# Notes on Probability Theory 

## ASTP 611-01: Statistical Methods for Astrophysics*

Fall Semester 2017

## Contents

1 Fundamentals of Probability ..... 2
1.1 Probability and Logic ..... 2
1.2 Bayes's Theorem ..... 3
1.3 Bayesian and Frequentist Inference ..... 4
2 Probability Distributions ..... 5
2.1 Expectation Values ..... 5
2.2 Change of Variables in Probability Distributions ..... 6
2.2.1 Example: Inclination ..... 7
2.3 Multivariate Distributions ..... 7
2.3.1 Expectation Values, Variance, Covariance and Correlation ..... 8
2.3.2 Change of Variables ..... 8
2.3.3 General Formula ..... 11
3 Probability Distributions ..... 11
3.1 Some Specific Probability Distributions ..... 11
3.1.1 The Binomial Distribution (discrete) ..... 11
3.1.2 Other Related Distributions ..... 13
3.1.3 The Poisson Distribution (discrete) ..... 13
3.1.4 The Exponential Distribution (continuous) ..... 15

[^0]| 3.1 .5 | The Gamma Distribution (continuous) | 16 |
| :---: | :---: | :---: |
| 3.1.6 | The Gaussian (aka Normal) Distribution |  |
|  | (continuous) | 18 |
| 3.1.7 | The Chi-Square Distribution (continuous) | 20 |
| 3.1.8 | Summary of Properties of Gamma Distri- |  |
|  | bution | 21 |

4 Sums of Random Variables ..... 21
4.1 Mean and Variance ..... 21
4.2 IID Random Variables (Random Samples) ..... 23
4.3 PDF of a Sum of Random Variables ..... 23
4.4 The Central Limit Theorem ..... 24
5 Multivariate Normal Distribution ..... 25
5.1 Linear Algebra: Reminders and Notation ..... 25
5.2 Special Case: Independent Gaussian Random ..... 275.3 Multivariate Distributions5.4 General Multivariate Normal Distribution5.5 Sample Mean and Sample Variance30
(Student's Theorem) ..... 32
$\frac{5.6 \text { Reduced Chi-Squared . . . . . . . . . }}{5 \text { Minimizing } \chi^{2} \text { over parameters }}$ ..... 36
5.6.1 Minimizing $\chi^{2}$ over parameters ..... 36

## Tuesday, September 12, 2017

## 1 Fundamentals of Probability

### 1.1 Probability and Logic

See Gregory, Chapters 1 and 2; see also section 1 of http://ccrg.rit. edu/~whelan/ courses/2017_1sp_ STAT_489/notes_parameters.pdf

There are numerous interpretations of probability, but one which applies well to observational science is that of an extended logic. Let $A$ be a proposition which could be either true or false, e.g., "The orbital period of Mars is between 686 and 687 days," "The student in question receives an A in stat methods," or "My detector will collect 427 photons in the next two hours." We may know, given the information at hand, that $A$ is definitely true or definitely false, or we may be uncertain about the answer, either because our knowledge of the situation is incomplete, or because it refers to the outcome of an experiment with a random element, which has not occurred yet. The probability of the proposition $A$ (which we also call an "event") is a number between 0 and 1 which quantifies our degree of certainty, given the information at hand. We write this as $P(A \mid I)$, where $I$ represents some state of knowledge, to emphasize that the probability we assign always depends on the information we have, the assumption that a model is correct, etc. If $A$ is definitely true, in the context of $I$, then $P(A \mid I)=1$. If it's definitely false, $P(A \mid I)=0$.

If $A$ represents the outcome of an experiment which we could somehow arrange to repeat under identical circumstances, then $P(A \mid I)$ will be approximately equal to the long-term frequency of the event $A$. I.e., if we do some large number $N$ of repetitions of the experiment, at the beginning of which we recreate the situation described by $I$, the approximate number of experiments in which $A$ will turn out to be true is $N \times P(A \mid I)$. In the
classical or "frequentist" approach to statistics, this is the only sort of event to which we're allowed to assign a probability, but in the more general "Bayesian" framework we are free to assign probabilities to any logical proposition.

Several basic operations can be used to combine logical propositions:

- Negation. $\bar{A}$ is true if $A$ is false, and vice-versa. In words, we can think of $\bar{A}$ as "not $A$ ". (Other notations include $A^{\prime}$ and $\neg A$.)
- Intersection. $A, B$ is true if $A$ and $B$ are both true. In words, this is " $A$ and $B$ ". (Other notations include $A B$, $A \cap B$ and $A \wedge B$.) The advantage of the comma is that $P(A, B \mid I)$ is the probability that both $A$ and $B$ are true, given $I$.
- Union. $A+B$ is true if either $A$ or $B$ (or both) is true. In words, this is " $A$ or $B$ ". (Other notations include $A \cup B$ and $A \vee B$.) Note the unfortunate aspect of this notation that + is to be read as "or" rather than "and".

The relationships between logical propositions are easier to see with so-called truth tables, where you make a list of all of the possible combinations of truth and falsehood for different propositions:

| $A$ | $B$ | $A, B$ | $A, \bar{B}$ | $\bar{A}, B$ | $\bar{A}, \bar{B}$ | $A+B$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| T | T | T | F | F | F | T |
| T | F | F | T | F | F | T |
| F | T | F | F | T | F | T |
| F | F | F | F | F | T | F |

Two propositions are considered to be equivalent if one is true whenever the other is true and false whenever the other is false. For instance, we can show that $A=A, B+A, \bar{B}$ with the following truth table:

| $A$ | $B$ | $A, B$ | $A, \bar{B}$ | $A, B+A, \bar{B}$ |
| :---: | :---: | :---: | :---: | :---: |
| T | T | T | F | T |
| T | F | F | T | T |
| F | T | F | F | F |
| F | F | F | F | F |

There are basic rules of probability corresponding to these logical operations:

- $P(A \mid I)+P(\bar{A} \mid I)=1$
- The product rule: $P(A, B \mid I)=P(A \mid B, I) P(B \mid I)$
- The sum rule: if $A$ and $B$ are mutually exclusive, i.e., if $P(A, B \mid I)=0$, then $P(A+B \mid I)=P(A \mid I)+P(B \mid I)$.

Note that in this approach, where all probabilities are conditional, the product rule is really what's fundamental. Classical approaches to probability instead define the conditional probability as $P(A \mid B)=\frac{P(A, B)}{P(B)}$, and therefore only entertain consideration of the conditional probability $P(A \mid B)$ if $B$ is not only something to which they're allowed assign a probability, but for which that probability is nonzero.

It is possible to show that any extension of logic which allows for a quantitative characterization of the plausibility of logical statements, obeying a few desiderata, is necessarily equivalent to probability theory. I encourage you to read Chapter Two of Gregory (which is based on a presentation in Jaynes) for details.

### 1.2 Bayes's Theorem

Because the logical "and" and "or" operations are symmetrical, i.e., $A, B$ is equivalent to $B, A$ and $A+B$ is equivalent to $B+A$, we can write the product rule in two different ways:

$$
\begin{equation*}
P(A, B \mid I)=P(A \mid B, I) P(B \mid I)=P(B \mid A, I) P(A \mid I) \tag{1.1}
\end{equation*}
$$

this can be rearranged into Bayes's Theorem, which says that

$$
\begin{equation*}
P(A \mid B, I)=\frac{P(B \mid A, I) P(A \mid I)}{P(B \mid I)} \tag{1.2}
\end{equation*}
$$

which is incredibly useful when you naturally know $P(B \mid A, I)$ but would like to know $P(A \mid B, I)$. For instance, suppose $A$ refers to "I have terrible-disease-of-the-year (TDY)", $B$ refers to "I test positive for TDY", and $I$ represents the information that I had no extra risk factors or symptoms for TDY but was routinely tested, $0.1 \%$ of people in such a group have TDY, the test has a $2 \%$ false positive rate ( $2 \%$ of people without TDY will test positive for it) and a $1 \%$ false negative rate ( $1 \%$ of people with TDY will test negative for it). This information tells us that:

- $P(A \mid I)=0.001$ so $P(\bar{A} \mid I)=0.999$.
- $P(\bar{B} \mid A, I)=0.01$ so $P(B \mid A, I)=0.99$.
- $P(B \mid \bar{A}, I)=0.02$ so $P(\bar{B} \mid \bar{A}, I)=0.98$.

Additionally, since $B=B, A+B, \bar{A}$,

$$
\begin{align*}
P(B \mid I) & =P(B, A \mid I)+P(B, \bar{A} \mid I) \\
& =P(B \mid A, I) P(A \mid I)+P(B \mid \bar{A}, I) P(\bar{A} \mid I) \\
& =0.99 \times 0.001+0.02 \times 0.999=0.00099+0.01998 \\
& =0.02097 \tag{1.3}
\end{align*}
$$

We can then use Bayes's theorem to show that

$$
\begin{equation*}
P(A \mid B, I)=\frac{0.00099}{0.02097} \approx 0.04721 \tag{1.4}
\end{equation*}
$$

I.e., if I test positive for TDY, I have about a $4.7 \%$ chance of actually having the disease. This is a lot less than $P(B \mid A, I)$, which is $99 \%$ !

In the context of observational science, Bayes's theorem is most commonly applied to a situation where $H$ is a hypothesis which I'd like to evaluate and $D$ is a particular set of data I've collected. It's usually straightforward to work out $P(D \mid H, I)$, the probability of observing a particular set of data values given a model, but I generally want to answer the question, what is my degree of belief in the hypothesis $H$ after the observation. The answer, according to Bayes's Theorem, is

$$
\begin{equation*}
P(H \mid D, I)=\frac{P(D \mid H, I) P(H \mid I)}{P(D \mid I)} \tag{1.5}
\end{equation*}
$$

As a bit of terminology:

- $P(H \mid I)$ is called the prior probability of hypothesis $H$.
- $P(H \mid D, I)$ is called the posterior probability of $H$. (This is a bit of a misnomer since it implies they are chronological; really we're considering the probabilities we assign with and without the knowledge of $D$.)
- $P(D \mid H, I)$ is called the sampling distribution if we view it as a function of $D$ or the likelihood if we view it as a function of $H$.
- $P(D \mid I)$ is called the evidence, but as we'll see, it's usually considered a normalization constant, since it doesn't depend on $H$.


### 1.3 Bayesian and Frequentist Inference

Now a brief preview of how probabilities can be used to make statements about different hypotheses in light of observational data. We have a hypothesis $H$ (e.g., a model, with certain values for its parameters) and it makes predictions about data $D$, i.e., the results of an observation or experiment. If these predictions had no randomness or uncertainty, i.e., $H$ always led to a particular $D$, we could know for sure that e.g., $\bar{D}$ implied
$\bar{H}$. Instead, we need to take the probability $P(D \mid H, I)$ for the outcome of an experiment given a hypothesis, and use it to say something about the hypothesis.

The Bayesian approach uses Bayes's theorem (1.5) to define $P(H \mid D, I)$ and evaluate it for the actual data observation that we made. In particular if we want to compare two hypotheses $H_{1}$ and $H_{2}$, we can consider the ratid ${ }^{1}$

$$
\begin{equation*}
\frac{P\left(H_{1} \mid D, I\right)}{P\left(H_{2} \mid D, I\right)}=\frac{P\left(D \mid H_{1}, I\right)}{P\left(D \mid H_{2}, I\right)} \frac{P\left(H_{1} \mid I\right)}{P\left(H_{2} \mid I\right)} \tag{1.6}
\end{equation*}
$$

Now, different people with different background information $I$ would likely assign different values to the prior odds ratio $\frac{P\left(H_{1} \mid I\right)}{P\left(H_{2} \mid I\right)}$, but much less information is used to calculate the likelihood ratio $\frac{P\left(D \mid H_{1}, I\right)}{P\left(D \mid H_{2}, I\right)}$. When we think of it as the factor by which we multiply the prior odds ratio $\frac{P\left(H_{1} \mid I\right)}{P\left(H_{2} \mid I\right)}$ to get the posterior odds ratio $\frac{P\left(H_{1} \mid D, I\right)}{P\left(H_{2} \mid D, I\right)}$, we call it the Bayes Factor.

In the frequentist approach, we're not allowed to assign probabilities to different hypotheses, so all we can consider is how the likelihood ${ }^{2} P(D \mid H, I)$ looks for different choices of the hypothesis $H$ and evaluate it for the particular data $D$ observed. The problem is that we can't address the question about the relative likelihood of different hypotheses and instead have to compare the data observed to different data that could have been observed for each hypothesis. In exchange for not having to say anything about $P(H \mid I)$ we can't actually say anything about $P(H \mid D, I)$.

[^1]
## Thursday, September 14, 2017

## 2 Probability Distributions

In what follows, we will often suppress the explicit mention of the background information $I$ on which all of our probabilities are conditional. The logical propositions to which we often assign probabilities involve the values of some random or otherwise unknown quantities. So for example, $N_{\text {counts }}=37$ or $70 \mathrm{~km} / \mathrm{s} / \mathrm{Mpc}<H_{0}<75 \mathrm{~km} / \mathrm{s} / \mathrm{Mpc}$. Sometimes the notation gets a bit confused between a quantity and its value, and you'll see things like $X$ for a "random variable" and $x$ for a value it can take on. You'd like to be able to specify the probability that $X=x$, as a function of $x$. In practice, this is slightly complicated by whether we think of $X$ as taking on only discrete values, or if it can take on any value in a continuous range.

If $X$ is discrete, we can talk about its probability mass function $p_{X}(x)=P(X=x)$. This is often just written $p(x)$ or $P(x)$. For instance, if $X$ is the number of events in a particular interval from a stationary process in which the events are independent of one another, and the average number of events expected in the interval given the long-term event rate is $\mu$ it is described by the Poisson distribution

$$
p(x)=P(X=x)= \begin{cases}\mu^{x} & e^{-\mu}  \tag{2.1}\\ 0 & x=0,1,2, \ldots \\ 0 & \text { otherwise }\end{cases}
$$

However, it often happens that $X$ is continuous, so that it is vanishingly unlikely that it takes on one specific value. For instance, the height of a randomly chosen person will not be exactly 175 cm . If you measure it to more significant figures, it will turn out to be 175.25 cm or 175.24732 cm etc. So instead we want to talk about the probability for $X$ to be in a small
interval, which we call the probability density function

$$
\begin{equation*}
f(x)=\lim _{d x \rightarrow 0} \frac{P(x<X<x+d x)}{d x} \tag{2.2}
\end{equation*}
$$

so that

$$
\begin{equation*}
P(a<X<b)=\int_{a}^{b} f(x) d x \tag{2.3}
\end{equation*}
$$

The pdf might be called $\operatorname{pdf}(x)$ or even $P(x)$. A useful notation for the pdf is $\frac{d P}{d x}$, which tends to make the impact of changes of variables more obvious. In the end, it's a bit hopeless to try to stick to one letter, since you might want to talk about the joint probability distribution associated with some discrete and some continuous random variables. To give a concrete example, a common probability distribution is the Gaussian distribution with parameters $\mu$ and $\sigma$, which has pdf

$$
\begin{equation*}
f(x)=\frac{1}{\sigma \sqrt{2 \pi}} e^{-(x-\mu)^{2} / 2 \sigma^{2}}, \quad-\infty<x<\infty \tag{2.4}
\end{equation*}
$$

### 2.1 Expectation Values

In either case, you can define an operation known as the expectation value

$$
E[g(X)]= \begin{cases}\sum_{x} g(x) p(x) & X \text { discrete }  \tag{2.5}\\ \int_{-\infty}^{\infty} g(x) f(x) d x & X \text { continuous }\end{cases}
$$

with the mean $\mu_{X}=E[X]$ as a special case, and also the variance

$$
\begin{equation*}
\operatorname{Var}(X)=E\left[\left(X-\mu_{X}\right)^{2}\right] \tag{2.6}
\end{equation*}
$$

Note that the linearity of the expectation value means that

$$
\begin{align*}
\operatorname{Var}(X) & =E\left[\left(X-\mu_{X}\right)^{2}\right]=E\left[X^{2}\right]-2 \mu_{X} E[X]+\mu_{X}^{2} \\
& =E\left[X^{2}\right]-2 \mu_{X}^{2}+\mu_{X}^{2}=E\left[X^{2}\right]-\mu_{X}^{2} \tag{2.7}
\end{align*}
$$

To have a sensible probability distribution, we should satisfy a normalization condition $\sum_{x} p(x)=1$ or $\int_{-\infty}^{\infty} f(x) d x=1$.

One particularly useful expectation value is known as the moment generating function, defined as

$$
\begin{equation*}
M(t)=E\left[e^{t X}\right] \tag{2.8}
\end{equation*}
$$

If we use the Taylor series for the exponential, we can see that

$$
\begin{equation*}
M(t)=E\left[\sum_{k=0}^{\infty} \frac{(t X)^{k}}{k!}\right]=\sum_{k=0}^{\infty} \frac{t^{k}}{k!} E\left[X^{k}\right] \tag{2.9}
\end{equation*}
$$

Considered as a function of $t$, we can think of this in terms of a McLaurin series whose $k$ th coëfficient $M^{(k)}(0) / k$ ! equals the $k$ th moment divided by $k$ !. I.e., if we take the $k$ th derivative of the mgf, evaluated at $t=0$, we get the $k$ th moment:

$$
\begin{equation*}
M^{(k)}(0)=E\left(X^{k}\right) \tag{2.10}
\end{equation*}
$$

Note that the moment generating function doesn't exist for every probability distribution. In general, the expectation value $E[h(X)]$ only exists if $E[|h(X)|]$ is finite. Consider, for example, the Cauchy distribution, which has the probability density function

$$
\begin{equation*}
f(x)=\frac{1}{\pi} \frac{1}{1+x^{2}} \quad-\infty<x<\infty \tag{2.11}
\end{equation*}
$$

It is properly normalized, since

$$
\begin{equation*}
\int_{-\infty}^{\infty} \frac{d x}{1+x^{2}}=\pi \tag{2.12}
\end{equation*}
$$

(which you can show with the trigonometric substitution $x=$ $\tan \theta$, but in any event, the integral is well-defined because the integrand goes like $x^{-2}$ as $x$ goes to $\left.\pm \infty\right)$. This distribution doesn't have a mean or in fact any moments, because

$$
\begin{equation*}
E\left[X^{k}\right]=\int_{-\infty}^{\infty}|x|^{k} \frac{d x}{1+x^{2}} \tag{2.13}
\end{equation*}
$$

diverges because the integrand goes to $|x|^{k-2}$ for large $|x|$. Even if $k=1$, we get logarithmic divergence because the integrand goes like $|x|^{-1}$, which doesn't go to zero fast enough.

A closely related function is the characteristic function $\Phi_{X}(\xi)=E\left[e^{i \xi X}\right]$. For a continuous distribution with pdf $f(x)$, this is

$$
\begin{equation*}
\Phi(\xi)=\int_{-\infty}^{\infty} e^{i \xi x} f(x) d x \tag{2.14}
\end{equation*}
$$

which is, up to conventions about $2 \pi$ and $i$ vs $-i$, the Fourier transform of the pdf. The nice thing about the characteristic function is that it will always exist, because $\int_{-\infty}^{\infty} f(x) d x=1<$ $\infty$. If it is an analytic function, we can use analytic continuation ${ }^{3}$ to define $M(t)=\Phi(-i t)$. In some cases the moment generating function won't exist, which typically means some of the moments are not defined. But if it does, it uniquely determines the probability distribution, basically because the inverse Fourier transform is unique.

### 2.2 Change of Variables in Probability Distributions

Imagine you have a random variable $X$ and another random variable $Y=h(X)$ whose value is given by acting on the random value of $X$ with the deterministic function $h()$. How can we determine the probability distribution for $Y$ from the probability distribution for $X$ ?

Well, if they're discrete random variables, things are pretty straightforward:

$$
\begin{equation*}
p_{Y}(h(x))=P(Y=h(x))=P(X=x)=p_{X}(x) \tag{2.15}
\end{equation*}
$$

[^2]The pmf for $Y$ has the same value as the pmf for $X$; you just have to evaluate it at the appropriate value.

