$ which arise from $N(\mu = 0, \sigma^2 = 1)$ will have a power spectrum which is flat (because each member is independent of any other member of the sequence; thus the auto-correlation is a delta

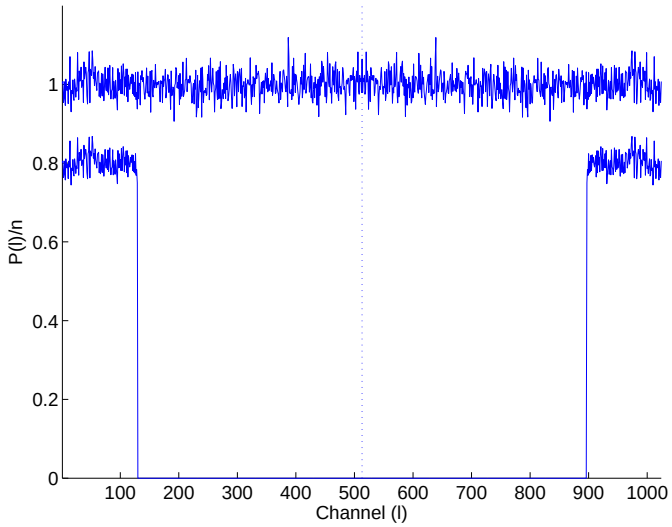

Figure 14.1: (Top) Mean power spectrum of a sequence $h(k)$ with $k = 1, 2, ..., n$ with $n = 1024$. The $h(k)$ were Gaussian variates drawn from $N(0, 1)$ and $m = 1024$ such runs. Each of the power spectrum was scaled down by n so that the average value across the power spectrum would be unity. The horizontal axis is the bin index with the Nyquist bin ($n = 513$) marked by a vertical dotted line. The bins to the right of the Nyquist bin are negative frequency components. (Bottom) The mean power spectrum of $y(k)$ where $y(k)$ is the low pass filtered version of h (see text). The following low pass filter was applied, $B(l) = 1$ for $l < nf$ and zero otherwise; here, $f = 1/8$. The mean power spectrum has been scaled down by an additional factor of 0.8 (for the purpose of plotting) so that the filtered version does not overlap the original power spectrum.

function and from this the assertion of a flat spectrum follows). Traditionally, noise with a flat spectrum is called as "White Noise".

The following prescription generates Gaussian noise with a specific power spectrum, say $B(f)$.

1. Generate $h(k)$ of desired length n.

2. Fourier transform to obtain $H(l)$.

3. Filter the signal per specification, $Y(l) = H(l)B(l)$.

4. Inverse Fourier transform $Y(l)$ to obtain $y(l)$. By construction $y(l)$ will be Gaussian noise with a power spectrum specified by $B(f)$.

The basic idea described above is quite simple. The usual difficulty lies in implementation and the usual mistake is not correctly treating the negative frequencies. Hermetian symmetry must be preserved in step (3) above and this should be borne in mind when treating the negative frequency components. A simple check is to verify that y is purely real.

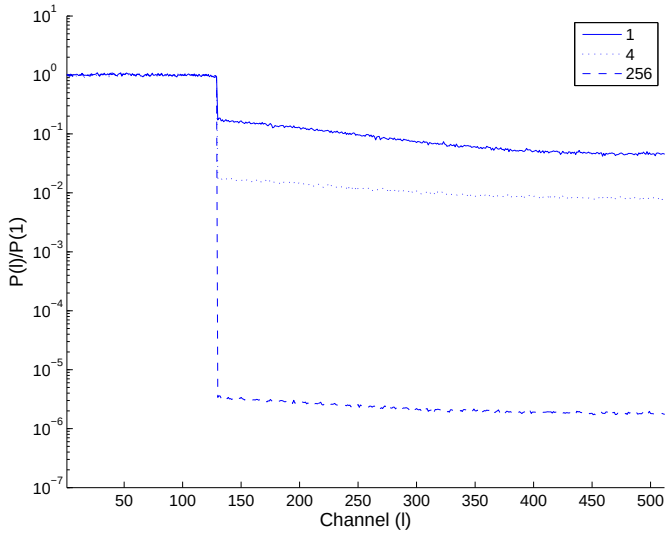

Figure 14.2: The power spectrum of quantized white noise source. The filtered white noise, $y(k)$, is generated in a block of $n = 1024$ and in a manner described in the caption to Figure 14.1. The quantization is done as follows: $z = \text{nint}(yS)$ where $S = [1, 4, 256]$ and corresponds (in spirit) to 1-bit, 2-bit and 16-bit digitization. The power spectrum for $m = 1024$ such blocks is averaged and normalized by the average of the dc channel. Only channels with positive frequencies ($l = 1, ..., 512$) are shown.

An example for the need for filtered noise is shown in Figure 14.2. The filtered time series described in Figure 14.1, $y(l)$, is quantized as follows: $z = \text{nint}(yS)$ where $S = 1, 4, 256$ and nint is the nearest integer function. Quantization produces a long tail of essentially noise. Thus quantization takes a band-limited signal and makes it not band limited. The signal-to-noise ratio of the quantized signal is worse than that of the unquantized signal for two reasons. First is the loss of power (the power in the tail is at the expense of the principal signal) and the second is increased noise from the tail (as it folds back into the main band).

14.4 Spectral Resolution & Fine Binning

In astronomy a major application of Fourier Transforms is to search periodic signals. Examples abound: searches for pulsars at radio wavelengths, searches for pulsations at optical and X-ray wavelengths, searches for orbital period via radial velocity and photometry and so on. `Other examples`. Here, we will explore this important application in some detail. We will start by assuming that the periodic signal is the simplest possible, namely, a sine wave with frequency f_0 (and period $P_0 = 1/f_0$).