Things get more interesting, though, for continuous random variables, since the pdf is a density, and not the probability of a specific value. It's fundamentally related to the fact that the probability for an event like $x_{1}<X<x_{2}$ has to be the same as that for an equivalent event, say $y_{1}<Y<y_{2}$ where $Y=h(X)$ and either $y_{1}=h\left(x_{1}\right)$ and $y_{2}=h\left(x_{2}\right)$ (if $h(x)$ is monotonically increasing) or $y_{1}=h\left(x_{2}\right)$ and $y_{2}=h\left(x_{1}\right)$ (if $h(x)$ is monotonically decreasing):

$$
\begin{equation*}
\int_{x_{1}}^{x_{2}} f_{X}(x) d x=\int_{y_{1}}^{y_{2}} f_{Y}(y) d y \tag{2.16}
\end{equation*}
$$

You can derive the formula in detail, but the appropriate transformation is suggested by the notation:

$$
\begin{equation*}
\frac{d P}{d y}=\frac{\frac{d P}{d x}}{\left|\frac{d y}{d x}\right|} \tag{2.17}
\end{equation*}
$$

i.e.,

$$
\begin{equation*}
f_{Y}(h(x))=\frac{f_{X}(x)}{\left|h^{\prime}(x)\right|} \tag{2.18}
\end{equation*}
$$

Why the absolute value? Because the probability density for $X$ and $Y$ is defined to be positive, even if the transformation is such that $Y$ decreases with increasing $X$. (Basically, it's a property of the way densities transform.)

### 2.2.1 Example: Inclination

As an example of the importance of a change of variables, consider the inclination $\iota$ between an arbitrary direction (say the normal to the orbital plane of a binary star system) and our line
of sight. This is uniformly distributed in $\chi=\cos \iota$, from $\chi=-1$ to $\chi=1$, so

$$
\begin{equation*}
f(\chi)=\frac{d P}{d \chi}=\frac{1}{2} \quad-1 \leq \chi \leq 1 \tag{2.19}
\end{equation*}
$$

If we change variables to $\iota$, we have to use

$$
\begin{equation*}
\frac{d \chi}{d \iota}=\frac{d}{d \iota} \cos \iota=-\sin \iota \tag{2.20}
\end{equation*}
$$

to find

$$
\begin{equation*}
f(\iota)=\frac{d P}{d \iota}=\frac{d P}{d \chi}\left|\frac{d \chi}{d \iota}\right|=\frac{1}{2} \sin \iota \quad 0 \leq \iota \leq \pi \tag{2.21}
\end{equation*}
$$

which is not uniform in $\iota$.
Of course, this can also be extended to a probability density in the inclination and an azimuthal angle $\psi$, and leads to a direction uniformly distributed over the sphere:

$$
\begin{equation*}
\frac{d^{2} P}{d^{2} \Omega}=\frac{d^{2} P}{\sin \iota d \iota d \psi}=\frac{d^{2} P}{d \chi d \psi}=\frac{1}{4 \pi} \tag{2.22}
\end{equation*}
$$

## Tuesday, September 19, 2017

### 2.3 Multivariate Distributions

Of course, you need not be dealing with only one unknown or random quantity at a time, and it's often useful to consider the joint probability distribution for multiple quantities at once. If they are discrete, this is pretty simple, for example

$$
\begin{equation*}
p\left(x_{1}, x_{2}\right)=P\left(X_{1}=x_{1}, X_{2}=x_{2}\right) \tag{2.23}
\end{equation*}
$$

If they're continuous, you can again define a probability density

$$
\begin{align*}
& f\left(x_{1}, x_{2}\right)=\frac{d^{2} P}{d x_{1} d x_{2}} \\
& =\lim _{\substack{d x_{1} \rightarrow 0 \\
d x_{2} \rightarrow 0}} \frac{P\left(x_{1}<X_{1}<x_{1}+d x_{1}, x_{2}<X_{2}<x_{2}+d x_{2}\right)}{d x_{1} d x_{2}} \tag{2.24}
\end{align*}
$$

We can use a version of the probability sum rule which says that if $\left\{B_{1}, B_{2}, \ldots, B_{n}\right\}$ are a set of mutually exclusive exhaustive alternatives, so that $A=A, B_{1}+A, B_{2}+\cdots+A, B_{n}$,

$$
\begin{equation*}
P(A \mid I)=\sum_{k=1}^{n} P\left(A, B_{k} \mid I\right) \tag{2.25}
\end{equation*}
$$

This process is known as marginalization, and is useful if we really don't care which of the alternatives $B_{k}$ is true. Likewise, if we want to focus on the probability distribution for $X_{1}$ and ignore the value of $X_{2}$, we can integrate out the joint probability density

$$
\begin{equation*}
f_{X_{1}}\left(x_{1}\right)=\int_{-\infty}^{\infty} f\left(x_{1}, x_{2}\right) d x_{2} \tag{2.26}
\end{equation*}
$$

This can be combined with the product rule $f\left(x_{1}, x_{2}\right)=$ $f_{X_{1} \mid X_{2}}\left(x_{1} \mid x_{2}\right) f_{X_{2}}\left(x_{2}\right)$ to write

$$
\begin{equation*}
f_{X_{1}}\left(x_{1}\right)=\int_{-\infty}^{\infty} f\left(x_{1}, x_{2}\right) d x_{2}=\int_{-\infty}^{\infty} f_{X_{1} \mid X_{2}}\left(x_{1} \mid x_{2}\right) f_{X_{2}}\left(x_{2}\right) d x_{2} \tag{2.27}
\end{equation*}
$$

where $f_{X_{1} \mid X_{2}}\left(x_{1} \mid x_{2}\right)$ is the conditional probability distribution for $X_{1}$, assuming that $X_{2}$ takes on the value $x_{2}$.

### 2.3.1 Expectation Values, Variance, Covariance and Correlation

If we have a multivariate probability distribution, we can define the expectation value of a function of the variables with the appropriate sum or integral, e.g., for a continuous distribution with density $f(x, y)$, the expectation value of a function $g(X, Y)$ is

$$
\begin{equation*}
E[g(X, Y)]=\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y) f(x, y) d x d y \tag{2.28}
\end{equation*}
$$

In addition to the usual means and variances:

$$
\begin{equation*}
\mu_{X}=E[X]=\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x f(x, y) d x d y=\int_{-\infty}^{\infty} x f_{X}(x) d x \tag{2.29}
\end{equation*}
$$

and

$$
\begin{equation*}
\operatorname{Var}(X)=\sigma_{X}^{2}=E\left[\left(X-\mu_{X}\right)^{2}\right]=E\left[X^{2}\right]-\mu_{X}^{2} \tag{2.30}
\end{equation*}
$$

we can define the covariance

$$
\begin{equation*}
\operatorname{Cov}(X, Y)=E\left[\left(X-\mu_{X}\right)\left(Y-\mu_{Y}\right)\right]=E[X Y]-\mu_{X} \mu_{Y} \tag{2.31}
\end{equation*}
$$

and the dimensionless correlation

$$
\begin{equation*}
\operatorname{Corr}(X, Y)=\frac{\operatorname{Cov}(X, Y)}{\sigma_{X} \sigma_{Y}} \tag{2.32}
\end{equation*}
$$

We can also define the moment generating function for a series of variables, e.g.,

$$
\begin{equation*}
M\left(t_{1}, t_{2}\right)=E\left[\exp \left(t_{1} X_{1}+t_{2} X_{2}\right)\right] \tag{2.33}
\end{equation*}
$$

where the moments are written in terms of partial derivatives

$$
\begin{equation*}
E\left[X_{1}^{m_{1}} X_{2}^{m_{2}}\right]=\left.\frac{\partial^{m_{1}+m_{2}}}{\partial^{m_{1}} t_{1} \partial^{m_{2}} t_{2}} M\left(t_{1}, t_{2}\right)\right|_{\left(t_{1}, t_{2}\right)=(0,0)} \tag{2.34}
\end{equation*}
$$

### 2.3.2 Change of Variables

For the case of the transformation of continuous random variables we have to deal with the fact that $f_{X_{1}, X_{2}}\left(x_{1}, x_{2}\right)$ and $f_{Y_{1}, Y_{2}}\left(y_{1}, y_{2}\right)$ are probability densities and the volume (area) element has to be transformed from one set of variables to the other. If we write $f_{X_{1}, X_{2}}\left(x_{1}, x_{2}\right) \sim \frac{d^{2} P}{d x_{1} d x_{2}}$ and $f_{Y_{1}, Y_{2}}\left(y_{1}, y_{2}\right) \sim$ $\frac{d^{2} P}{d y_{1} d y_{2}}$, the transformation we'll need is

$$
\begin{equation*}
\frac{d^{2} P}{d y_{1} d y_{2}} \sim\left|\operatorname{det} \frac{\partial\left(x_{1}, x_{2}\right)}{\partial\left(y_{1}, y_{2}\right)}\right| \frac{d^{2} P}{d x_{1} d x_{2}} \tag{2.35}
\end{equation*}
$$

where we use the determinant of the Jacobian matrix

$$
\frac{\partial\left(y_{1}, y_{2}\right)}{\partial\left(x_{1}, x_{2}\right)}=\left(\begin{array}{ll}
\frac{\partial y_{1}}{\partial x_{1}} & \frac{\partial y_{1}}{\partial x_{2}}  \tag{2.36}\\
\frac{\partial y_{2}}{\partial x_{1}} & \frac{\partial y_{2}}{\partial x_{2}}
\end{array}\right)
$$

which may be familiar from the transformation of the volume element

$$
\begin{equation*}
d y_{1} d y_{2}=\left|\operatorname{det} \frac{\partial\left(y_{1}, y_{2}\right)}{\partial\left(x_{1}, x_{2}\right)}\right| d x_{1} d x_{2} \tag{2.37}
\end{equation*}
$$

if we change variables in a double integral.
To get a concrete handle on this, consider an example. Let $X$ and $Y$ be continuous random variables with a joint pdf

$$
f_{X, Y}(x, y)= \begin{cases}\frac{4}{\pi} e^{-x^{2}-y^{2}} & 0<x<\infty ; 0<y<\infty  \tag{2.38}\\ 0 & \text { otherwise }\end{cases}
$$

If we want to calculate the probability that $X^{2}+Y^{2}<a^{2}$ we have to integrate over the part of this disc which lies in the first quadrant $x>0, y>0$ (where the pdf is non-zero):


The limits of the $x$ integral are determined by $0<x$ and $x^{2}+$ $y^{2}<a$, i.e., $x<\sqrt{a^{2}-y^{2}}$; the range of $y$ values represented can be seen from the figure to be $0<y<a$, so we can write the probability as

$$
\begin{equation*}
P\left(X^{2}+Y^{2}<a^{2}\right)=\int_{0}^{a} \int_{0}^{\sqrt{a^{2}-y^{2}}} \frac{4}{\pi} e^{-x^{2}-y^{2}} d x d y \tag{2.39}
\end{equation*}
$$

but we can't really do the integral in this form. However, if we define random variables $R=\sqrt{X^{2}+Y^{2}}$ and ${ }^{4} \Phi=\tan ^{-1}(Y / X)$, so that $X=R \cos \Phi$ and $Y=R \sin \Phi$, we can write the probability as

$$
\begin{equation*}
P\left(X^{2}+Y^{2}<a^{2}\right)=P(R<a)=\int_{0}^{\pi / 2} \int_{0}^{a} f_{R, \Phi}(r, \phi) d r d \phi \tag{2.41}
\end{equation*}
$$

if we have the transformed $\operatorname{pdf} f_{R, \Phi}(r, \phi)$. On the other hand, we know that we can write the volume element $d x d y=r d r d \phi$. We can get this either from geometry in this case, or more generally by differentiating the transformation

$$
\begin{equation*}
\binom{x}{y}=\binom{r \cos \phi}{r \sin \phi} \tag{2.42}
\end{equation*}
$$

${ }^{4}$ Note that we can only get away with using the arctangent $\tan ^{-1}(y / x)$ as an expression for $\phi$ because $x$ and $y$ are both positive. In general, we need to be careful; $(x, y)=(-1,-1)$ corresponds to $\phi=-3 \pi / 4$ even though $\tan ^{-1}([-1] /[-1])=\tan ^{-1}(1)=\pi / 4$ if we use the principal branch of the arctangent. For a general point in the $(x, y)$ plane, we'd need to use the function

$$
\operatorname{atan} 2(y, x)= \begin{cases}\tan ^{-1}(y / x)-\pi & x<0 \text { and } y<0  \tag{2.40}\\ -\pi / 2 & x=0 \text { and } y<0 \\ \tan ^{-1}(y / x) & x>0 \\ \pi / 2 & x=0 \text { and } y>0 \\ \tan ^{-1}(y / x)+\pi & x<0 \text { and } y \geq 0\end{cases}
$$

to get

$$
\binom{d x}{d y}=\binom{\cos \phi d r-r \sin \phi d \phi}{\sin \phi d r+r \cos \phi d \phi}=\left(\begin{array}{cc}
\cos \phi & -r \sin \phi  \tag{2.43}\\
\sin \phi & r \cos \phi
\end{array}\right)\binom{d r}{d \phi}
$$

and taking the determinant of the Jacobian matrix:

$$
\operatorname{det} \frac{\partial(x, y)}{\partial(r, \phi)}=\left|\begin{array}{cc}
\cos \phi & -r \sin \phi  \tag{2.44}\\
\sin \phi & r \cos \phi
\end{array}\right|=r \cos ^{2} \phi+r \sin ^{2} \phi=r
$$

so the volume element transforms like

$$
\begin{equation*}
d x d y=\left|\operatorname{det} \frac{\partial(x, y)}{\partial(r, \phi)}\right| d r d \phi=r d r d \phi \tag{2.45}
\end{equation*}
$$

Even if we knew nothing about the transformation of random variables, we could use this to change variables in the integral (2.39) to get

$$
\begin{equation*}
\int_{0}^{a} \int_{0}^{\sqrt{a^{2}-y^{2}}} \frac{4}{\pi} e^{-x^{2}-y^{2}} d x d y=\int_{0}^{\pi / 2} \int_{0}^{a} \frac{4}{\pi} e^{-r^{2}} r d r d \phi \tag{2.46}
\end{equation*}
$$

If we compare the integrands of 2.46 and 2.46 we can see that the transformed pdf must be

$$
f_{R, \Phi}(r, \phi)= \begin{cases}r e^{-r^{2}} & 0<r<\infty ; 0<\phi<\pi / 2  \tag{2.47}\\ 0 & \text { otherwise }\end{cases}
$$

Incidentally, we can calculate the probability as

$$
\begin{align*}
P(R<a) & =\int_{0}^{\pi / 2} \int_{0}^{a} \frac{4}{\pi} e^{-r^{2}} r d r d \phi=\int_{0}^{a} e^{-r^{2}} 2 r d r=-\left.e^{-r^{2}}\right|_{0} ^{a} \\
& =1-e^{-a^{2}} \tag{2.48}
\end{align*}
$$

To return to the general case, we see there are basically two things to worry about: one is the Jacobian determinant relating
the volume elements in the two sets of variables, and the other is transforming the ranges of variables used to describe the event, as well as the allowed range of variables. In general terms, if $\mathcal{S}$ is the support of the random variables $X_{1}$ and $X_{2}$, i.e., the smallest region of $\mathbb{R}^{2}$ such that $P\left[\left(X_{1}, X_{2}\right) \in \mathcal{S}\right]=1$ and $\mathcal{T}$ is the support of $Y_{1}$ and $Y_{2}$, we need a transformation of the pdf $f_{X_{1}, X_{2}}\left(x_{1}, x_{2}\right)$ defined on $\mathcal{S}$ such that

$$
\begin{align*}
P\left[\left(X_{1}, X_{2}\right)\right. & \in A]=\iint_{A} f_{X_{1}, X_{2}}\left(x_{1}, x_{2}\right) d x_{1} d x_{2} \\
& =\iint_{B} f_{Y_{1}, Y_{2}}\left(y_{1}, y_{2}\right) d y_{1} d y_{2}=P\left[\left(Y_{1}, Y_{2}\right) \in B\right] \tag{2.49}
\end{align*}
$$

where $B$ is the image of $A$ under the transformation, i.e., $\left(x_{1}, x_{2}\right) \in A$ is equivalent to $\left\{u_{1}\left(x_{1}, x_{2}\right), u_{2}\left(x_{1}, x_{2}\right)\right\} \in B$. Since a change of variables in the integral gives us

$$
\begin{align*}
& \iint_{A} f_{X_{1}, X_{2}}\left(x_{1}, x_{2}\right) d x_{1} d x_{2} \\
& \quad=\iint_{B} f_{X_{1}, X_{2}}\left(w_{1}\left(y_{1}, y_{2}\right), w_{2}\left(y_{1}, y_{2}\right)\right)\left|\operatorname{det} \frac{\partial\left(x_{1}, x_{2}\right)}{\partial\left(y_{1}, y_{2}\right)}\right| d y_{1} d y_{2} \tag{2.50}
\end{align*}
$$

we must have, in general,

$$
\begin{array}{r}
f_{Y_{1}, Y_{2}}\left(y_{1}, y_{2}\right)=\left|\operatorname{det} \frac{\partial\left(x_{1}, x_{2}\right)}{\partial\left(y_{1}, y_{2}\right)}\right| f_{X_{1}, X_{2}}\left(w_{1}\left(y_{1}, y_{2}\right), w_{2}\left(y_{1}, y_{2}\right)\right) \\
\left(y_{1}, y_{2}\right) \in \mathcal{T} \quad(2 . \tag{2.51}
\end{array}
$$

which is the more careful way of writing the easier-to-remember formula we started with:

$$
\begin{equation*}
\frac{d^{2} P}{d y_{1} d y_{2}} \sim\left|\operatorname{det} \frac{\partial\left(x_{1}, x_{2}\right)}{\partial\left(y_{1}, y_{2}\right)}\right| \frac{d^{2} P}{d x_{1} d x_{2}} \tag{2.52}
\end{equation*}
$$

### 2.3.3 General Formula

You can also perform this a change of variables on a joint probability density to write it in terms of another set of variables. Suppose you have $N$ random variables $\left\{X_{i}\right\}$ from which you can determine the values of $N$ random variables $\left\{Y_{i}\right\}$. Then the transformation uses the Jacobian determinant:

$$
\begin{equation*}
\frac{d^{N} P}{d^{N} y}=\left(\frac{d^{N} P}{d^{N} x}\right) /\left|\operatorname{det}\left\{\frac{\partial y_{i}}{\partial x_{j}}\right\}\right| \tag{2.53}
\end{equation*}
$$

You can see this is the right thing to do because the Jacobian determinant is used to transform the measure of a multiple integral:

$$
\begin{equation*}
d^{N} y=\left|\operatorname{det}\left\{\frac{\partial y_{i}}{\partial x_{j}}\right\}\right| d^{N} x \tag{2.54}
\end{equation*}
$$

and these probability densities are meant to be put under multiple integrals. Written in more standard notation, if we define

$$
\begin{equation*}
\mathbf{x} \equiv\left\{x_{i}\right\}, \quad \mathbf{y} \equiv\left\{y_{i}\right\} \tag{2.55}
\end{equation*}
$$

and

$$
\begin{equation*}
J_{\mathbf{y x}}(\mathbf{x})=\operatorname{det}\left\{\frac{\partial y_{i}}{\partial x_{j}}\right\} \tag{2.56}
\end{equation*}
$$

then

$$
\begin{equation*}
f_{\mathbf{Y}}(\mathbf{h}(\mathbf{x}))=\frac{f_{\mathbf{X}}(\mathbf{x})}{\left|J_{\mathbf{y x}}(\mathbf{x})\right|} \tag{2.57}
\end{equation*}
$$

Thursday, September 21, 2017

## 3 Probability Distributions

### 3.1 Some Specific Probability Distributions

See Gregory, Chapter 5

Before we delve into frequentist and Bayesian applications of probability distributions, it's useful to consider some specific random variables, the sort of physical situations to which they're relevant, and what their probability distributions look like.

### 3.1.1 The Binomial Distribution (discrete)

Consider a sequence of identical independent yes/no questions, for example, the outcomes of a series of $n$ flips of a coin, where the probability of heads on each flip is $\theta$. Let $H_{i}$ denote the event "the $i$ th flip is heads" and $T_{i}=\overline{H_{i}}$ denote "the $i$ th flip is tails", so that $P\left(H_{i} \mid \theta, I\right)=\theta$ and $P\left(T_{i} \mid \theta, I\right)=1-\theta$ If we want to get the probability of a particular sequence of heads and tails, we can use the product rule, e.g., the probability of a head on the first flip followed by a tail on the second flip is

$$
\begin{equation*}
P\left(H_{1}, T_{2} \mid \theta, I\right)=P\left(H_{1} \mid \theta, I\right) P\left(T_{2} \mid H_{1}, \theta, I\right) \tag{3.1}
\end{equation*}
$$

If the information $I$ includes the fact that coins have no "memory", i.e., we're not any more or less likely to flip a head on the second flip because we've got a head on the first flip, then $P\left(T_{2} \mid H_{1}, \theta, I\right)=P\left(T_{2} \mid \theta, I\right)=1-\theta$, so that

$$
\begin{equation*}
P\left(H_{1}, T_{2} \mid \theta, I\right)=P\left(H_{1} \mid \theta, I\right) P\left(T_{2} \mid \theta, I\right)=\theta(1-\theta) \tag{3.2}
\end{equation*}
$$

Note, this is a special case of independence of events. Saying that two events $A$ and $B$ are independent in the context of information $I$ means that we can use a specialized product rule

$$
\begin{equation*}
P(A, B \mid I)=P(A \mid I) P(B \mid I) \quad \text { iff } A \& B \text { independent } \tag{3.3}
\end{equation*}
$$

If we consider the event $H_{1}, H_{2}, T_{3}$, the probability is

$$
\begin{align*}
P\left(H_{1}, H_{2}, T_{3} \mid \theta, I\right) & =P\left(H_{1} \mid \theta, I\right) P\left(H_{2} \mid \theta, I\right) P\left(T_{3} \mid \theta, I\right) \\
& =(\theta)(\theta)(1-\theta)=\theta^{2}(1-\theta)^{1} \tag{3.4a}
\end{align*}
$$

Similarly

$$
\begin{align*}
P\left(H_{1}, T_{2}, H_{3} \mid \theta, I\right) & =P\left(H_{1} \mid \theta, I\right) P\left(T_{2} \mid \theta, I\right) P\left(H_{3} \mid \theta, I\right) \\
& =(\theta)(1-\theta)(\theta)=\theta^{2}(1-\theta)^{1}  \tag{3.4b}\\
P\left(T_{1}, H_{2}, H_{3} \mid \theta, I\right) & =P\left(T_{1} \mid \theta, I\right) P\left(H_{2} \mid \theta, I\right) P\left(H_{3} \mid \theta, I\right) \\
& =(1-\theta)(\theta)(\theta)=\theta^{2}(1-\theta)^{1} \tag{3.4c}
\end{align*}
$$

So each possible sequence of two heads and a tail has a probability of $\theta^{2}(1-\theta)^{1}$. In general, the probability for each possible sequence of $k$ heads and $n-k$ tails is $\theta^{k}(1-\theta)^{n-k}$. To get the overall probability of $k$ heads and $n-k$ tails in $n$ flips, we need to add up the probabilities for all of the relevant sequence; the result is the probability mass function for a binomial random variable $K, p(k \mid n, \theta)$. For example, we see there are three possible sequences of two heads and a tail, so

$$
\begin{equation*}
p(2 \mid 3, \theta)=3 \theta^{2}(1-\theta)^{1} \tag{3.5}
\end{equation*}
$$

In general, the number of ways to pick $k$ of the $n$ flips to be heads is " $n$ choose $k$ ", $\binom{n}{k}=\frac{n!}{k!(n-k)!}$, also known as ${ }_{n} C_{k}$, or the "binomial coëfficient" ${ }^{5}$ so the pmf for a binomial distribution is

$$
\begin{equation*}
p(k \mid n, \theta)=\binom{n}{k} \theta^{k}(1-\theta)^{n-k} \quad k=0,1, \ldots, n \tag{3.6}
\end{equation*}
$$

We can show that this probability distribution is normalized using the binomial expansion formula

$$
\begin{equation*}
(a+b)^{n}=\sum_{k=0}^{n}\binom{n}{k} a^{k} b^{n-k} \tag{3.7}
\end{equation*}
$$

[^3]so that
\[

$$
\begin{equation*}
\sum_{k=0}^{n} p(k \mid n, \theta)=\sum_{k=0}^{n}\binom{n}{k} \theta^{k}(1-\theta)^{n-k}=(\theta+[1-\theta])^{n}=1^{n}=1 \tag{3.8}
\end{equation*}
$$

\]

The same sort of calculation can be used to find the moment generating function and use it to find the mean and variance relatively quickly:

$$
\begin{align*}
M(t) & =E\left[e^{t K}\right]=\sum_{k=0}^{n} e^{t k} p(k \mid n, \theta)=\binom{n}{k}\left(\theta e^{t}\right)^{k}(1-\theta)^{n-k} \\
& =\left(\theta e^{t}+[1-\theta]\right)^{n} \tag{3.9}
\end{align*}
$$

It is often easier to work with the logarithm of the mgf, known as the cumulant generating function, which you studied on the most recent homework:

$$
\begin{equation*}
\psi(t)=\ln M(t)=n \ln \left(\theta e^{t}+[1-\theta]\right) \tag{3.10}
\end{equation*}
$$

We use the chain rule to take its derivative:

$$
\begin{equation*}
\psi^{\prime}(t)=\frac{n \theta e^{t}}{\theta e^{t}+[1-\theta]}=\frac{n \theta}{\theta+[1-\theta] e^{-t}} \tag{3.11}
\end{equation*}
$$

from which we get the mean

$$
\begin{equation*}
E(K)=\psi^{\prime}(0)=\frac{n \theta}{\theta+[1-\theta]}=n \theta \tag{3.12}
\end{equation*}
$$

and differentiating again gives

$$
\begin{equation*}
\psi^{\prime \prime}(t)=-\frac{-n \theta(1-\theta) e^{-t}}{\left(\theta+[1-\theta] e^{-t}\right)^{2}} \tag{3.13}
\end{equation*}
$$

from which we get the variance:

$$
\begin{equation*}
\operatorname{Var}(K)=\psi^{\prime \prime}(0)=\frac{n \theta(1-\theta)}{(\theta+[1-\theta])^{2}}=n \theta(1-\theta) \tag{3.14}
\end{equation*}
$$

### 3.1.2 Other Related Distributions

We mention briefly a few other distributions here, which are related to the binomial distribution, and which we will consider more closely as we need them for applications.

The hypergeometric distribution describes "sampling without replacement", where we have e.g., $N$ balls, of which $M$ are red, and we draw $n$ of them without replacement, so the population of available balls changes with each draw; the number of red balls in our sample will be a hypergeometric random variable $x$ whose pmf is

$$
\begin{equation*}
p(x)=\frac{\binom{M}{x}\binom{N-M}{n-x}}{\binom{N}{n}} \tag{3.15}
\end{equation*}
$$

(The situation described by the binomial distribution can be thought of as "sampling with replacement", where you put the each ball back after drawing it, and mix things up again, so the probability of drawing a red ball remains $M / N$.

The negative binomial distribution describes a situation where instead of flipping the coin a fixed number $n$ of times, we decide to keep flipping until we have $r$ heads. The random variable $X$ is the number of tails before the $r$ th head. Using the fact that there must be a string of $r-1$ heads and $x$ tails followed by one final failure, we can show that probability mass function is

$$
\begin{equation*}
p(x)=\binom{x+r-1}{r-1} p^{r-1}(1-p)^{x} p=\binom{x+r-1}{r-1} p^{r}(1-p)^{x} \tag{3.16}
\end{equation*}
$$

The negative binomial distribution is important when considering the so-called "optional stopping problem", where questions of frequentist inference depend not just on the sequence of results seen, but on when you were planning to stop the experiment.

The multinomial distribution is a generalization of the binomial distribution, where instead of a series of coin flips, each with 2 possible outcomes, we have a series of trials which each have $m$ possible outcomes. (E.g., we roll an $m$-sided die $n$ times.) The parameters are $\left\{\theta_{1}, \theta_{2}, \ldots, \theta_{m}\right\}$, the probabilities for each of the $m$ outcomes. Not only must each $\theta_{i}$ lie in the range $0 \leq \theta_{i} \leq 1$, they must also sum to one: $\sum_{i=1}^{m} \theta_{i}=1$. We then have $m$ random variables $K_{1}, K_{2}, \ldots, K_{m}$ with the joint pmf

$$
\begin{align*}
& p\left(k_{1}, k_{2}, \ldots, k_{m}\right)=\frac{n!}{k_{1}!k_{2}!\cdots k_{m}!} \theta_{1}^{k_{1}} \theta_{2}^{k_{2}} \cdots \theta_{m}^{k_{m}} \\
& k_{1}=0,1, \ldots n ; \quad k_{2}=0,1, \ldots, n-k_{1} ; \\
& \cdots ; \quad k_{m}=n-k_{1}-k_{2}-\cdots-k_{m-1} \tag{3.17}
\end{align*}
$$

### 3.1.3 The Poisson Distribution (discrete)

Consider a random process in which discrete events happen independently of each other with an average rate of $r$. If we count the number of events in an interval of duration $T$, that number of events is a random quantity with an expected value of $\mu=r T$. We usually think of the rate $r$ as having units of inverse time and the duration $T$ as having units of time. Examples of processes with rates in time include popcorn kernels popping, clicks on a Geiger counter, or gamma-ray bursts observed. But the interval could also be in space, e.g., we could be counting the number of cosmic rays collected in a detector of a certain area, or the number of galaxies within a certain redshift range found in a patch of sky of a given solid angle. The number $K$ of events is a random variable with a probability mass function

$$
\begin{equation*}
P(K=k)=p(k \mid \mu) \tag{3.18}
\end{equation*}
$$

Such a process is called a Poisson process if we can sub-divide the interval into smaller intervals (in time, space, sky position,
or whatever) and then the number of events in each sub-interval is an independent random variable with the same properties (but a smaller rate, obviously).

One way to wrap our head around a Poisson process is to think about collecting a large number $M$ of identical intervals, in which the expected total number of events is $M \mu$, and randomly assigning each of those $M \mu$ events to one of the $M$ intervals, without regard to where the other events have been placed. The number of events landing in any interval will be approximately $\sqrt{6}^{6}$ a Poisson random variable with Poisson parameter $\mu$. See the ipython notebook http://ccrg.rit.edu/~whelan/courses/2017_3fa_ASTP_ 611/data/notes_probability_poisson.ipynb for an exploration of this. (As a further investigation, try removing the artificial factor of 200 scale-up and see how the histogram generated only approximately tracks the Poisson pmf.)

We can calculate the pmf for the Poisson random variable $K$ as follows: subdivide the interval into some large number $N$ of identically-sized sub-intervals. Each one has an expected number of events of $\mu / N$. If we choose $N$ large enough, we can make this number very very small. This means that in any one sub-interval, the odds are pretty good that there will be no events. There is some small chance (of order $\mu / N$ ) of seeing one event, and a vanishingly small chance (of order $[\mu / N]^{2}$ ) of seeing

[^4]two or more events in this sub-interval:
\[

$$
\begin{align*}
p(0 \mid \mu / N) & =1-\mu / N+\mathcal{O}\left([\mu / N]^{2}\right)  \tag{3.19a}\\
p(1 \mid \mu / N) & =\mu / N+\mathcal{O}\left([\mu / N]^{2}\right)  \tag{3.19b}\\
\sum_{k=2}^{\infty} p(k \mid \mu / N) & =\mathcal{O}\left([\mu / N]^{2}\right) \tag{3.19c}
\end{align*}
$$
\]

but this is basically a single trial which can have a yes (there is an event) or no (there is not an event) result, and the total number $K$ of events in the larger interval can be approximated by a binomial random variable with $N$ trials and a probability for success of $\mu / N$ for each trial. That means

$$
\begin{align*}
p(k \mid \mu) & =\lim _{N \rightarrow \infty} b(k \mid \mu / N, N)=\lim _{N \rightarrow \infty} \frac{N!}{(N-k)!k!}\left(\frac{\mu}{N}\right)^{k}\left(1-\frac{\mu}{N}\right)^{N-k} \\
& =\frac{(\mu)^{k}}{k!} \lim _{N \rightarrow \infty}\left(1-\frac{\mu}{N}\right)^{N} \frac{N!}{(N-k)!}(N-\mu)^{-k} \tag{3.20}
\end{align*}
$$

Now,

$$
\begin{equation*}
\frac{N!}{(N-k)!}=N(N-1) \ldots(N-k+1)=\prod_{\ell=0}^{k-1}(N-\ell) \tag{3.21}
\end{equation*}
$$

and of course

$$
\begin{equation*}
(N-\mu)^{-k}=\prod_{\ell=0}^{k-1} \frac{1}{N-\mu} \tag{3.22}
\end{equation*}
$$

SO

$$
\begin{equation*}
\frac{N!}{(N-k)!}(N-\mu)^{-k}=\prod_{\ell=0}^{k-1} \frac{N-\ell}{N-\mu}=\prod_{\ell=0}^{k-1} \frac{1-\ell / N}{1-\mu / N} \tag{3.23}
\end{equation*}
$$

but for finite $k$ this is the product of a finite number of things, each of which goes to 1 as $N \rightarrow \infty$, so

$$
\begin{equation*}
p(k \mid r, T)=\frac{(\mu)^{k}}{k!} \lim _{N \rightarrow \infty}\left(1-\frac{\mu}{N}\right)^{N}=\frac{(\mu)^{k}}{k!} e^{-\mu} . \tag{3.24}
\end{equation*}
$$

This is the Poisson distribution. It's easy to check that it's normalized, i.e.,

$$
\begin{equation*}
\sum_{k=0}^{\infty} p(k \mid \mu)=e^{-\mu} \sum_{k=0}^{\infty} \frac{(\mu)^{k}}{k!}=e^{-\mu} e^{\mu}=1 \tag{3.25}
\end{equation*}
$$

Note that as a consistency check we can go back and verify our assumptions (3.19):

$$
\begin{gather*}
p(0 \mid \mu / N)=e^{-\mu / N}=1-\mu / N+\mathcal{O}\left([\mu / N]^{2}\right)  \tag{3.26a}\\
p(1 \mid \mu / N)=\frac{\mu}{N} e^{-\mu / N}=\frac{\mu}{N}+\mathcal{O}\left([\mu / N]^{2}\right)  \tag{3.26b}\\
\sum_{k=2}^{\infty} p(k \mid \mu / N)=1-p(0 \mid \mu / N)-p(1 \mid \mu / N)=\mathcal{O}\left([\mu / N]^{2}\right) \tag{3.26c}
\end{gather*}
$$

We can get the mean and variance of the Poisson distribution either by taking the limiting forms of those for the binomial distribution, or by calculating the moment generating function:

$$
\begin{equation*}
M(t)=\sum_{k=0}^{\infty} e^{t k} \frac{\mu^{k}}{k!} e^{-\mu}=\frac{\left(\mu e^{t}\right)^{k}}{k!} e^{-\mu}=e^{\mu e^{t}} e^{-\mu}=e^{\mu\left(e^{t}-1\right)} \tag{3.27}
\end{equation*}
$$

To find the mean and variance, it's useful once again to use the cumulant generating function $\psi(t)=\ln M(t)$ :

$$
\begin{equation*}
\psi(t)=\mu\left(e^{t}-1\right) \tag{3.28}
\end{equation*}
$$

Differentiating gives us

$$
\begin{equation*}
\psi^{\prime}(t)=\mu e^{t} \tag{3.29}
\end{equation*}
$$

and

$$
\begin{equation*}
\psi^{\prime \prime}(t)=\mu e^{t} \tag{3.30}
\end{equation*}
$$

so the mean is

$$
\begin{equation*}
E(K)=\psi^{\prime}(0)=\mu \tag{3.31}
\end{equation*}
$$

(which was the definition we started with) and the variance is

$$
\begin{equation*}
\operatorname{Var}(K)=\psi^{\prime \prime}(0)=\mu \tag{3.32}
\end{equation*}
$$

## Tuesday, September 26, 2017

### 3.1.4 The Exponential Distribution (continuous)

Let's consider further a Poisson process. The Poisson distribution gives the pmf for the total number of events in an interval, which is a discrete random variable. Now consider another question. Suppose we are observing a Poisson process with an event rate $r$. Let's assume the intervals are in time, so that $r$ has units of inverse time, and the number of events in a time $\Delta t$ will be a Poisson random variable with parameter $r \Delta t$. Now suppose we start watching at a given time and see how long we have to wait for the next event. This waiting time $T$ will itself be a random variable, with a probability density function $f_{T}(t \mid r)$ which depends on the rate $r$. Note that this is a continuous random variable. We can actually work out the pdf from our knowledge of the Poisson process. Consider the probability that $T$ is longer than some value $t$ :

$$
\begin{equation*}
P(T>t)=\int_{t}^{\infty} f_{T}\left(t^{\prime} \mid r\right) d t^{\prime} \tag{3.33}
\end{equation*}
$$

This is the probability that in the interval of length $t$, beginning when we start watching, there are no events. But we know how to write the probability that there are no events from a Poisson process in an interval of a given length. It is the probability that the corresponding Poisson random variable (which has parameter $r t$ ) will take on the value 0 :

$$
\begin{equation*}
P(K=0)=p(0 \mid r t)=\frac{(r t)^{0}}{0!} e^{-r t}=e^{-r t} \tag{3.34}
\end{equation*}
$$

Equating the two expressions for this probability gives

$$
\begin{equation*}
\int_{t}^{\infty} f_{T}\left(t^{\prime} \mid r\right) d t^{\prime}=e^{-r t} \tag{3.35}
\end{equation*}
$$

We can differentiate both sides with respect to $t$ (not $t^{\prime}$ ) and find

$$
\begin{equation*}
-f_{T}(t \mid r)=-r e^{-r t} \tag{3.36}
\end{equation*}
$$

which gives us the pdf for $T$, the exponential distribution:

$$
\begin{equation*}
f(t \mid r)=r e^{-r t} \quad t \geq 0 \tag{3.37}
\end{equation*}
$$

It is often conventional to refer to the rate parameter as $\lambda$ rather than $r$, so the pdf is

$$
\begin{equation*}
f(t \mid \lambda)=\lambda e^{-\lambda t} \quad t \geq 0 \tag{3.38}
\end{equation*}
$$

The pdf looks like this:


Note that this derivation made use of an integral of the pdf. For a general continuous random variable, we define the cumulative distribution function

$$
\begin{equation*}
F_{X}(x):=P(X \leq x)=\int_{-\infty}^{x} f_{X}\left(x^{\prime}\right) d x^{\prime} \tag{3.39}
\end{equation*}
$$

The derivative of the cdf is the pdf:

$$
\begin{equation*}
\frac{d F_{X}}{d x}(x)=f_{X}(x) \tag{3.40}
\end{equation*}
$$

### 3.1.5 The Gamma Distribution (continuous)

The plot of the exponential distribution above has the same shape regardless of the value of the parameter $\lambda$; changing its value just changes the scales of the axes. It is, however, one member of a family of distributions known as the Gamma distribution, with two parameters $\alpha>0$ and $\beta>0$, and the pdf

$$
\begin{equation*}
f(x \mid \alpha, \beta)=\frac{1}{\Gamma(\alpha) \beta^{\alpha}} x^{\alpha-1} e^{-x / \beta} \quad 0<x<\infty \tag{3.41}
\end{equation*}
$$

where

$$
\begin{equation*}
\Gamma(\alpha)=\int_{0}^{\infty} u^{\alpha-1} e^{-u} d u \tag{3.42}
\end{equation*}
$$

is the Gamma function. If $n$ is a non-negative integer, $\Gamma(n+1)=$ $n!$, and in general $\Gamma(\alpha+1)=\alpha \Gamma(\alpha)$. Comparing the pdfs, we see that an exponential distribution with rate $\lambda$ is a Gamma distribution with shape parameter $\alpha=1$ and scale parameter $\beta=1 / \lambda]^{7}$ Here's the shape of the pdf for different choices of $\alpha$.

[^5]

We can check that the pdf is normalized by taking the integral

$$
\begin{array}{r}
\frac{1}{\Gamma(\alpha) \beta^{\alpha}} \int_{0}^{\infty} x^{\alpha-1} e^{-x / \beta} d x=\frac{1}{\Gamma(\alpha)} \int_{0}^{\infty}\left(\frac{x}{\beta}\right)^{\alpha-1} e^{-x / \beta} \frac{d x}{\beta} \\
=\frac{1}{\Gamma(\alpha)} \int_{0}^{\infty} u^{\alpha-1} e^{-u} d u=\frac{1}{\Gamma(\alpha)} \Gamma(\alpha)=1 \tag{3.43}
\end{array}
$$

where we have made the change of variables $u=x / \beta$. The same calculation can be done to find the mgf:

$$
\begin{align*}
M(t) & =E\left(e^{t X}\right)=\frac{1}{\Gamma(\alpha) \beta^{\alpha}} \int_{0}^{\infty} x^{\alpha-1} e^{-x / \beta} e^{t x} d x \\
& =\frac{1}{\Gamma(\alpha) \beta^{\alpha}} \int_{0}^{\infty} x^{\alpha-1} \exp \left(-\left[\frac{1}{\beta}-t\right] x\right) d x \tag{3.44}
\end{align*}
$$

If we require $t<\frac{1}{\beta}$ and make the substitution $u=\left(\frac{1}{\beta}-t\right) x$ so $x=\frac{\beta}{1-\beta t} u$, this becomes

$$
\begin{equation*}
M(t)=\frac{1}{\Gamma(\alpha) \beta^{\alpha}}\left(\frac{\beta}{1-\beta t}\right)^{\alpha} \int_{0}^{\infty} u^{\alpha-1} e^{-u} d u=(1-\beta t)^{-\alpha} \tag{3.45}
\end{equation*}
$$

If we again construct the cumulant generating function

$$
\begin{equation*}
\psi(t)=\ln M(t)=-\alpha \ln (1-\beta t) \tag{3.46}
\end{equation*}
$$

we find the derivative

$$
\begin{equation*}
\psi^{\prime}(t)=\alpha \beta(1-\beta t)^{-1} \tag{3.47}
\end{equation*}
$$

so the mean is

$$
\begin{equation*}
E(X)=\psi^{\prime}(0)=\alpha \beta \tag{3.48}
\end{equation*}
$$

and the second derivative

$$
\begin{equation*}
\psi^{\prime \prime}(t)=\alpha \beta^{2}(1-\beta t)^{-2} \tag{3.49}
\end{equation*}
$$

so the variance is

$$
\begin{equation*}
\operatorname{Var}(X)=\psi^{\prime \prime}(0)=\alpha \beta^{2} \tag{3.50}
\end{equation*}
$$

This means that in particular, and exponential distribution (for which $\alpha=1$ and $\beta=\lambda^{-1}$. has a mean of $\lambda^{-1}$, a variance of $\lambda^{-2}$, and a standard deviation of $\lambda^{-1}$.