The DFT as implemented by a standard FFT package produces the following response to our putative period signal of unit amplitude:

$$H(l) = \frac{1}{2}\text{sinc}\Big(\frac{\pi(l - l_0)}{n}\Big) \tag{14.8}$$

where $l_0 = f_0/\Delta f$.

`Exemplars.` Need astronomical exemplars.

`Exercise FFT-1:` Most package FFTs accept a sequence of real measurements and return complex output (both positive and negative frequencies). However, given the Hemetian symmetry the computation of (say) the negative frequencies is unnecessary. Bearing this in mind devise an algorithm to undertake an efficient FFT of two real sequences of the same length.

`Exercise FFT-2:` As discussed in the text, zero padding allows for computation of Fourier transform on a finer grid of frequencies. What is the effect of "pre-padding" instead of "post-padding"?

`Exercise FFT-3:` Say that you have an FFT "Machine", a device that can undertake rapid n-point transforms (but not more than n) attached to your

machine as an accessory. How would you go about making a $2n$ point transform on your computer?

`Exercise FFT-4:` Reproduce Figure 14.2. This means generate $h(k)$, for $k = 1, 2, ..., n$, Gaussian variates drawn from $N(0, 1)$. Low pass filter this time series as noted in the caption of Figure 14.1. Next digitize the filter series to 1-bit, 2-bits and 16-bits. [The digitization is in units of σ of the filtered series]. Obtain the power spectrum. Repeat the process say $m = 1024$ and plot the mean power spectrum of the digitized data. [For every step pay attention to the variance of the signal; you need to understand what determines the variance].

`Exercise FFT-5:` Redo the above exercise except that the filter function is say $B(f) \propto f^{\alpha/2}$. Undertake the exercise for $n = 1024$ and plot a single run of $y(l)$. Compute the variance of $y(l)$ and (preferably predict) explain the value. Compare the character of the time series of white noise and the red noise with $\alpha = -1$ (which means that the power spectrum has a power law slope of 2α).

`Exercise FFT-6:` [Interesting problem but not sure if there is a solution]. Can one use the FFT to compute an octave of frequency components (at a finer grid)?

`Literature:` The Wikipedia ("Fast Fourier Transform") provides a good starting point for a summary of the FFT algorithm. See also Numerical Recipes.

`Solutions:`

`Exercise FFT-1:`

`Exercise FFT-2:`

`Exercise FFT-3:`

`Exercise FFT-4:`
See text.

`Exercise FFT-5:`

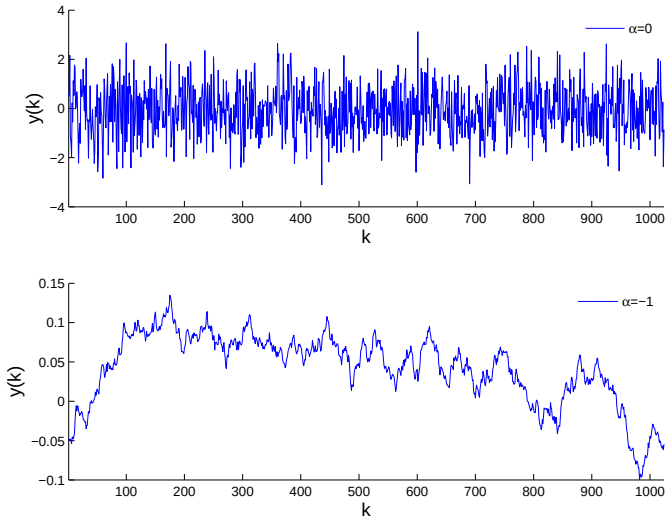

Figure 14.3: (Top) Time series generated from $k = 1, 2, ..., n = 1024$ measurements drawn from $N(0, 1)$. (Bottom) The above series filtered to $l^{2\alpha}$ where $\alpha = -1$ and l is the channel number.

Chapter 15

Filters

15.1 FIR

15.2 IIR

15.3 Polyphase

Chapter 16

Kalman Filtering

Chapter 17

Structure Function & Allan Variance

Chapter 18

Semi-classical Theory of Detection

Chapter 19

Quantum Theory of Detection

Chapter 20

Amplitude Interferometry

20.0.1 Radio

20.0.2 Optical

Chapter 21

Intensity Interferometry

21.0.1 Radio

21.0.2 Optical

Chapter 22

FringePhasor/FringePhasor

Acknowledgements

First and foremost I would like to thank the taxpayers of India who gave me the finest undegraduate education, the taxpayers of the state of California who gave me the finest graduate education and the taxpayers of the United States of America who gave me opportunties unparalleled anywhere else in this world to develop as a scientist and contribute to humanity's growth of knowledge.

Students and post-doctoral fellows who read and corrected this book:

Colleagues who shared their courseware and suggested topics:

Bibliography

This is a pedagogical book and as such clarity trumps over most issues. In particular, I relegated the references to this chapter.

I have profited from the following textbooks:

1. *The Fast Fourier Transform* and its Application by E. O. Brigham (Prentice Hall). A very clear and practical exposition of Fourier Transforms.

2. *Probability and Statistics* by

3. *An Introduction to the Theory of Random Signals and Noise* by W. B. Davenport & W. L. Root (IEEE Press).

4. *Data Analysis: A Bayesian Tutorial* by D. Sivia & J. Skilling (Oxford University Press).

I would like to acknowledge Eric Weisstein's Mathworld website (`http://mathworld.wolfram.com`
It is a very reliable reference source. What surprised me was the vast amount of excellent material that was available on the `wikipedia`. I intend to donate a 10% of the earnings from this book to `wikipedia`.