Even more useful than the mean and variance, the mgf can be used to show what happens when we add two independent Gamma random variables with the same scale parameter, say $X_{1}$ with parameters $\left(\alpha_{1}, \beta\right)$ and $X_{2}$ with parameters $\left(\alpha_{2}, \beta\right)$.

First a few words about independence, as it applies to random variables. Just as events $A$ and $B$ are independent if and only if $P(A, B)=P(A) P(B)$, two continuous random variables are independent if and only if their joint pdf is $f\left(x_{1}, x_{2}\right)=$
$f_{X_{1}}\left(x_{1}\right) f_{X_{2}}\left(x_{2}\right)$. This then means that if we take the expectation value of some function of $X_{1}$ times some function of $X_{2}$, we have

$$
\begin{align*}
E\left[h_{1}\left(X_{1}\right) h_{2}\left(X_{2}\right)\right] & =\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h\left(x_{1}\right) h\left(x_{2}\right) f\left(x_{1}, x_{2}\right) d x_{1} d x_{2} \\
& =\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h\left(x_{1}\right) h\left(x_{2}\right) f_{X_{1}}\left(x_{1}\right) f_{X_{2}}\left(x_{2}\right) d x_{1} d x_{2} \\
& =\int_{-\infty}^{\infty} h\left(x_{1}\right) f_{X_{1}}\left(x_{1}\right) d x_{1} \int_{-\infty}^{\infty} h\left(x_{2}\right) f_{X_{2}}\left(x_{2}\right) d x_{2} \\
& =E\left[h_{1}\left(X_{1}\right)\right] E\left[h_{2}\left(X_{2}\right)\right] \quad \text { if } X_{1} \& X_{2} \text { independent } \tag{3.51}
\end{align*}
$$

In particular, if $Y=X_{1}+X_{2}$, where $X_{1}$ and $X_{2}$ are independent, the mgf of $Y$ is

$$
\begin{align*}
M_{Y}(t) & =E\left[e^{t Y}\right]=E\left[e^{t\left(X_{1}+X_{2}\right)}\right]=E\left[e^{t X_{1}} e^{t X_{2}}\right]=E\left[e^{t X_{1}}\right] E\left[e^{t X_{2}}\right] \\
& =M_{X_{1}}(t) M_{X_{2}}(t) \quad \text { if } X_{1} \& X_{2} \text { independent } \tag{3.52}
\end{align*}
$$

Specializing to the case described above where $X_{1}$ and $X_{2}$ are two Gamma random variables with the same scale parameter $\beta$, the mgf of their sum is
$M(t)=M_{1}(t) M_{2}(t)=(1-\beta t)^{-\alpha_{1}}(1-\beta t)^{-\alpha_{2}}=(1-\beta t)^{-\left(\alpha_{1}+\alpha_{2}\right)}$
which we see is the mgf of a Gamma random variable with parameters $\alpha_{1}+\alpha_{2}$ and $\beta$. So, if you add Gamma rvs with the same scale parameter, their sum is another Gamma rv whose scale parameter is the sum of the individual scale parameters.

For example, if we add $n$ independent exponential random variables with the same rate parameter $\lambda$, each of which is a $\operatorname{Gamma}(1,1 / \lambda)$ random variable, their sum is a $\operatorname{Gamma}(n, 1 / \lambda)$ random variable. This special case of the Gamma distribution (where $\alpha$ is an integer) is called an Erlang distribution.

## Thursday, September 28, 2017 <br> Review for First Prelim Exam

## Tuesday, October 3, 2017 <br> First Prelim Exam

## Thursday, October 5, 2017

### 3.1.6 The Gaussian (aka Normal) Distribution (continuous)

The Gaussian distribution, also known as the normal distribution, is a distribution with location parameter $\mu$ and scale parameter $\sigma>0$. We refer to this as a $N\left(\mu, \sigma^{2}\right)$ distribution, which has pdf

$$
\begin{equation*}
f\left(x \mid \mu, \sigma^{2}\right)=\frac{1}{\sigma \sqrt{2 \pi}} \exp \left(-\frac{(x-\mu)^{2}}{2 \sigma^{2}}\right) \quad-\infty<x<\infty \tag{3.54}
\end{equation*}
$$

This turns out to be a good approximation in many situations, for reasons we'll delve into as the course goes on, but in brief

1. It is the limiting form of many distributions, as you'll investigate on the homework, which is mostly due to a result we'll see later known as the Central Limit Theorem
2. It is the natural result of a truncated Taylor expansion either of the logarithm of a pdf about a local maximum, or of an mgf about $t=0$.
3. It is the Maximum Entropy distribution appropriate when you know nothing but the mean and variance of a distribution defined for all $x$.
To show that this is normalized, we take the integral

$$
\begin{equation*}
\int_{-\infty}^{\infty} f(x) d x=\frac{1}{\sigma \sqrt{2 \pi}} \int_{-\infty}^{\infty} e^{-\frac{(x-\mu)^{2}}{2 \sigma^{2}}} d x=\frac{1}{\sqrt{2 \pi}} \int_{-\infty}^{\infty} e^{-z^{2} / 2} d z \tag{3.55}
\end{equation*}
$$

where we have made the substitution $z=(x-\mu) / \sigma$. The integral

$$
\begin{equation*}
I=\int_{-\infty}^{\infty} e^{-z^{2} / 2} d z \tag{3.56}
\end{equation*}
$$

is a bit tricky; there's no ordinary function whose derivative is $e^{-z^{2} / 2}$, so we can't just do an indefinite integral and evaluate at the endpoints. But we can do the definite integral by writing

$$
\begin{align*}
I^{2} & =\left(\int_{-\infty}^{\infty} e^{-x^{2} / 2} d x\right)\left(\int_{-\infty}^{\infty} e^{-y^{2} / 2} d y\right)  \tag{3.57}\\
& =\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-\left(x^{2}+y^{2}\right) / 2} d x d y
\end{align*}
$$

If we interpret this as a double integral in Cartesian coördinates, we can change to polar coördinates $r$ and $\phi$, and write

$$
\begin{align*}
I^{2} & =\int_{0}^{2 \pi} \int_{0}^{\infty} e^{-r^{2} / 2} r d r d \phi=2 \pi \int_{0}^{\infty} e^{-r^{2} / 2} r d r  \tag{3.58}\\
& =-\left.2 \pi e^{-r^{2} / 2}\right|_{0} ^{\infty}=2 \pi
\end{align*}
$$

SO

$$
\begin{equation*}
I=\int_{-\infty}^{\infty} e^{-z^{2} / 2} d z=\sqrt{2 \pi} \tag{3.59}
\end{equation*}
$$

and the pdf does integrate to one.
To get the mgf, we have to take the integral

$$
\begin{equation*}
M(t)=\frac{1}{\sigma \sqrt{2 \pi}} \int_{-\infty}^{\infty} \exp \left(-\frac{(x-\mu)^{2}}{2 \sigma^{2}}+t x\right) d x \tag{3.60}
\end{equation*}
$$

If we complete the square in the exponent, we get

$$
\begin{align*}
-\frac{(x-\mu)^{2}}{2 \sigma^{2}}+t x & =-\frac{1}{2 \sigma^{2}}\left(x-\left[\mu+t \sigma^{2}\right]\right)^{2}+\frac{1}{2 \sigma^{2}}\left(2 \mu t \sigma^{2}+t^{2} \sigma^{4}\right) \\
& =-\frac{1}{2 \sigma^{2}}\left(x-\left[\mu+t \sigma^{2}\right]\right)^{2}+\mu t+\frac{t^{2} \sigma^{2}}{2} \tag{3.61}
\end{align*}
$$

SO

$$
\begin{align*}
M(t) & =\frac{1}{\sigma \sqrt{2 \pi}} e^{\mu t+\frac{t^{2} \sigma^{2}}{2}} \int_{-\infty}^{\infty} \exp \left(-\frac{\left(x-\left[\mu+t \sigma^{2}\right]\right)^{2}}{2 \sigma^{2}}\right) d x \\
& =\frac{1}{\sigma \sqrt{2 \pi}} e^{\mu t+\frac{t^{2} \sigma^{2}}{2}} \int_{-\infty}^{\infty} e^{-z^{2} / 2} d z=e^{\mu t+\frac{t^{2} \sigma^{2}}{2}} \tag{3.62}
\end{align*}
$$

This means that the cumulant generating function is

$$
\begin{equation*}
\psi(t)=\ln M(t)=\mu t+\frac{\sigma^{2} t^{2}}{2} \tag{3.63}
\end{equation*}
$$

taking derivatives gives

$$
\begin{equation*}
\psi^{\prime}(t)=\mu+\sigma^{2} t \tag{3.64}
\end{equation*}
$$

so the mean is

$$
\begin{equation*}
E(X)=\psi^{\prime}(0)=\mu \tag{3.65}
\end{equation*}
$$

and

$$
\begin{equation*}
\psi^{\prime \prime}(t)=\sigma^{2} t \tag{3.66}
\end{equation*}
$$

so the variance is

$$
\begin{equation*}
\operatorname{Var}(X)=\psi^{\prime \prime}(0)=\sigma^{2} \tag{3.67}
\end{equation*}
$$

which means that the parameters $\mu$ and $\sigma$ are the mean and standard deviation of the distribution, as their names suggest. Note that this is in some sense the "simplest" possible distribution with a given mean and variance. For a general random variable, since $\psi(0)=0, \psi^{\prime}(0)=E(X)$ and $\psi^{\prime \prime}(0)=\operatorname{Var}(X)$, the first few terms of the Maclaurin series for $\psi(t)$ must be

$$
\begin{equation*}
\psi(t)=t E(X)+\frac{t^{2}}{2} \operatorname{Var}(X)+\mathcal{O}\left(t^{3}\right) \tag{3.68}
\end{equation*}
$$

Given a random variable $X$ which follows a $N\left(\mu, \sigma^{2}\right)$ distribution, we can define $Z=\frac{X-\mu}{\sigma}$. Its pdf will be

$$
\begin{equation*}
f_{Z}(z)=\sigma f_{X}(\mu+z \sigma)=\frac{1}{\sqrt{2 \pi}} e^{-z^{2} / 2} \tag{3.69}
\end{equation*}
$$

which is a $N(1,0)$ distribution, also known as a standard normal distribution.

The cdf of a $N(\mu, \sigma)$ random variable will be

$$
\begin{equation*}
F(x)=\frac{1}{\sigma \sqrt{2 \pi}} \int_{-\infty}^{x} e^{-(u-\mu)^{2} /\left(2 \sigma^{2}\right)} d u=\frac{1}{\sqrt{2 \pi}} \int_{-\infty}^{(x-\mu) / \sigma} e^{-t^{2} / 2} d t \tag{3.70}
\end{equation*}
$$

again, $e^{-t^{2} / 2}$ is not the derivative of any known function, but it's useful enough that we define a function

$$
\begin{equation*}
\Phi(z)=\frac{1}{\sqrt{2 \pi}} \int_{-\infty}^{z} e^{-t^{2} / 2} d t \tag{3.71}
\end{equation*}
$$

which is tabulated in lots of places. In terms of this, the cdf for a $N\left(\mu, \sigma^{2}\right) \mathrm{rv}$ is

$$
\begin{equation*}
P(X \leq x)=\Phi\left(\frac{x-\mu}{\sigma}\right) \tag{3.72}
\end{equation*}
$$

If we add two independent Gaussian random variables, $X_{1}$ and $X_{2}$, following $N\left(\mu_{1}, \sigma_{1}^{2}\right)$ and $N\left(\mu_{2}, \sigma_{2}^{2}\right)$ distributions, respectively, their sum has the mgf

$$
\begin{align*}
M(t) & =M_{1}(t) M_{2}(t)=\exp \left(t \mu_{1}+\frac{t^{2} \sigma_{1}^{2}}{2}\right) \exp \left(t \mu_{2}+\frac{t^{2} \sigma_{2}^{2}}{2}\right) \\
& =\exp \left(t\left(\mu_{1}+\mu_{2}\right)+\frac{t^{2}\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)}{2}\right) \tag{3.73}
\end{align*}
$$

which is the mgf of a $N\left(\mu_{1}+\mu_{2}, \sigma_{1}^{2}+\sigma_{2}^{2}\right)$ distribution. (In general, if we add independent random variables, their sum has a mean
which is the sum of the means and a variance which is the sum of the variances, but what's notable here is the sum obeys a normal distribution, so it's characterized only by its mean and variance.

### 3.1.7 The Chi-Square Distribution (continuous)

Finally, suppose we have $r$ independent Gaussian random variables $\left\{X_{i}\right\}$ with means $\mu_{i}$ and variances $\sigma_{i}^{2}$. Consider the combination

$$
\begin{equation*}
Y=\sum_{i=1}^{r}\left(\frac{X_{i}-\mu_{i}}{\sigma_{i}}\right)^{2} \tag{3.74}
\end{equation*}
$$

We can show that this obeys a $\chi^{2}(r)$ distribution

$$
\begin{equation*}
f_{Y}(y)=\frac{1}{\Gamma(r / 2) 2^{r / 2}} y^{\frac{r}{2}-1} e^{-y / 2} \tag{3.75}
\end{equation*}
$$

which is a special case of a Gamma distribution with $\alpha=r$ and $\beta=2$.

First note that the sum of $r$ independent $\chi^{2}(1)$ random variables is a $\chi^{2}(r)$ random variable. This is because a $\chi^{2}(1)$ is the same as a $\operatorname{Gamma}\left(\frac{1}{2}, 2\right)$, and if we add $r$ independent $\operatorname{Gamma}\left(\frac{1}{2}, 2\right)$ random variables, we get a $\operatorname{Gamma}\left(\frac{r}{2}, 2\right)$ random variable, which is a $\chi^{2}(r)$ random variable.
Thus all we need to do to prove that $\sum_{i=1}^{r}\left(\frac{X_{i}-\mu_{i}}{\sigma_{i}}\right)^{2}$ is a $\chi^{2}(r)$ is to show that if $X$ is $N\left(\mu, \sigma^{2}\right)$, so that $Z=\frac{X-\mu}{\sigma}$ is $N(0,1)$, $Y=\frac{(X-\mu)^{2}}{\sigma}=Z^{2}$ is $\chi^{2}(1)$. Now, since

$$
\begin{equation*}
f_{Z}(z)=\frac{1}{\sqrt{2 \pi}} e^{-z^{2} / 2} \quad-\infty<z<\infty \tag{3.76}
\end{equation*}
$$

we can't quite use the usual formalism for transformation of pdfs, since the transformation $Y=Z^{2}$ is not invertible. But
since $f_{Z}(-z)=f_{Z}(z)$ it's not hard to see that if we define a rv $W=|Z|$, it must have a pdf
$f_{W}(w)=f_{Z}(-w)+f_{Z}(w)=2 f_{Z}(z)=\frac{2}{\sqrt{2 \pi}} e^{-w^{2} / 2} \quad 0<w<\infty$
and then we can use the transformation $y=w^{2}, w=y^{1 / 2}$ to work out

$$
\begin{align*}
f_{Y}(y) & =\frac{d P}{d y}=\frac{d P}{d w} \frac{d w}{d y}=\frac{1}{2} y^{-1 / 2} f_{W}\left(y^{1 / 2}\right) \\
& =\frac{1}{\sqrt{2 \pi}} y^{-1 / 2} e^{-y / 2} \quad 0<y<\infty \tag{3.78}
\end{align*}
$$

If we recall the $\chi^{2}(1) \mathrm{pdf}$, it was

$$
\begin{equation*}
f(y)=\frac{1}{\Gamma(1 / 2) 2^{1 / 2}} y^{-1 / 2} e^{-y / 2} \quad 0<y<\infty \tag{3.79}
\end{equation*}
$$

so the pdf of $Y=Z^{2}$ is the $\chi^{2}(1)$ pdf, if the value of $\Gamma(1 / 2)$ is $\sqrt{\pi}$. Now, if we think about it, that has to be the case, in order for the two pdfs to be normalized, but we can work out the value directly. Recall the Gamma function

$$
\begin{equation*}
\Gamma(\alpha)=\int_{0}^{\infty} t^{\alpha-1} e^{-t} d t \tag{3.80}
\end{equation*}
$$

which is a finite positive number for any $\alpha>0$. (For positive integer $n$, we know that $\Gamma(n)=(n-1)$ !.) Thus

$$
\begin{equation*}
\Gamma(1 / 2)=\int_{0}^{\infty} t^{-1 / 2} e^{-t} d t \tag{3.81}
\end{equation*}
$$

We can show that the integral is well-behaved at the lower limit of $t=0$, and evaluate it, by changing variables to $u=\sqrt{2 t}$ so that $t=u^{2} / 2$ and $d u=2^{1 / 2} t^{-1 / 2} d t$; thus

$$
\begin{equation*}
\Gamma(1 / 2)=\frac{1}{\sqrt{2}} \int_{0}^{\infty} e^{-u^{2}} d u=\frac{1}{\sqrt{2}} \frac{\sqrt{2 \pi}}{2}=\sqrt{\pi} \tag{3.82}
\end{equation*}
$$

as expected. We've used the symmetry of the integrand to say that

$$
\begin{equation*}
\int_{0}^{\infty} e^{-u^{2}} d u=\frac{1}{2} \int_{-\infty}^{\infty} e^{-u^{2}} d u=\frac{\sqrt{2 \pi}}{2} \tag{3.83}
\end{equation*}
$$

### 3.1.8 Summary of Properties of Gamma Distribution

- If $X \sim \operatorname{Gamma}(\alpha, \beta)$, then $E[X]=\alpha \beta$ and $\operatorname{Var}(X)=$ $\alpha \beta^{2}$. The shape parameter $\alpha$ is dimensionless and the scale parameter $\beta$ has the same units as $X$.
- If we add independent random variables, a $\operatorname{Gamma}\left(\alpha_{1}, \beta\right)$ and a $\operatorname{Gamma}\left(\alpha_{2}, \beta\right)$, their sum is a $\operatorname{Gamma}\left(\alpha_{1}+\alpha_{2}, \beta\right)$.
- If $X$ is a $\operatorname{Gamma}\left(\alpha_{1}, \beta\right), a X$ is a $\operatorname{Gamma}\left(\alpha_{1}, a \beta\right) \square^{8}$
- The sum of the squares of $r$ standard normal random variables is a $\chi^{2}(r)$, which is the same as a $\operatorname{Gamma}\left(\frac{r}{2}, 2\right)$.
- The waiting time for the next event in a Poisson process with rate $\lambda$ is an $\operatorname{Exp}(\lambda)$ (a random variable obeying an exponential distribution), which is a $\Gamma\left(1, \frac{1}{\lambda}\right)$.
- The waiting time for the $n$th event in a Poisson process with rate $\lambda$ is a $\Gamma\left(n, \frac{1}{\lambda}\right)$. (This special case of a Gamma distribution with integer $\alpha$ is also called an Erlang distribution.)


## 4 Sums of Random Variables

### 4.1 Mean and Variance

Consider a situation where there are two random variables $X_{1}$ and $X_{2}$, and we construct a new random variable which is their

[^6]sum,
\[

$$
\begin{equation*}
T=X_{1}+X_{2} \tag{4.1}
\end{equation*}
$$

\]

If the expectation values of the random variables are

$$
\begin{equation*}
\mu_{1}=E\left[X_{1}\right] \quad \text { and } \quad \mu_{2}=E\left[X_{2}\right] \tag{4.2}
\end{equation*}
$$

then the linearity of the expectation value operation means that the expectation value of their sum is

$$
\begin{equation*}
\mu_{T}=E[T]=E\left[X_{1}\right]+E\left[X_{2}\right]=\mu_{1}+\mu_{2} \tag{4.3}
\end{equation*}
$$

If the random variables ave standard deviations $\sigma_{1}$ and $\sigma_{2}$ and covariance $\operatorname{Cov}\left(X_{1}, X_{2}\right)$, so that

$$
\begin{align*}
E\left[\left(X_{1}-\mu_{1}\right)^{2}\right] & =\sigma_{1}^{2}  \tag{4.4a}\\
E\left[\left(X_{2}-\mu_{2}\right)^{2}\right] & =\sigma_{2}^{2}  \tag{4.4b}\\
E\left[\left(X_{1}-\mu_{1}\right)\left(X_{2}-\mu_{2}\right)\right] & =\operatorname{Cov}\left(X_{1}, X_{2}\right) \tag{4.4c}
\end{align*}
$$

then the variance of their sum is

$$
\begin{align*}
\sigma_{T}^{2} & =E\left[\left(T-\mu_{T}\right)^{2}\right]=E\left[\left(X_{1}+X_{2}-\mu_{1}-\mu_{2}\right)^{2}\right] \\
& =E\left[\left(\left[X_{1}-\mu_{1}\right]+\left[X_{2}-\mu_{2}\right]\right)^{2}\right] \\
& =E\left[\left(X_{1}-\mu_{1}\right)^{2}\right]+2 E\left[\left(X_{1}-\mu_{1}\right)\left(X_{2}-\mu_{2}\right)\right]+E\left[\left(X_{2}-\mu_{2}\right)^{2}\right] \\
& =\sigma_{1}^{2}+\sigma_{2}^{2}+2 \operatorname{Cov}\left(X_{1}, X_{2}\right) \tag{4.5}
\end{align*}
$$

In particular, if $X_{1}$ and $X_{2}$ are independent or otherwise uncorrelated, then the variance of their sum is equal to the sum of their variances:

$$
\begin{equation*}
\sigma_{T}^{2}=\sigma_{1}^{2}+\sigma_{2}^{2} \quad \text { if } X_{1} \text { and } X_{2} \text { uncorrelated } \tag{4.6}
\end{equation*}
$$

Note that this means the standard deviations are "added in quadrature":

$$
\begin{equation*}
\sigma_{T}=\sqrt{\sigma_{1}^{2}+\sigma_{2}^{2}} \quad \text { if } X_{1} \text { and } X_{2} \text { uncorrelated } \tag{4.7}
\end{equation*}
$$

This is the standard way in which uncorrelated random errors are combined.

Note that in the special case where $X_{1}$ and $X_{2}$ are independent Gaussian random variables, we can show their sum is also a Gaussian random variable. This is pretty quick to show using the moment generating functions, since in this case

$$
\begin{align*}
M(t) & =E\left[e^{X_{1}+X_{2}}\right]=E\left[e^{X_{1}}\right] E\left[e^{X_{2}}\right]=M_{1}(t) M_{2}(t) \\
& =\exp \left(t \mu_{1}+\frac{1}{2} t^{2} \sigma_{1}^{2}\right) \exp \left(t \mu_{2}+\frac{1}{2} t^{2} \sigma_{2}^{2}\right)  \tag{4.8}\\
& =\exp \left(t\left(\mu_{1}+\mu_{2}\right)+\frac{1}{2} t^{2}\left(\sigma_{1}^{2}+\sigma_{2}^{2}\right)\right)
\end{align*}
$$

Note also that this is a special case of a general linear combination of random variables

$$
\begin{equation*}
Y=\sum_{i=1}^{n} a_{i} X_{i} \tag{4.9}
\end{equation*}
$$

for which

$$
\begin{equation*}
E[Y]=\sum_{i=1}^{n} a_{i} E\left[X_{i}\right] \tag{4.10}
\end{equation*}
$$

and

$$
\begin{equation*}
\operatorname{Var}(Y)=\sum_{i=1}^{n} \sum_{j=1}^{n} a_{i} a_{j} \operatorname{Cov}\left(X_{i}, X_{j}\right) \tag{4.11}
\end{equation*}
$$

For example,
$\operatorname{Var}\left(a_{1} X_{1}+a_{2} X_{2}\right)=a_{1}{ }^{2} V\left(X_{1}\right)+2 a_{1} a_{2} \operatorname{Cov}\left(X_{1}, X_{2}\right)+a_{2}{ }^{2} V\left(X_{2}\right)$

### 4.2 IID Random Variables (Random Samples)

Considered now the example of $N$ independent, identically distributed (iid) random variables $\left\{X_{i}\right\}$ with expectation values

$$
\begin{equation*}
E\left[X_{i}\right]=\mu \quad \text { and } \quad E\left[\left(X_{i}-\mu\right)\left(X_{j}-\mu\right)\right]=\delta_{i j} \sigma^{2} \tag{4.13}
\end{equation*}
$$

This is known as a random sample, and it can be used to estimate the properties of the underlying distribution. If we construct the sum

$$
\begin{equation*}
T=\sum_{i=1}^{N} X_{i} \tag{4.14}
\end{equation*}
$$

then an extension of the results for a pair of random variables shows that its mean is

$$
\begin{equation*}
\mu_{T}=E[T]=\sum_{i=1}^{N} \mu=N \mu \tag{4.15}
\end{equation*}
$$

and its variance is

$$
\begin{equation*}
\sigma_{T}^{2}=E\left[\left(T-\mu_{T}\right)^{2}\right]=\sum_{i=1}^{N} \sigma^{2}=N \sigma^{2} \tag{4.16}
\end{equation*}
$$

so its standard deviation is

$$
\begin{equation*}
\sigma_{T}=\sqrt{N} \sigma \tag{4.17}
\end{equation*}
$$

If we take the average of the $N$ random variables

$$
\begin{equation*}
\bar{X}=\frac{1}{N} \sum_{k=0}^{N-1} X_{i}=\frac{T}{N} \tag{4.18}
\end{equation*}
$$

which is itself a random variable, we can see

$$
\begin{equation*}
\mu_{\bar{X}}=E[\bar{X}]=\frac{E[T]}{N}=\mu \tag{4.19}
\end{equation*}
$$

and

$$
\begin{equation*}
\sigma_{\bar{X}}^{2}=\operatorname{Var}(\bar{X})=\operatorname{Var}\left(\frac{T}{N}\right)=\frac{\operatorname{Var}(T)}{N^{2}}=\frac{\sigma_{T}^{2}}{N^{2}}=\frac{N \sigma^{2}}{N} \tag{4.20}
\end{equation*}
$$

which means

$$
\begin{equation*}
\sigma_{\bar{X}}=\frac{\sigma}{\sqrt{N}} \tag{4.21}
\end{equation*}
$$

I.e., if you take the average of $N$ iid random variables, the standard deviation is $1 / \sqrt{N}$ times theie individual standard deviation.

## Thursday, October 12, 2017

### 4.3 PDF of a Sum of Random Variables

If we consider two independent random variables $X_{1}$ and $X_{2}$ with (not necessarily identical) pdfs $f_{1}\left(x_{1}\right)$ and $f_{2}\left(x_{2}\right)$, so that their joint pdf is

$$
\begin{equation*}
f\left(x_{1}, x_{2}\right)=f_{1}\left(x_{1}\right) f_{2}\left(x_{2}\right) \tag{4.22}
\end{equation*}
$$

and write their sum as

$$
\begin{equation*}
T=X_{1}+X_{2} \tag{4.23}
\end{equation*}
$$

we can ask what the pdf $f_{T}(t)$ is. One way to approach this ${ }^{9}$ is to consider the joint pdf $p\left(t, x_{2}\right)$. We can do this by changing variables from $x_{1}$ to $t=x_{1}+x_{2}$; since we're actually changing

[^7]from $\left\{x_{1}, x_{2}\right\}$ to $\left\{t, x_{2}\right\}$, we can treat $x_{2}$ as a constant ${ }^{10}$ and since $d t=d x_{1}$ in that case,
\[

$$
\begin{equation*}
f\left(t, x_{2}\right)=f_{1}\left(t-x_{2}\right) f_{2}\left(x_{2}\right) \tag{4.24}
\end{equation*}
$$

\]

if we then marginalize over $x_{2}$, we find

$$
\begin{equation*}
f(t)=\int_{-\infty}^{\infty} f\left(t-x_{2}\right) f_{2}\left(x_{2}\right) d x_{2} \tag{4.25}
\end{equation*}
$$

and so we see that the pdf of a sum of variables is the convolution of their individual pdfs.

Of course, we also know that when we add independent random variables, we multiply their mgfs. These things all tie together when you recall that the mgf is closely related to the characteristic function, which is the Fourier transform of the pdf. This means we're looking at an example of the convolution theorem: when you convolve pdfs, that's the same thing as multiplying their Fourier transforms.

### 4.4 The Central Limit Theorem

The central limit theorem states that if $\left\{X_{i}\right\}$ is a sample of size $n$ drawn from a distribution with mean $E\left(X_{i}\right)=\mu$ and variance $\operatorname{Var}\left(X_{i}\right)=\sigma^{2}$, then

$$
\begin{equation*}
Z_{n}=\frac{\sum_{i=1}^{n} X_{i}-n \mu}{\sigma \sqrt{n}}=\frac{\bar{X}-\mu}{\sigma / \sqrt{n}} \longrightarrow N(0,1) \tag{4.26}
\end{equation*}
$$

${ }^{10}$ Alternatively, we can consider the Jacobian determinant

$$
\left\|\frac{\partial\left(t, x_{2}\right)}{\partial\left(x_{1}, x_{2}\right)}\right\|=\left|\operatorname{det}\left(\begin{array}{ll}
\left(\frac{\partial t}{\partial x_{1}}\right)_{x_{2}} & \left(\frac{\partial t}{\partial x_{1}}\right)_{x_{1}} \\
\left(\frac{\partial x_{1}}{\partial x_{1}}\right)_{x_{2}} & \left(\frac{\partial x_{1}}{\partial x_{1}}\right)_{x_{1}}
\end{array}\right)\right|=\left|\operatorname{det}\left(\begin{array}{ll}
1 & 1 \\
1 & 0
\end{array}\right)\right|=1
$$

We can prove this using the moment generating function. I'll show this using the cumulant generating function

$$
\begin{equation*}
\psi(t)=\ln E\left(\exp \left[t X_{i}\right]\right) \tag{4.27}
\end{equation*}
$$

for the distribution from which the sample is drawn. We know that $\psi(0)=0, \psi^{\prime}(0)=E\left(X_{i}\right)=\mu$, and $\psi^{\prime \prime}(0)=\operatorname{Var}\left(X_{i}\right)=\sigma^{2}$, so we can write the truncated McLaurin series for $\psi(t)$ as
$\psi(t)=\psi(0)+t \psi^{\prime}(0)+\frac{t^{2}}{2} \psi^{\prime \prime}(0)+o\left(t^{2}\right)=\psi(0)+\mu t+\frac{\sigma^{2} t^{2}}{2}+o\left(t^{2}\right)$
where $o\left(t^{2}\right)$ is some expression such that $\lim _{t \rightarrow 0} \frac{o\left(t^{2}\right)}{t^{2}} \rightarrow 0$; we also sometimes write this as $\mathcal{O}\left(t^{3}\right)$. We can write that more precisely as
$\psi(t)=\psi(0)+t \psi^{\prime}(0)+\frac{t^{2}}{2} \psi^{\prime \prime}(\xi(t))=\mu t+\frac{\sigma^{2} t^{2}}{2}+\frac{\psi^{\prime \prime}(\xi(t))-\sigma^{2}}{2} t^{2}$
where $\xi(t)$ is some number between 0 and $t$. Since

$$
\lim _{t \rightarrow 0} \psi^{\prime \prime}(\xi(t))=\psi^{\prime \prime}(0)=\sigma^{2}
$$

the last term is indeed $o\left(t^{2}\right)$. Now we consider the cumulant generating function

$$
\begin{align*}
\Psi(t ; n) & =\ln E\left(\exp \left[t Z_{n}\right]\right)=\ln E\left(\exp \left[\frac{t}{\sigma \sqrt{n}}\left\{\sum_{i=1}^{n}\left(X_{i}-\mu\right)\right\}\right]\right) \\
& =\ln \prod_{i=1}^{n} E\left(\exp \left[\frac{t}{\sigma \sqrt{n}}\left\{X_{i}-\mu\right\}\right]\right) \\
& =\sum_{i=1}^{n} \ln E\left(\exp \left[\frac{t}{\sigma \sqrt{n}}\left\{X_{i}-\mu\right\}\right]\right) \\
& =n\left[\psi\left(\frac{t}{\sigma \sqrt{n}}\right)-\frac{t}{\sigma \sqrt{n}} \mu\right] \tag{4.31}
\end{align*}
$$

If we use the Taylor expansion (4.29) we have
$\Psi(t ; n)=n\left[\frac{1}{2}\left(\frac{t}{\sigma \sqrt{n}}\right)^{2} \psi^{\prime \prime}\left(\xi\left(\frac{t}{\sigma \sqrt{n}}\right)\right)\right]=\frac{t^{2}}{2 \sigma^{2}} \psi^{\prime \prime}\left(\xi\left(\frac{t}{\sigma \sqrt{n}}\right)\right)$
When we take the limit $n \rightarrow \infty$, the argument of $\psi^{\prime \prime}()$ goes to zero, and thus

$$
\begin{equation*}
\lim _{n \rightarrow \infty} \Psi(t ; n)=\frac{t^{2}}{2 \sigma^{2}} \psi^{\prime \prime}(0)=\frac{t^{2}}{2} \tag{4.33}
\end{equation*}
$$

which is the natural log of the mgf of a $N(0,1)$ random variable, so by the mgf method we've shown $Z_{n} \longrightarrow N(0,1)$.

## 5 Multivariate Normal Distribution

### 5.1 Linear Algebra: Reminders and Notation

If $\mathbf{A}$ is an $m \times n$ matrix:

$$
\mathbf{A}=\left(\begin{array}{cccc}
A_{11} & A_{12} & \cdots & A_{1 n}  \tag{5.1}\\
A_{21} & A_{22} & \cdots & A_{2 n} \\
\vdots & \vdots & \ddots & \vdots \\
A_{m 1} & A_{m 2} & \cdots & A_{m n}
\end{array}\right)
$$

and $\mathbf{B}$ is an $n \times p$ matrix,

$$
\mathbf{B}=\left(\begin{array}{cccc}
B_{11} & B_{12} & \cdots & B_{1 p}  \tag{5.2}\\
B_{21} & B_{22} & \cdots & B_{2 p} \\
\vdots & \vdots & \ddots & \vdots \\
B_{n 1} & B_{n 2} & \cdots & B_{n p}
\end{array}\right)
$$

then their product $\mathbf{C}=\mathbf{A B}$ is an $m \times p$ matrix as shown in Figure 1 so that $C_{i k}=\sum_{j=1}^{n} A_{i j} B_{j k}$.

If $\mathbf{A}$ is an $m \times n$ matrix, $\mathbf{B}=\mathbf{A}^{\mathrm{T}}$ is an $n \times m$ matrix with elements $B_{i j}=A_{j i}$ :

$$
\left(\begin{array}{cccc}
B_{11} & B_{12} & \cdots & B_{1 m}  \tag{5.4}\\
B_{21} & B_{22} & \cdots & B_{2 m} \\
\vdots & \vdots & \ddots & \vdots \\
B_{n 1} & B_{n 2} & \cdots & B_{n m}
\end{array}\right)=\mathbf{B}=\mathbf{A}^{\mathrm{T}}=\left(\begin{array}{cccc}
A_{11} & A_{21} & \cdots & A_{n 1} \\
A_{12} & A_{22} & \cdots & A_{n 2} \\
\vdots & \vdots & \ddots & \vdots \\
A_{1 n} & A_{2 n} & \cdots & A_{n m}
\end{array}\right)
$$

If $\mathbf{v}$ is an $n$-element column vector (which is an $n \times 1$ matrix) and $\mathbf{A}$ is an $m \times n$ matrix, $\mathbf{w}=\mathbf{A v}$ is an $m$-element column vector (i.e., an $m \times 1$ matrix):

$$
\begin{align*}
\left(\begin{array}{c}
w_{1} \\
w_{2} \\
\vdots \\
w_{m}
\end{array}\right)= & \mathbf{w}=\mathbf{A v}=\left(\begin{array}{cccc}
A_{11} & A_{12} & \cdots & A_{1 n} \\
A_{21} & A_{22} & \cdots & A_{2 n} \\
\vdots & \vdots & \ddots & \vdots \\
A_{m 1} & A_{m 2} & \cdots & A_{m n}
\end{array}\right)\left(\begin{array}{c}
v_{1} \\
v_{2} \\
\vdots \\
v_{n}
\end{array}\right)  \tag{5.5}\\
= & \left(\begin{array}{c}
A_{11} v_{1}+A_{12} v_{2}+\cdots+A_{1 n} v_{n} \\
A_{21} v_{1}+A_{22} v_{2}+\cdots+A_{2 n} v_{n} \\
\vdots \\
A_{m 1} v_{1}+A_{m 2} v_{2}+\cdots+A_{m n} v_{n}
\end{array}\right)
\end{align*}
$$

so that $w_{i}=\sum_{j=1}^{n} A_{i j} v_{j}$.
If $\mathbf{u}$ is an $n$-element column vector, then $\mathbf{u}^{\mathrm{T}}$ is an $n$-element row vector (a $1 \times n$ matrix):

$$
\mathbf{u}^{\mathrm{T}}=\left(\begin{array}{c}
u_{1}  \tag{5.6}\\
u_{2} \\
\vdots \\
u_{n}
\end{array}\right)
$$

If $\mathbf{u}$ and $\mathbf{v}$ are $n$-element column vectors, $\mathbf{u}^{\mathrm{T}} \mathbf{v}$ is a number,

$$
\begin{align*}
\mathbf{C}= & \mathbf{A B}=\left(\begin{array}{cccc}
C_{11} & C_{12} & \cdots & C_{1 p} \\
C_{21} & C_{22} & \cdots & C_{2 p} \\
\vdots & \vdots & \ddots & \vdots \\
C_{m 1} & C_{m 2} & \cdots & C_{m p}
\end{array}\right)=\left(\begin{array}{cccc}
A_{11} & A_{12} & \cdots & A_{1 n} \\
A_{21} & A_{22} & \cdots & A_{2 n} \\
\vdots & \vdots & \ddots & \vdots \\
A_{m 1} & A_{m 2} & \cdots & A_{m n}
\end{array}\right)\left(\begin{array}{cccc}
B_{11} & B_{12} & \cdots & B_{1 p} \\
B_{21} & B_{22} & \cdots & B_{2 p} \\
\vdots & \vdots & \ddots & \vdots \\
B_{n 1} & B_{n 2} & \cdots & B_{n p}
\end{array}\right)  \tag{5.3}\\
& =\left(\begin{array}{cccc|}
A_{11} B_{11}+A_{12} B_{21}+\cdots+A_{1 n} B_{n 1} & A_{11} B_{12}+A_{12} B_{22}+\cdots+A_{1 n} B_{n 2} & \cdots & A_{11} B_{1 p}+A_{12} B_{2 p}+\cdots+A_{1 n} B_{n p} \\
A_{21} B_{11}+A_{22} B_{21}+\cdots+A_{2 n} B_{n 1} & A_{21} B_{12}+A_{22} B_{22}+\cdots+A_{2 n} B_{n 2} & \cdots & A_{21} B_{1 p}+A_{22} B_{2 p}+\cdots+A_{2 n} B_{n p} \\
\vdots & \vdots & \ddots & \vdots \\
A_{m 1} B_{11}+A_{m 2} B_{21}+\cdots+A_{m n} B_{n 1} & A_{m 1} B_{12}+A_{m 2} B_{22}+\cdots+A_{m n} B_{n 2} & \cdots & A_{m 1} B_{1 p}+A_{m 2} B_{2 p}+\cdots+A_{m n} B_{n p}
\end{array}\right)
\end{align*}
$$

Figure 1: Expansion of the product $\mathbf{C}=\mathbf{A B}$ to show $C_{i k}=\sum_{j=1}^{n} A_{i j} B_{j k}$.
known as the inner product:

$$
\begin{align*}
\mathbf{u}^{\mathrm{T}} \mathbf{v} & =\left(\begin{array}{llll}
u_{1} & u_{2} & \vdots & u_{n}
\end{array}\right)\left(\begin{array}{c}
v_{1} \\
v_{2} \\
\cdots \\
v_{n}
\end{array}\right)  \tag{5.7}\\
& =u_{1} v_{1}+u_{2} v_{2}+\cdots+u_{n} v_{n}=\sum_{i=1}^{n} u_{i} v_{i}
\end{align*}
$$

If $\mathbf{v}$ is an $m$-element column vector, and $\mathbf{w}$ is an $n$-element
column vector, $\mathbf{A}=\mathbf{v w}^{\mathrm{T}}$ is an $m \times n$ matrix

$$
\begin{align*}
& \left(\begin{array}{cccc}
A_{11} & A_{12} & \cdots & A_{1 n} \\
A_{21} & A_{22} & \cdots & A_{2 n} \\
\vdots & \vdots & \ddots & \vdots \\
A_{m 1} & A_{m 2} & \cdots & A_{m n}
\end{array}\right)=\mathbf{A}=\mathbf{v w}^{\mathrm{T}} \\
& =\left(\begin{array}{c}
v_{1} \\
v_{2} \\
\vdots \\
v_{m}
\end{array}\right)\left(\begin{array}{llll}
w_{1} & w_{2} & \cdots & w_{m}
\end{array}\right)=\left(\begin{array}{cccc}
v_{1} w_{1} & v_{1} w_{2} & \cdots & v_{1} w_{n} \\
v_{2} w_{1} & v_{2} w_{2} & \cdots & v_{2} w_{n} \\
\vdots & \vdots & \ddots & \vdots \\
v_{m} w_{1} & v_{m} w_{2} & \cdots & v_{m} w_{n}
\end{array}\right) \tag{5.8}
\end{align*}
$$

so that $A_{i j}=v_{i} w_{j}$.
If $\mathbf{M}$ and $\mathbf{N}$ are $n \times n$ matrices, the determinant $\operatorname{det}(\mathbf{M N})=$ $\operatorname{det}(\mathbf{M}) \operatorname{det}(\mathbf{N})$.
If $\mathbf{M}$ is an $n \times n$ matrix (known as a square matrix), the inverse matrix $\mathbf{M}^{-1}$ is defined by $\mathbf{M}^{-1} \mathbf{M}=\mathbf{1}_{n \times n}=\mathbf{M} \mathbf{M}^{-1}$ where $\mathbf{1}_{n \times n}$
is the identity matrix

$$
\mathbf{1}_{n \times n}=\left(\begin{array}{cccc}
1 & 0 & \cdots & 0  \tag{5.9}\\
0 & 1 & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & 1
\end{array}\right)
$$

If $\mathbf{M}^{-1}$ exists, we say $\mathbf{M}$ is invertible.
If $\mathbf{M}$ is a real, symmetric $n \times n$ matrix, so that $\mathbf{M}^{\mathrm{T}}=\mathbf{M}$, i.e., $M_{j i}=M_{i j}$, there is a set of $n$ orthonormal eigenvectors $\left\{\mathbf{v}_{1}, \mathbf{v}_{2}, \ldots, \mathbf{v}_{n}\right\}$ with real eigenvalues $\left\{\lambda_{1}, \lambda_{2}, \ldots, \lambda_{n}\right\}$, so that $\mathbf{M v}_{i}=\lambda_{i} \mathbf{v}_{i}$. Orthonormal means

$$
\mathbf{v}_{i}^{\mathrm{T}} \mathbf{v}_{j}=\delta_{i j}= \begin{cases}0 & i \neq j  \tag{5.10}\\ 1 & i=j\end{cases}
$$

where we have introduced the Kronecker delta symbol $\delta_{i j}$. The eigenvalue decomposition means

$$
\begin{equation*}
\mathbf{M}=\sum_{i=1}^{n} \lambda_{i} \mathbf{v}_{i} \mathbf{v}_{i}^{\mathrm{T}} \tag{5.11}
\end{equation*}
$$

The determinant is $\operatorname{det}(\mathbf{M})=\prod_{i=1}^{n} \lambda_{i}$. If none of the eigenvalues $\left\{\lambda_{i}\right\}$ are zero, $\mathbf{M}$ is invertible, and the inverse matrix is

$$
\begin{equation*}
\mathbf{M}^{-1}=\sum_{i=1}^{n} \frac{1}{\lambda_{i}} \mathbf{v}_{i} \mathbf{v}_{i}^{\mathrm{T}} \tag{5.12}
\end{equation*}
$$

If all of the eigenvalues $\left\{\lambda_{i}\right\}$ are positive, we say $\mathbf{M}$ is positive definite. If none of the eigenvalues $\left\{\lambda_{i}\right\}$ are negative, we say $\mathbf{M}$ is positive semi-definite.

## Tuesday, October 17, 2017

### 5.2 Special Case: Independent Gaussian Random Variables

Before considering the general multivariate normal distribution, consider the case of $n$ independent normally-distributed random variables $\left\{X_{i}\right\}$ with means $\left\{\mu_{i}\right\}$ and variances $\left\{\sigma_{i}\right\}$. The pdf for $X_{i}$, the $i$ th random variable, is

$$
\begin{equation*}
f_{i}\left(x_{i}\right)=\frac{1}{\sigma_{i} \sqrt{2 \pi}} \exp \left(-\frac{\left(x_{i}-\mu_{i}\right)^{2}}{2 \sigma_{i}^{2}}\right) \tag{5.13}
\end{equation*}
$$

and its mgf is

$$
\begin{equation*}
M_{i}\left(t_{i}\right)=\exp \left(t_{i} \mu_{i}+\frac{1}{2} t_{i}^{2} \sigma_{i}^{2}\right) \tag{5.14}
\end{equation*}
$$

If we consider the random variables $X_{i}$ to be the elements of a random vector

$$
\mathbf{X}=\left(\begin{array}{c}
X_{1}  \tag{5.15}\\
X_{2} \\
\vdots \\
X_{n}
\end{array}\right)
$$

its expectation value is

$$
E(\mathbf{X})=\boldsymbol{\mu}=\left(\begin{array}{c}
\mu_{1}  \tag{5.16}\\
\mu_{2} \\
\vdots \\
\mu_{n}
\end{array}\right)
$$

and its variance-covariance matrix is

$$
\operatorname{Cov}(\mathbf{X})=\boldsymbol{\sigma}^{2}=\left(\begin{array}{cccc}
\sigma_{1}^{2} & 0 & \cdots & 0  \tag{5.17}\\
0 & \sigma_{2}^{2} & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & \sigma_{n}^{2}
\end{array}\right)
$$

which is diagonal because the different $X_{i} \mathrm{~S}$ are independent of each other and therefore have zero covariance. We can thus write the joint mgf for these random variables as

$$
\begin{align*}
M(\mathbf{t}) & =\prod_{i=1}^{n} M_{i}\left(t_{i}\right)=\exp \left(\sum_{i=1}^{n}\left[t_{i} \mu_{i}+\frac{1}{2} t_{i}^{2} \sigma_{i}^{2}\right]\right)  \tag{5.18}\\
& =\exp \left(\mathbf{t}^{\mathrm{T}} \boldsymbol{\mu}+\frac{1}{2} \mathbf{t}^{\mathrm{T}} \boldsymbol{\sigma}^{2} \mathbf{t}\right)
\end{align*}
$$

We can also write the joint pdf

$$
\begin{equation*}
f(\mathbf{x})=\prod_{i=1}^{n} f_{i}\left(x_{i}\right)=\frac{1}{\sqrt{\prod_{i=1}^{n} 2 \pi \sigma_{i}^{2}}} \exp \left(-\frac{1}{2} \sum_{i=1}^{n} \frac{\left(x_{i}-\mu_{i}\right)^{2}}{\sigma_{i}^{2}}\right) \tag{5.19}
\end{equation*}
$$

in matrix form if we consider a few operations on the matrix $\boldsymbol{\sigma}^{2}$. First, since it's a diagonal matrix, its determinant is just the product of its diagonal entries:

$$
\begin{equation*}
\operatorname{det} \boldsymbol{\sigma}^{2}=\prod_{i=1}^{n} \sigma_{i}^{2} \tag{5.20}
\end{equation*}
$$

and, for that matter,

$$
\begin{equation*}
\operatorname{det}\left(2 \pi \boldsymbol{\sigma}^{2}\right)=\prod_{i=1}^{n} 2 \pi \sigma_{i}^{2} \tag{5.21}
\end{equation*}
$$

Also, we can invert the matrix to get

$$
\boldsymbol{\sigma}^{-2}=\left(\begin{array}{cccc}
\frac{1}{\sigma_{1}^{2}} & 0 & \cdots & 0  \tag{5.22}\\
0 & \frac{1}{\sigma_{2}^{2}} & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & \frac{1}{\sigma_{n}^{2}}
\end{array}\right)
$$

so

$$
\begin{equation*}
\sum_{i=1}^{n} \frac{\left(x_{i}-\mu_{i}\right)^{2}}{\sigma_{i}^{2}}=(\mathbf{x}-\boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\sigma}^{-2}(\mathbf{x}-\boldsymbol{\mu}) \tag{5.23}
\end{equation*}
$$

which makes the pdf for the random vector $X$

$$
\begin{equation*}
f(\mathbf{x})=\frac{1}{\sqrt{\operatorname{det}\left(2 \pi \boldsymbol{\sigma}^{2}\right)}} \exp \left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\sigma}^{-2}(\mathbf{x}-\boldsymbol{\mu})\right) \tag{5.24}
\end{equation*}
$$

The generalization from $n$ independent normal random variables to an $n$-dimensional multivariate normal distribution is to use the same matrix form for $M(\mathbf{t})$ and just to replace $\boldsymbol{\sigma}^{2}$, which was a diagonal matrix with positive diagonal entries, with a general symmetric positive semi-definite matrix. So one change is to allow $\sigma^{2}$ to have off-diagonal entries, and another is to allow it to have zero eigenvalues. If $\sigma^{2}$ is positive definite, i.e., its eigenvalues are all positive, we can use the matrix expression for the $\operatorname{pdf} f(\mathbf{x})$ as well. As we'll see, if $\boldsymbol{\sigma}^{2}$ has some zero eigenvalues, we won't be able to define a pdf for the random vector $\mathbf{X}$.

### 5.3 Multivariate Distributions

Recall that a set of $n$ random variables $X_{1}, X_{2}, \ldots X_{n}$ can be combined into a random vector

$$
\mathbf{X}=\left(\begin{array}{c}
X_{1}  \tag{5.25}\\
X_{2} \\
\vdots \\
X_{n}
\end{array}\right)
$$

with expectation value

$$
E(\mathbf{X})=\boldsymbol{\mu}=\left(\begin{array}{c}
\mu_{1}  \tag{5.26}\\
\mu_{2} \\
\vdots \\
\mu_{n}
\end{array}\right)
$$

and variance-covariance matrix

$$
\begin{align*}
\boldsymbol{\sigma}^{2}= & \operatorname{Cov}(\mathbf{X})=E\left([\mathbf{X}-\boldsymbol{\mu}]^{\mathrm{T}}[\mathbf{X}-\boldsymbol{\mu}]\right) \\
& =\left(\begin{array}{cccc}
\operatorname{Var}\left(X_{1}\right) & \operatorname{Cov}\left(X_{1}, X_{2}\right) & \cdots & \operatorname{Cov}\left(X_{1}, X_{n}\right) \\
\operatorname{Cov}\left(X_{1}, X_{2}\right) & \operatorname{Var}\left(X_{2}\right) & \cdots & \operatorname{Cov}\left(X_{2}, X_{n}\right) \\
\vdots & \vdots & \ddots & \vdots \\
\operatorname{Cov}\left(X_{1}, X_{n}\right) & \operatorname{Cov}\left(X_{2}, X_{n}\right) & \cdots & \operatorname{Var}\left(X_{n}\right)
\end{array}\right) \tag{5.27}
\end{align*}
$$

The variance-covariance matrix must be positive semi-definite, i.e., have no negative eigenvalues. To see why that is the case, let $\left\{\lambda_{i}\right\}$ be the eigenvalues and $\left\{\mathbf{v}_{i}\right\}$ be the orthonormal eigenvectors, so that $\boldsymbol{\sigma}^{2}=\sum_{i=1}^{n} \mathbf{v}_{i} \lambda_{i} \mathbf{v}^{\mathrm{T}}{ }_{i}$. For each $i$ define a random variable $\mathcal{X}_{i}=\mathbf{v}_{i}^{\mathrm{T}} \mathbf{X}$. It has mean $E\left(\mathcal{X}_{i}\right)=\mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}$ and variance

$$
\begin{align*}
\operatorname{Var}\left(\mathcal{X}_{i}\right) & =E\left(\left[\mathbf{v}_{i}^{\mathrm{T}}(\mathbf{X}-\boldsymbol{\mu})\right]^{2}\right)=E\left[\mathbf{v}_{i}^{\mathrm{T}}(\mathbf{X}-\boldsymbol{\mu})(\mathbf{X}-\boldsymbol{\mu})^{\mathrm{T}} \mathbf{v}_{i}\right] \\
& =\mathbf{v}_{i}^{\mathrm{T}} E\left[(\mathbf{X}-\boldsymbol{\mu})(\mathbf{X}-\boldsymbol{\mu})^{\mathrm{T}}\right] \mathbf{v}_{i}=\mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\sigma}^{2} \mathbf{v}_{i}=\lambda_{i} \mathbf{v}_{i}^{\mathrm{T}} \mathbf{v}_{i}=\lambda_{i} \tag{5.28}
\end{align*}
$$

Since the variance of a random variable must be non-negative, $\sigma^{2}$ cannot have any negative eigenvalues.

Incidentally, we can see that the different random variables $\left\{\mathcal{X}_{i}\right\}$ are uncorrelated, since

$$
\begin{align*}
& \operatorname{Cov}\left(\mathcal{X}_{i}, \mathcal{X}_{j}\right)=E\left[\mathbf{v}_{i}^{\mathrm{T}}(\mathbf{X}-\boldsymbol{\mu})(\mathbf{X}-\boldsymbol{\mu})^{\mathrm{T}} \mathbf{v}_{j}\right] \\
& \quad=\mathbf{v}_{i}^{\mathrm{T}} E\left[(\mathbf{X}-\boldsymbol{\mu})(\mathbf{X}-\boldsymbol{\mu})^{\mathrm{T}}\right] \mathbf{v}_{j}=\mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\sigma}^{2} \mathbf{v}_{j}=\lambda_{j} \mathbf{v}_{i}^{\mathrm{T}} \mathbf{v}_{j}=\lambda_{i} \delta_{i j} \tag{5.29}
\end{align*}
$$

Remember that $\operatorname{Cov}\left(\mathcal{X}_{i}, \mathcal{X}_{j}\right)=0$ does not necessarily imply that $\mathcal{X}_{i}$ and $\mathcal{X}_{j}$ are independent. (It will, however, turn out to be the case for normally distributed random variables.)

Note that if we assemble the $\left\{\mathcal{X}_{i}\right\}$ into a column vector

$$
\mathcal{X}=\left(\begin{array}{c}
\mathcal{X}_{1}  \tag{5.30}\\
\mathcal{X}_{2} \\
\vdots \\
\mathcal{X}_{n}
\end{array}\right)=\left(\begin{array}{c}
\mathbf{v}_{1}^{\mathrm{T}} \\
\mathbf{v}_{2}^{\mathrm{T}} \\
\vdots \\
\mathbf{v}_{n}^{\mathrm{T}}
\end{array}\right) \mathbf{X}=\boldsymbol{\Gamma} \mathbf{X}
$$

where the matrix $\boldsymbol{\Gamma}$ is made up out of the components of the orthonormal eigenvectors $\left\{\mathbf{v}_{i}\right\}$ :

$$
\boldsymbol{\Gamma}=\left(\begin{array}{c}
\mathbf{v}_{1}^{\mathrm{T}}  \tag{5.31}\\
\mathbf{v}_{2}^{\mathrm{T}} \\
\vdots \\
\mathbf{v}_{n}^{\mathrm{T}}
\end{array}\right)=\left(\begin{array}{cccc}
\left(\mathbf{v}_{1}\right)_{1} & \left(\mathbf{v}_{1}\right)_{2} & \cdots & \left(\mathbf{v}_{1}\right)_{n} \\
\left(\mathbf{v}_{2}\right)_{1} & \left(\mathbf{v}_{2}\right)_{2} & \cdots & \left(\mathbf{v}_{2}\right)_{n} \\
\vdots & \vdots & \ddots & \vdots \\
\left(\mathbf{v}_{n}\right)_{1} & \left(\mathbf{v}_{n}\right)_{2} & \cdots & \left(\mathbf{v}_{n}\right)_{n}
\end{array}\right)
$$

This matrix is not symmetric, but it is orthogonal, meaning that $\boldsymbol{\Gamma}^{\mathrm{T}}=\boldsymbol{\Gamma}^{-1}$. We can see this from

$$
\begin{align*}
\boldsymbol{\Gamma} \boldsymbol{\Gamma}^{\mathrm{T}} & =\left(\begin{array}{c}
\mathbf{v}_{1}^{\mathrm{T}} \\
\mathbf{v}_{2}^{\mathrm{T}} \\
\vdots \\
\mathbf{v}_{n}^{\mathrm{T}}
\end{array}\right)\left(\begin{array}{llll}
\mathbf{v}_{1} & \mathbf{v}_{2}^{\mathrm{T}} & \cdots & \mathbf{v}_{n}^{\mathrm{T}}
\end{array}\right) \\
& =\left(\begin{array}{cccc}
\mathbf{v}_{1}^{\mathrm{T}} \mathbf{v}_{1} & \mathbf{v}_{1}^{\mathrm{T}} \mathbf{v}_{2} & \cdots & \mathbf{v}_{1}^{\mathrm{T}} \mathbf{v}_{n} \\
\mathbf{v}_{2}^{\mathrm{T}} \mathbf{v}_{1} & \mathbf{v}_{2}^{\mathrm{T}} \mathbf{v}_{2} & \cdots & \mathbf{v}_{2}^{\mathrm{T}} \mathbf{v}_{n} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{v}_{n}^{\mathrm{T}} \mathbf{v}_{1} & \mathbf{v}_{n}^{\mathrm{T}} \mathbf{v}_{2} & \cdots & \mathbf{v}_{n}^{\mathrm{T}} \mathbf{v}_{n}
\end{array}\right)=\left(\begin{array}{cccc}
1 & 0 & \cdots & 0 \\
0 & 1 & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & 1
\end{array}\right)=\mathbf{1} \tag{5.32}
\end{align*}
$$

This matrix $\boldsymbol{\Gamma}$ can be thought of a transformation from the original basis to the eigenbasis for $\boldsymbol{\sigma}^{2}$. One effect of this is that it diagonalizes $\boldsymbol{\sigma}^{2}$ :

$$
\boldsymbol{\Gamma} \boldsymbol{\sigma}^{2} \boldsymbol{\Gamma}^{\mathrm{T}}=\left(\begin{array}{cccc}
\lambda_{1} & 0 & \cdots & 0  \tag{5.33}\\
0 & \lambda_{2} & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & \lambda_{n}
\end{array}\right)=\boldsymbol{\Lambda}
$$

Finally, recall that the moment generating function is defined as

$$
\begin{equation*}
M(\mathbf{t})=E\left(e^{\mathbf{t} \mathbf{X}}\right) \tag{5.34}
\end{equation*}
$$

and that if we define $\psi(\mathbf{t})=\ln M(\mathbf{t})$,

$$
\begin{equation*}
\left.\frac{\partial \psi}{\partial t_{i}}\right|_{\mathbf{t}=\mathbf{0}}=\mu_{i} \tag{5.35}
\end{equation*}
$$

and

$$
\begin{equation*}
\left.\frac{\partial^{2} \psi}{\partial t_{i} \partial t_{j}}\right|_{\mathbf{t}=\mathbf{0}}=\operatorname{Cov}\left(X_{1}, X_{2}\right) \tag{5.36}
\end{equation*}
$$

This means that if we do a Maclaurin expansion of $\psi(\mathbf{t})$ we get, in general,

$$
\begin{equation*}
\psi(\mathbf{t})=\mathbf{t}^{\mathrm{T}} \boldsymbol{\mu}+\frac{1}{2} \mathbf{t}^{\mathrm{T}} \boldsymbol{\sigma}^{2} \mathbf{t}+\ldots \tag{5.37}
\end{equation*}
$$

where the terms indicated by ... have three or more powers of t.

### 5.4 General Multivariate Normal Distribution

We define a multivariate normal random vector $X$ as a random vector having the moment generating function

$$
\begin{equation*}
M(\mathbf{t})=\exp \left(\mathbf{t}^{\mathrm{T}} \boldsymbol{\mu}+\frac{1}{2} \mathbf{t}^{\mathrm{T}} \boldsymbol{\sigma}^{2} \mathbf{t}\right) \tag{5.38}
\end{equation*}
$$

We refer to the distribution as $N_{n}\left(\boldsymbol{\mu}, \boldsymbol{\sigma}^{2}\right)$. Note that this is equivalent to starting with the Maclaurin series for $\psi(\mathbf{t})=$ $\ln M(\mathbf{t})$ and cutting it off after the quadratic term.

We start with the mgf rather than the pdf because it applies whether the variance-covariance matrix $\sigma^{2}$ is positive definite or only positive semi-definite, i.e., whether it has one or more zero eigenvalues. To see what happens if one or more of the
eigenvalues is zero, we use the orthonormal eigenvectors $\left\{\mathbf{v}_{i}\right\}$ of $\boldsymbol{\sigma}^{2}$ to combine the random variables in $\mathbf{X}$ into $n$ uncorrelated random variables $\left\{\mathcal{X}_{i}\right\}$, where $\mathcal{X}_{i}=\mathbf{v}_{i}^{\mathrm{T}} \mathbf{X}$, which have means $E\left(\mathcal{X}_{i}\right)=\mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}$ and variances $\operatorname{Var}\left(\mathcal{X}_{i}\right)=\lambda_{i}$. If we combine the $\left\{\mathcal{X}_{i}\right\}$ into a random vector

$$
\begin{equation*}
\mathcal{X}=\Gamma \mathrm{X} \tag{5.39}
\end{equation*}
$$

where

$$
\boldsymbol{\Gamma}=\left(\begin{array}{c}
\mathbf{v}_{1}^{\mathrm{T}}  \tag{5.40}\\
\mathbf{v}_{2}^{\mathrm{T}} \\
\vdots \\
\mathbf{v}_{n}^{\mathrm{T}}
\end{array}\right)
$$

is the orthogonal matrix made up out of eigenvector components, $\mathcal{X}$ has mean $E(\mathcal{X})=\boldsymbol{\Gamma} \boldsymbol{\mu}$ and variance-covariance matrix

$$
\operatorname{Cov}(\boldsymbol{\mathcal { X }})=\boldsymbol{\Lambda}=\boldsymbol{\Gamma} \boldsymbol{\sigma}^{2} \boldsymbol{\Gamma}^{\mathrm{T}}=\left(\begin{array}{cccc}
\lambda_{1} & 0 & \cdots & 0  \tag{5.41}\\
0 & \lambda_{2} & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & \lambda_{n}
\end{array}\right)
$$

The random vector $\boldsymbol{\mathcal { X }}$ also follows a multivariate normal distribution, in this case $N_{n}(\boldsymbol{\Gamma} \boldsymbol{\mu}, \boldsymbol{\Lambda})$. To show this, we'll show the more general result, that if $\mathbf{X}$ is a $N_{n}\left(\boldsymbol{\mu}, \boldsymbol{\sigma}^{2}\right)$ multivariate normal random vector, $\mathbf{A}$ is an $m \times n$ constant matrix and $\mathbf{b}$ is an $m$-element column vector, the random vector $\mathbf{Y}=\mathbf{A X}+\mathbf{b}$ also obeys a multivariate normal distribution. (Note that this works whether $m$ is equal to, less than, or greater than $n!$ ) We prove this using the mgf. The mgf for $\mathbf{Y}$ is

$$
\begin{align*}
M_{Y}(\mathbf{t}) & =E\left[\exp \left(\mathbf{t}^{\mathrm{T}} \mathbf{Y}\right)\right]=E\left[\exp \left(\mathbf{t}^{\mathrm{T}} \mathbf{A} \mathbf{X}+\mathbf{t}^{\mathrm{T}} \mathbf{b}\right)\right] \\
& =e^{\mathrm{t}^{\mathrm{T}} \mathbf{b}} E\left[\exp \left(\left[\mathbf{A}^{\mathrm{T}} \mathbf{t}\right]^{\mathrm{T}} \mathbf{X}\right)\right] \tag{5.42}
\end{align*}
$$

Now here is the key step. $\mathbf{t}$ is an $m$-element column vector. A is an $m \times n$ matrix, so its transpose $\mathbf{A}^{\mathrm{T}}$ is an $n \times m$ matrix,
and the combination $\mathbf{A}^{\mathrm{T}} \mathbf{t}$ is an $n$-element column vector, whose transpose is the $n$-element row vector

$$
\begin{equation*}
\left[\mathbf{A}^{\mathrm{T}} \mathbf{t}\right]^{\mathrm{T}}=\mathbf{t}^{\mathrm{T}} \mathbf{A} \tag{5.43}
\end{equation*}
$$

Therefore, the expectation value in the last line above is just the mgf for the original multivariate normal random vector $\mathbf{X}$ evaluated at the argument $\mathbf{A}^{\mathrm{T}} \mathbf{t}$ :

$$
\begin{align*}
E\left[\exp \left(\left[\mathbf{A}^{\mathrm{T}} \mathbf{t}\right]^{\mathrm{T}} \mathbf{X}\right)\right] & =M_{X}\left(\mathbf{A}^{\mathrm{T}} \mathbf{t}\right) \\
& =\exp \left(\left[\mathbf{A}^{\mathrm{T}} \mathbf{t}\right]^{\mathrm{T}} \boldsymbol{\mu}+\frac{1}{2}\left[\mathbf{A}^{\mathrm{T}} \mathbf{t}\right]^{\mathrm{T}} \boldsymbol{\sigma}^{2}\left[\mathbf{A}^{\mathrm{T}} \mathbf{t}\right]\right) \\
& =\exp \left(\mathbf{t}^{\mathrm{T}} \mathbf{A} \boldsymbol{\mu}+\frac{1}{2} \mathbf{t}^{\mathrm{T}} \mathbf{A} \boldsymbol{\sigma}^{2} \mathbf{A}^{\mathrm{T}} \mathbf{t}\right) \tag{5.44}
\end{align*}
$$

This makes the mgf for $\mathbf{Y}$ equal to $e^{\mathbf{t}^{\mathbf{T}} \mathbf{b}}$ times this, or

$$
\begin{equation*}
M_{Y}(\mathbf{t})=\exp \left(\mathrm{t}^{\mathrm{T}}[\mathbf{A} \boldsymbol{\mu}+\mathbf{b}]+\frac{1}{2} \mathbf{t}^{\mathrm{T}} \mathbf{A} \boldsymbol{\sigma}^{2} \mathbf{A}^{\mathrm{T}} \mathbf{t}\right) \tag{5.45}
\end{equation*}
$$

which is the mgf for a normal random vector with mean $\mathbf{A} \boldsymbol{\mu}+\mathbf{b}$ and variance-covariance matrix $\mathbf{A} \boldsymbol{\sigma}^{2} \mathbf{A}^{\mathrm{T}}$, i.e., one that obeys a $N_{m}\left(\mathbf{A} \boldsymbol{\mu}+\mathbf{b}, \mathbf{A} \boldsymbol{\sigma}^{2} \mathbf{A}^{\mathrm{T}}\right)$ distribution.

So, we return to the random vector $\boldsymbol{\mathcal { X }}=\boldsymbol{\Gamma} \mathbf{X}$, which we now see is a multivariate normal random vector with mean $\boldsymbol{\Gamma} \boldsymbol{\mu}$ and diagonal variance-covariance matrix $\boldsymbol{\Lambda}$. Its mgf is

$$
\begin{align*}
M_{\mathcal{X}}(\mathbf{t}) & =\exp \left(\mathbf{t}^{\mathrm{T}} \boldsymbol{\Gamma} \boldsymbol{\mu}+\frac{1}{2} \mathbf{t}^{\mathrm{T}} \boldsymbol{\Lambda} \mathbf{t}\right)=\exp \left(\sum_{i=1}^{n}\left[t_{i} \mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}+\frac{1}{2} \lambda_{i} t_{i}^{2}\right]\right) \\
& =\prod_{i=1}^{n} \exp \left(t_{i} \mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}+\frac{1}{2} \lambda_{i} t_{i}^{2}\right)=\prod_{i=1}^{n} M_{\mathcal{X}_{i}}\left(t_{i}\right) \tag{5.46}
\end{align*}
$$

which is the mgf of $n$ independent random variables. There are two possibilities: either $\boldsymbol{\sigma}^{2}$ (and thus $\boldsymbol{\Lambda}$ ) is positive definite, which means all of the $\left\{\lambda_{i}\right\}$ are positive, or one or more of the $\left\{\lambda_{i}\right\}$ are zero.

In the first case, we have the special case we considered before, $n$ independent normally-distributed random variables, so the joint pdf is

$$
\begin{equation*}
f_{\mathcal{X}}(\boldsymbol{\xi})=\prod_{i=1}^{n} f_{\mathcal{X}_{i}}\left(\xi_{i}\right)=\frac{1}{\sqrt{\operatorname{det}(2 \pi \boldsymbol{\Lambda})}} \exp \left(-\frac{1}{2}(\boldsymbol{\xi}-\boldsymbol{\Gamma} \boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\Lambda}^{-1}(\boldsymbol{\xi}-\boldsymbol{\Gamma} \boldsymbol{\mu})\right) \tag{5.47}
\end{equation*}
$$

We can then do a multivariate transformation to get the pdf for $\mathbf{X}=\boldsymbol{\Gamma}^{-1} \mathcal{X}=\Gamma^{\mathrm{T}} \mathcal{X}$. The Jacobian of the transformation is $\Gamma^{\mathrm{T}}$, whose determinant is either 1 or -1 , because

$$
\begin{equation*}
1=\operatorname{det} \mathbf{1}=\operatorname{det}\left(\boldsymbol{\Gamma} \boldsymbol{\Gamma}^{\mathrm{T}}\right)=(\operatorname{det} \boldsymbol{\Gamma})^{2} \tag{5.48}
\end{equation*}
$$

(This is true for any orthogonal matrix.) This means $|\operatorname{det} \boldsymbol{\Gamma}|=1$, and the pdf for $\mathbf{X}$ is simply

$$
\begin{equation*}
f_{\mathbf{X}}(\mathbf{x})=f_{\mathcal{X}}(\boldsymbol{\Gamma} \mathbf{x}) \tag{5.49}
\end{equation*}
$$

If we also note that $\operatorname{det} \boldsymbol{\Lambda}=\prod_{i=1}^{n} \lambda_{i}=\operatorname{det} \boldsymbol{\sigma}^{2}$, we see that

$$
\begin{align*}
f_{\mathbf{X}}(\mathbf{x}) & =\frac{1}{\sqrt{\operatorname{det}(2 \pi \boldsymbol{\Lambda})}} \exp \left(-\frac{1}{2} \boldsymbol{\Gamma}(\mathbf{x}-\boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\Lambda}^{-1} \boldsymbol{\Gamma}(\mathbf{x}-\boldsymbol{\mu})\right) \\
& =\frac{1}{\sqrt{\operatorname{det}\left(2 \pi \boldsymbol{\sigma}^{2}\right)}} \exp \left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\Gamma}^{\mathrm{T}} \boldsymbol{\Lambda}^{-1} \boldsymbol{\Gamma}(\mathbf{x}-\boldsymbol{\mu})\right) \tag{5.50}
\end{align*}
$$

But the combination $\boldsymbol{\Gamma}^{\mathrm{T}} \boldsymbol{\Lambda}^{-1} \boldsymbol{\Gamma}$ is just the inverse of $\boldsymbol{\sigma}^{2}$, because

$$
\begin{equation*}
(\boldsymbol{\Lambda})^{-1}=\left(\boldsymbol{\Gamma} \boldsymbol{\sigma}^{2} \boldsymbol{\Gamma}^{\mathrm{T}}\right)^{-1}=\left(\boldsymbol{\Gamma}^{\mathrm{T}}\right)^{-1} \boldsymbol{\sigma}^{-2} \boldsymbol{\Gamma}^{-1} \tag{5.51}
\end{equation*}
$$

so we find that, for arbitrary positive definite symmetric $\boldsymbol{\sigma}^{2}$,

$$
\begin{equation*}
f_{\mathbf{X}}(\mathbf{x})=\frac{1}{\sqrt{\operatorname{det}\left(2 \pi \boldsymbol{\sigma}^{2}\right)}} \exp \left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\sigma}^{-2}(\mathbf{x}-\boldsymbol{\mu})\right) \tag{5.52}
\end{equation*}
$$

which is exactly the generalization we expected from the $n$ -independent-random-variable case.

On the other hand, if $\boldsymbol{\sigma}^{2}$ has one or more negative eigenvalues, so that $\operatorname{det} \boldsymbol{\sigma}^{2}=0$, and $\boldsymbol{\sigma}^{-2}$ is not defined, that pdf won't make sense. In that case, consider the random variable $\mathcal{X}_{i}$ corresponding to the zero eigenvalue $\lambda_{i}=0$. Its mgf is

$$
\begin{equation*}
E\left(e^{t_{i} \mathcal{X}_{i}}\right)=M_{\mathcal{X}_{i}}\left(t_{i}\right)=\exp \left(t_{i} \mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}+\frac{1}{2} \lambda_{i} t_{i}^{2}\right)=\exp \left(t_{i} \mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}\right) \tag{5.53}
\end{equation*}
$$

but the only way that is possible is if $\mathcal{X}_{i}$ is always equal to $\mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}$, i.e., $\mathcal{X}_{i}$ is actually a discrete random variable with pmf

$$
P\left(\mathcal{X}_{i}=\xi_{i}\right)= \begin{cases}1 & \xi_{i}=\mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}  \tag{5.54}\\ 0 & \text { otherwise }\end{cases}
$$

This is the limit of a normal distribution as its variance goes to zero.

Returning to the case where $\boldsymbol{\sigma}^{2}$ is positive definite, so that each $\mathcal{X}_{i}$ is an independent $N\left(\mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}, \lambda_{i}\right)$ random variable, we can construct the corresponding standard normal random variable

$$
\begin{equation*}
\mathcal{Z}_{i}=\left(\lambda_{i}\right)^{-1 / 2}\left(\mathcal{X}_{i}-\mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\mu}\right)=\left(\lambda_{i}\right)^{-1 / 2} \mathbf{v}_{i}^{\mathrm{T}}\left(X_{i}-\boldsymbol{\mu}\right) \tag{5.55}
\end{equation*}
$$

which we could combine into a $N_{n}(\mathbf{0}, \mathbf{1})$ random vector

$$
\mathcal{Z}=\left(\begin{array}{c}
\mathcal{Z}_{1}  \tag{5.56}\\
\mathcal{Z}_{2} \\
\vdots \\
\mathcal{Z}_{n}
\end{array}\right)
$$

However, it's actually more convenient to combine them into a different $N_{n}(\mathbf{0}, \mathbf{1})$ random vector

$$
\begin{equation*}
\boldsymbol{Z}=\sum_{i=1}^{n} \mathbf{v}_{i} \mathcal{Z}_{i}=\boldsymbol{\Gamma}^{\mathrm{T}} \mathcal{Z}=\sum_{i=1}^{n} \mathbf{v}_{i}\left(\lambda_{i}\right)^{-1 / 2} \mathbf{v}_{i}^{\mathrm{T}}\left(X_{i}-\boldsymbol{\mu}\right)=\boldsymbol{\sigma}^{-1}\left(X_{i}-\boldsymbol{\mu}\right) \tag{5.57}
\end{equation*}
$$

with pdf

$$
\begin{equation*}
f_{\mathbf{Z}}(\mathbf{z})=\frac{1}{(2 \pi)^{n / 2}} e^{-\mathbf{z}^{\mathrm{T}} \mathbf{z} / 2} \tag{5.58}
\end{equation*}
$$

## Thursday, October 19, 2017

### 5.5 Sample Mean and Sample Variance (Student's Theorem)

Recall the definition of a random sample $\left\{X_{i} \mid i=1, \ldots, n\right\}$, which is $n$ independent random variables drawn from the same distribution, and let $E\left[X_{i}\right]=\mu$ and $\operatorname{Cov}\left(X_{i}, X_{j}\right)=\delta_{i j} \sigma^{2}$. We showed in class that the sample mean

$$
\begin{equation*}
\bar{X}=\frac{1}{n} \sum_{i=1}^{n} X_{i} \tag{5.59}
\end{equation*}
$$

has $E[\bar{X}]=\mu$ and $\operatorname{Var}(\bar{X})=\sigma^{2} / n$. On the homework, you've shown that the sample variance

$$
\begin{equation*}
S^{2}=\frac{1}{n-1} \sum_{i=1}^{n}\left(X_{i}-\bar{X}\right)^{2} \tag{5.60}
\end{equation*}
$$

has expectation value $E\left[S^{2}\right]=\sigma^{2}$. In addition,

$$
\begin{equation*}
\Sigma^{2}=\frac{1}{n} \sum_{i=1}^{n}\left(X_{i}-\mu\right)^{2} \tag{5.61}
\end{equation*}
$$

has expectation value $E\left[\Sigma^{2}\right]=\sigma^{2}$. These are all true for any underlying distribution.

Now, let's consider the case where the underlying distribution is Gaussian, i.e., $X_{i} \sim N\left(\mu, \sigma^{2}\right)$, which means that $\mathbf{X} \sim N_{n}\left(\mu \mathbf{e}, \sigma^{2} \mathbf{1}_{n \times n}\right)$, where

$$
\mathbf{1}_{n \times n}=\left(\begin{array}{cccc}
1 & 0 & \cdots & 0  \tag{5.62}\\
0 & 1 & \cdots & 0 \\
\vdots & \vdots & \ddots & \vdots \\
0 & 0 & \cdots & 1
\end{array}\right)
$$

is the $n \times n$ identity matrix, and

$$
\mathbf{e}=\left(\begin{array}{c}
1  \tag{5.63}\\
1 \\
\vdots \\
1
\end{array}\right)
$$

Since $\bar{X}=\frac{1}{n} \mathbf{e}^{\mathrm{T}} \mathbf{X}$ is a linear transformation of the Gaussian random vector $\mathbf{X}$ (using the $1 \times n$ matrix $\frac{1}{n} \mathbf{e}^{\mathrm{T}}$ ), we know that $\bar{X}$ is Gaussian-distributed, $\bar{X} \sim N\left(\mu, \sigma^{2} / n\right)$, and that

$$
\begin{equation*}
\frac{\bar{X}-\mu}{\sigma / \sqrt{n}} \sim N(0,1) \tag{5.64}
\end{equation*}
$$

We also know that

$$
\begin{equation*}
\sum_{i=1}^{n}\left(\frac{X_{i}-\mu}{\sigma}\right)^{2}=\frac{n}{\sigma^{2}} \Sigma^{2} \sim \chi^{2}(n) \tag{5.65}
\end{equation*}
$$

so $\Sigma^{2} \sim \operatorname{Gamma}\left(\frac{n}{2}, \frac{2}{n} \sigma^{2}\right)$.
To understand the distribution of $S^{2}$ and its relationship to $\bar{X}$, we use several results which are known collectively as Student's theorem. ("Student" was the pseudonym of William S. Gosset, who found these results while working in the Guinness brewery in Dublin.) These results are:

1. $\bar{X}$ and $S^{2}$ are independent random variables
2. 

$$
\begin{equation*}
\sum_{i=1}^{n}\left(\frac{X_{i}-\bar{X}}{\sigma}\right)^{2}=\frac{n-1}{\sigma^{2}} S^{2} \sim \chi^{2}(n-1) \tag{5.66}
\end{equation*}
$$

which means that $S^{2} \sim \operatorname{Gamma}\left(\frac{n-1}{2}, \frac{2}{n-1} \sigma^{2}\right)$
3. The combination

$$
\begin{equation*}
T=\frac{\bar{X}-\mu}{\sqrt{S^{2} / n}} \tag{5.67}
\end{equation*}
$$

obeys what's known as a Student's $t$-distribution with $n-1$ degrees of freedom. The $t$-distribution with $\nu$ degrees of freedom is the distribution obeyed by the combination

$$
\begin{equation*}
\frac{Z}{\sqrt{W / \nu}} \tag{5.68}
\end{equation*}
$$

where $Z$ and $W$ are independent random variables with $Z \sim N(0,1)$ and $W \sim \chi^{2}(\nu)$. It is a straightforward but unenlightening calculation to show (by writing $f_{Z, W}(z, w)=$ $f_{Z}(z) f_{W}(w)$, changing variables to write $f_{T, W}(t, w)$ and marginalizing to get $\left.f_{T}(t)=\int_{0}^{\infty} f_{T, W}(t, w) d w\right)$ that the pdf of the Student $t$-distribution is

$$
\begin{equation*}
f_{T}(t ; \nu)=\frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\pi \nu} \Gamma\left(\frac{\nu}{2}\right)}\left(1+\frac{t^{2}}{\nu}\right)^{-\frac{\nu+1}{2}} \tag{5.69}
\end{equation*}
$$

If $\nu$ is very large, the $t$-distribution is approximately the same as the standard normal distribution. If $\nu=1$, it's the Cauchy distribution. Note that the moment generating function for the $t$-distribution doesn't exist, since for $k \geq \nu$, the integral

$$
\begin{equation*}
\int_{-\infty}^{\infty} t^{k}\left(1+\frac{t^{2}}{\nu}\right)^{-\frac{\nu+1}{2}} d t \tag{5.70}
\end{equation*}
$$

diverges, as its integrand becomes, up to a constant, $t^{k-\nu-1}$ for large $t$.

The demonstration that Student's theorem is true is a nice application of the multivariate normal distribution. First, we can write

$$
\begin{equation*}
\bar{X}=\frac{1}{n} \sum_{i=1}^{n} X_{i}=\frac{1}{n} \mathbf{e}^{\mathrm{T}} \mathbf{X} \tag{5.71}
\end{equation*}
$$

and

$$
\begin{equation*}
S^{2}=\frac{1}{n-1} \sum_{i=1}^{n}\left(X_{i}-\bar{X}\right)^{2}=\frac{1}{n-1}(\mathbf{X}-\mathbf{e} \bar{X})^{\mathrm{T}}(\mathbf{X}-\mathbf{e} \bar{X}) \tag{5.72}
\end{equation*}
$$

Note that $\bar{X}=\frac{1}{n} \mathbf{e}^{\mathrm{T}} \mathbf{X}$ and

$$
\begin{equation*}
\mathbf{Y}=\mathbf{X}-\mathbf{e} \bar{X}=\left(\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}\right) \mathbf{X} \tag{5.73}
\end{equation*}
$$

are both linear transformations of $\mathbf{X}$. We can combine them into an $n+1$-element random vector

$$
\begin{align*}
\left(\begin{array}{c}
\bar{X} \\
X_{1}-\bar{X} \\
X_{2}-\bar{X} \\
\vdots \\
X_{n}-\bar{X}
\end{array}\right) & =\left(\begin{array}{cccc}
\frac{1}{n} & \frac{1}{n} & \cdots & \frac{1}{n} \\
1-\frac{1}{n} & -\frac{1}{n} & \cdots & -\frac{1}{n} \\
-\frac{1}{n} & 1-\frac{1}{n} & \cdots & -\frac{1}{n} \\
\vdots & \vdots & \ddots & \vdots \\
-\frac{1}{n} & -\frac{1}{n} & \cdots & 1-\frac{1}{n}
\end{array}\right)\left(\begin{array}{c}
X_{1} \\
X_{2} \\
\vdots \\
X_{n}
\end{array}\right)  \tag{5.74}\\
& =\binom{\frac{1}{n} \mathbf{e}^{\mathrm{T}}}{\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}} \mathbf{X}=\mathbf{A X}
\end{align*}
$$

where $\mathbf{A}$ is a $(n+1) \times n$ matrix. Since the original random vector $\mathbf{X}$ has $\boldsymbol{\mu}=\mu \mathbf{e}$ and, the transformed vector $\mathbf{A X}$ has expectation value

$$
\begin{equation*}
E(\mathbf{A X})=\mathbf{A} \boldsymbol{\mu}=\binom{\frac{1}{n} \mathbf{e}^{\mathrm{T}}}{\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}} \mu \mathbf{e}=\binom{\mu}{\mathbf{0}} \tag{5.75}
\end{equation*}
$$

where we've used the fact that

$$
\mathbf{e}^{\mathrm{T}} \mathbf{e}=\left(\begin{array}{c}
1  \tag{5.76}\\
1 \\
\vdots \\
1
\end{array}\right)\left(\begin{array}{llll}
1 & 1 & \cdots & 1
\end{array}\right)=n
$$

and

$$
\begin{equation*}
\left(\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}\right) \mathbf{e}=\mathbf{e}-\frac{n}{n} \mathbf{e}=\mathbf{0} \tag{5.77}
\end{equation*}
$$

Since $\mathbf{X}$ has the variance-covariance matrix $\boldsymbol{\sigma}^{2}=\sigma^{2} \mathbf{1}$, the transformed random vector $\mathbf{A X}$ has variance-covariance matrix ${ }^{[1]}$

$$
\begin{align*}
& \operatorname{Cov}(\mathbf{A X})=\mathbf{A} \boldsymbol{\sigma}^{2} \mathbf{A}^{\mathrm{T}}=\binom{\frac{1}{n} \mathbf{e}^{\mathrm{T}}}{\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}} \sigma^{2} \mathbf{1}\left(\begin{array}{ll}
\frac{1}{n} \mathbf{e} & \left.\mathbf{1}-\frac{1}{n} \mathbf{e} \mathbf{e}^{\mathrm{T}}\right) \\
= & \sigma^{2}\left(\begin{array}{cc}
\frac{\mathbf{e}^{\mathrm{T}} \mathbf{e}}{n^{2}} & \mathbf{e}^{\mathrm{T}}\left(\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}\right) \\
\left(\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}\right) \mathbf{e} & \left(\mathbf{1}-\frac{1}{n} \mathbf{e} \mathbf{e}^{\mathrm{T}}\right)^{2}
\end{array}\right)=\sigma^{2}\left(\begin{array}{cc}
\frac{1}{n} & \mathbf{0}_{1 \times n} \\
\mathbf{0}_{n \times 1} & \left(\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}\right)
\end{array}\right)
\end{array}\right.
\end{align*}
$$

Where we've used the fact (previously noted) that $\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}$ annihilated $\mathbf{e}$ and also that

$$
\begin{equation*}
\left(\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}\right)^{2}=\mathbf{1}-\frac{1}{n} \mathbf{e} \mathbf{e}^{\mathrm{T}}-\frac{1}{n} \mathbf{e} \mathbf{e}^{\mathrm{T}}+\frac{1}{n^{2}} \mathbf{e} n \mathbf{e}^{\mathrm{T}}=\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}} \tag{5.79}
\end{equation*}
$$

We say that a matrix with this property is a projection matrix, in this case onto $n$-dimensional vectors perpendicular to $\mathbf{e}$.

Since the first row and column of $\operatorname{Cov}(\mathbf{A X})$ are all zeros, except for the diagonal element, it means that the random variable $\bar{X}$, which is the first element of $\mathbf{A X}$, and the random vector $\mathbf{Y}$ are independent, which in turn means $\bar{X}$ and $S^{2}=\frac{1}{n-1} \mathbf{Y}^{\mathrm{T}} \mathbf{Y}$ are independent random variables, which is part 2 of Student's

[^8]theorem. We've also seen part 1 , since $\bar{X}$ is a normal random variable whose mean is the first element of $E[\mathbf{A X}]$, i.e., $\mu$ and whose variance is the $(1,1)$ element of $\operatorname{Cov}(\mathbf{A X})$, which is $\sigma^{2} / n$.

To see that

$$
\begin{equation*}
\frac{(n-1) S^{2}}{\sigma^{2}}=\frac{\mathbf{Y}^{\mathrm{T}} \mathbf{Y}}{\sigma^{2}} \tag{5.80}
\end{equation*}
$$

is a chi-square random variable with $n-1$ degrees of freedom, consider the variance-covariance matrix

$$
\begin{equation*}
\operatorname{Cov}(\mathbf{Y})=\left(\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}\right) \sigma^{2} \tag{5.81}
\end{equation*}
$$

This is an $n \times n$ matrix, but it is not invertable, so we can't do the usual trick to construct a $\chi^{2}(n)$ random variable. It's actually not too hard to work out the eigenvalue decomposition of this matrix; since $\left(\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}\right) \mathbf{e}=\mathbf{0}$, we see that $\mathbf{e}$ is an eigenvector with eigenvalue zero. Because the matrix $\mathbf{1}-\frac{1}{n} \mathbf{e e}^{\mathrm{T}}$ is a projector onto the $n$-1-dimensional subspace perpendicular to $\mathbf{e}$, we can choose any $n-1$ orthonormal vectors in that subspace, and they will be eigenvectors with eigenvalue $\sigma^{2}$. For example, take

$$
\begin{gather*}
\mathbf{v}_{1}=\left(\begin{array}{c}
1 / \sqrt{2} \\
-1 \sqrt{2} \\
0 \\
\vdots \\
0 \\
0
\end{array}\right), \quad\left(\begin{array}{c}
1 / \sqrt{6} \\
1 / \sqrt{6} \\
-2 \sqrt{6} \\
\vdots \\
0 \\
0
\end{array}\right), \quad \cdots \\
\mathbf{v}_{2}=\left(\begin{array}{c}
1 / \sqrt{n(n-1)} \\
1 / \sqrt{n(n-1)} \\
1 / \sqrt{n(n-1)} \\
\vdots \\
1 / \sqrt{n(n-1)} \\
-(n-1) / \sqrt{n(n-1)}
\end{array}\right) \tag{5.82}
\end{gather*}
$$

If we let $\mathbf{v}_{n}=\mathbf{e} / \sqrt{n}$, we have our complete set of orthonormal eigenvectors, with $\lambda_{1}=\lambda_{2}=\cdots=\lambda_{n-1}=\sigma^{2}$ and $\lambda_{n}=0$. If we define $\mathcal{Y}_{i}=\mathbf{v}_{i}^{\mathrm{T}} \mathbf{Y}$, then $\mathcal{Y}_{1}, \mathcal{Y}_{2}, \ldots, \mathcal{Y}_{n-1}$ are $n-1$ independent normal random variables each with variance $\sigma^{2}$. (The last combination is trivial, since $\mathbf{v}_{n}^{\mathrm{T}} \mathbf{Y}=\frac{1}{\sqrt{n}} \mathbf{e}^{\mathrm{T}} \mathbf{Y}=\mathbf{0}$.) This means the following combination is a $\chi^{2}(n-1)$ random variable:

$$
\begin{align*}
\sum_{i=1}^{n-1}\left(\frac{\mathcal{Y}_{i}}{\sigma}\right) & =\frac{\mathbf{Y}^{\mathrm{T}} \mathbf{v}_{i} \mathbf{v}_{i}^{\mathrm{T}} \mathbf{Y}}{\sigma^{2}}=\frac{1}{\sigma^{2}} \mathbf{Y}^{\mathrm{T}}\left(\sum_{i=1}^{n-1} \mathbf{v}_{i} \mathbf{v}_{i}^{\mathrm{T}}\right) \mathbf{Y} \\
& =\frac{1}{\sigma^{2}} \mathbf{Y}^{\mathrm{T}}\left(\mathbf{1}-\frac{1}{n} \mathbf{e} \mathbf{e}^{\mathrm{T}}\right) \mathbf{Y}=\frac{1}{\sigma^{2}} \mathbf{Y}^{\mathrm{T}} \mathbf{Y}=\frac{n-1}{\sigma^{2}} S^{2} \tag{5.83}
\end{align*}
$$

This is point 2 of Student's theorem.
Finally, for point 3 we construct a $t$-distributed random variable.

$$
\begin{equation*}
\frac{\bar{X}-\mu}{\sigma / \sqrt{n}} \tag{5.84}
\end{equation*}
$$

is a standard normal random variable, and

$$
\begin{equation*}
\frac{n-1}{\sigma^{2}} S^{2} \tag{5.85}
\end{equation*}
$$

is a $\chi^{2}(n-1)$, we can take

$$
\begin{equation*}
T=\frac{\frac{\bar{X}-\mu}{\sigma / \sqrt{n}}}{\sqrt{\frac{n-1}{\sigma^{2}} S^{2} /(n-1)}}=\frac{\bar{X}-\mu}{\sqrt{S^{2} / n}} \tag{5.86}
\end{equation*}
$$

and it will obey a $t$-distribution with $n-1$ degrees of freedom, which completes the proof of Student's theorem.

## Tuesday, October 24, 2017

### 5.6 Reduced Chi-Squared

We know that if $\left\{X_{i}\right\}$ are iid Gaussian random variables, i.e., $X_{i} \sim N\left(\mu_{i}, \sigma_{i}^{2}\right)$, the combination

$$
\begin{equation*}
\sum_{i=1}^{n}\left(\frac{X_{i}-\mu_{i}}{\sigma_{i}}\right)^{2} \tag{5.87}
\end{equation*}
$$

is a chi-squared random variable with $n$ degrees of freedom, $\chi^{2}(n)$. This is often used as a goodness-of-fit test; if a model predicts that $X_{i} \sim N\left(\mu_{i}, \sigma_{i}^{2}\right)$, we can collect a realization of these data, call it $\left\{x_{i}\right\}$, and calculate

$$
\begin{equation*}
\sum_{i=1}^{n}\left(\frac{x_{i}-\mu_{i}}{\sigma_{i}}\right)^{2} \tag{5.88}
\end{equation*}
$$

if it's much larger than we expect for a chi-square with this number of degrees of freedom, this casts doubt upon the model, since

$$
\begin{equation*}
P\left(\sum_{i=1}^{n}\left(\frac{X_{i}-\mu_{i}}{\sigma_{i}}\right)^{2}>w\right)=\int_{w}^{\infty} f_{\chi^{2}}(u ; n) d u \tag{5.89}
\end{equation*}
$$

For example, if we have five degrees of freedom, we'd only expect to find a chi-square value greater than twenty $0.125 \%$ of the time:

In [1]: from scipy import stats
In [2]: stats.chi2.sf $(20,5)$
Out [2]: 0.0012497305630313482
We've seen that we can generalize this construction to the case of a multivariate Gaussian $\mathbf{X} \sim N\left(\boldsymbol{\mu}, \boldsymbol{\sigma}^{2}\right)$; if the variancecovariance matrix is non-singular, so that $\boldsymbol{\sigma}^{-2}$ exists, we have

$$
\begin{equation*}
(\mathbf{X}-\boldsymbol{\mu})^{\mathrm{T}} \boldsymbol{\sigma}^{-2}(\mathbf{X}-\boldsymbol{\mu}) \sim \chi^{2}(n) \tag{5.90}
\end{equation*}
$$

### 5.6.1 Minimizing $\chi^{2}$ over parameters

An important generalization of the chi-squared construction comes when the model depends on some set of parameters $\boldsymbol{\lambda} \equiv\left\{\lambda_{k} \mid k=1, \ldots, m\right\}$ where $m<n$. Typically, we assume that these parameters influence the means, not the variances, of the normal distribution, so that we have a family of models with $\boldsymbol{\mu}(\boldsymbol{\lambda})$. We can then construct

$$
\begin{equation*}
w(\mathbf{X} ; \boldsymbol{\lambda})=[\mathbf{X}-\boldsymbol{\mu}(\boldsymbol{\lambda})]^{\mathrm{T}} \boldsymbol{\sigma}^{-2}[\mathbf{X}-\boldsymbol{\mu}(\boldsymbol{\lambda})] \tag{5.91}
\end{equation*}
$$

for a particular realization $\mathbf{x}$ of the data, the chi-square $w(\mathbf{x} ; \boldsymbol{\lambda})$ will be a function of $\boldsymbol{\lambda}$, and we can find the $\boldsymbol{\lambda}$ which satisfies that by solving the set of $m$ equations

$$
\begin{equation*}
0=\frac{\partial w}{\partial \lambda_{k}}=-2 \sum_{i=1}^{n} \sum_{j=1}^{n} \frac{\partial \mu_{i}}{\partial \lambda_{k}}\left[\boldsymbol{\sigma}^{-2}\right]_{i j}\left[x_{j}-\mu_{j}(\boldsymbol{\lambda})\right] \tag{5.92}
\end{equation*}
$$

for the $m$ best-fit parameters $\left\{\hat{\lambda}_{k}(\mathbf{x})\right\}$. The minimized chi-square value is then $w(\mathbf{x} ; \widehat{\boldsymbol{\lambda}}(\mathbf{x}))$. In the simplifying case where the means are a linear function of the parameters, $\boldsymbol{\mu}=\mathbf{A} \boldsymbol{\lambda}+\mathrm{b}$, we can show that $w(\mathbf{X} ; \widehat{\boldsymbol{\lambda}}(\mathbf{X})) \sim \chi^{2}(n-m)$, i.e., the constructed quantity still obeys a chi-square distribution, but the number of degrees of freedom is the number of data points minus the number of parameters that we fit. It is generally assumed that this is approximately true even if the dependence on the parameters is not linear.
To show that it works in the case where the predicted $\boldsymbol{\mu}$ is linear in the parameters, $\boldsymbol{\mu}=\mathbf{A} \boldsymbol{\lambda}+\mathbf{b}$, note that in that case $\frac{\partial \mu_{i}}{\partial \lambda_{k}}=A_{i k}$, which is independent of $\boldsymbol{\lambda}$, and we can write the condition for minimizing the chi-square as a matrix equation (dividing by -2 for convenience):

$$
\begin{equation*}
\mathbf{0}=\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2}[\mathbf{x}-\boldsymbol{\mu}(\boldsymbol{\lambda})]=\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2}[\mathbf{x}-\mathbf{A} \boldsymbol{\lambda}-\mathbf{b}] \tag{5.93}
\end{equation*}
$$

$$
\begin{equation*}
\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A} \widehat{\boldsymbol{\lambda}}(\mathbf{x})=\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2}(\mathbf{x}-\mathbf{b}) \tag{5.94}
\end{equation*}
$$

Since $\mathbf{A}$ is an $n \times m$ matrix (it maps $m$ parameters into $n$ data values), the combination $\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}$ is an $m \times m$ matrix. We can assume it's invertible, since if it's not we don't actually need all $m$ parameters to describe the model. Thus the best-fit parameters are

$$
\begin{equation*}
\widehat{\boldsymbol{\lambda}}(\mathbf{x})=\left(\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\right)^{-1} \mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2}(\mathbf{x}-\mathbf{b}) \tag{5.95}
\end{equation*}
$$

and the corresponding expectation values for the data are

$$
\begin{equation*}
\boldsymbol{\mu}(\widehat{\boldsymbol{\lambda}}(\mathbf{x}))=\mathbf{A}\left(\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\right)^{-1} \mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2}(\mathbf{x}-\mathbf{b}) \tag{5.96}
\end{equation*}
$$

Note that we may be tempted to simplify this with something involving the matrix inverse of $\mathbf{A}$, but that doesn't exist because A is not a square matrix. What we can do, though, is define the matrix

$$
\begin{equation*}
\mathbf{P}=\boldsymbol{\sigma}^{-1} \mathbf{A}\left[\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\right]^{-1} \mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-1} \tag{5.97}
\end{equation*}
$$

so that

$$
\begin{equation*}
\boldsymbol{\mu}(\hat{\boldsymbol{\lambda}}(\mathrm{x}))=\boldsymbol{\sigma} \mathrm{P}^{-1} \mathrm{x} \tag{5.98}
\end{equation*}
$$

In fact, since

$$
\begin{equation*}
\mathbf{P} \boldsymbol{\sigma}^{-1} \mathbf{A}=\boldsymbol{\sigma}^{-1} \mathbf{A} \tag{5.99}
\end{equation*}
$$

we see that for any $\boldsymbol{\lambda}$,

$$
\begin{equation*}
\mathrm{P} \boldsymbol{\sigma}^{-1} \boldsymbol{\mu}(\boldsymbol{\lambda})=\mathrm{P} \boldsymbol{\sigma}^{-1} \mathrm{~A} \boldsymbol{\lambda}=\boldsymbol{\sigma}^{-1} \mathrm{~A} \boldsymbol{\lambda}=\boldsymbol{\sigma}^{-1} \boldsymbol{\mu}(\boldsymbol{\lambda}) \tag{5.100}
\end{equation*}
$$

The matrix $\mathbf{P}$ is not only symmetric, but it's a projection matrix, since

$$
\begin{align*}
\mathbf{P} \mathbf{P} & =\boldsymbol{\sigma}^{-1} \mathbf{A}\left[\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\right]^{-1} \mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\left[\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\right]^{-1} \mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-1} \\
& =\boldsymbol{\sigma}^{-1} \mathbf{A}\left[\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\right]^{-1} \mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-1}=\mathbf{P} \tag{5.101}
\end{align*}
$$

This $n \times n$ matrix is a projector onto an $m$-dimensional subspace, since

$$
\begin{align*}
\operatorname{tr} \mathbf{P} & =\operatorname{tr}\left(\boldsymbol{\sigma}^{-1} \mathbf{A}\left[\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\right]^{-1} \mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-1}\right)  \tag{5.102}\\
& =\operatorname{tr}\left(\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\left[\mathbf{A}^{\mathrm{T}} \boldsymbol{\sigma}^{-2} \mathbf{A}\right]^{-1}\right)=\operatorname{tr} \mathbf{1}_{m \times m}=m
\end{align*}
$$

That means that it has $m$ eigenvectors with unit eigenvalue and $n-m$ eigenvectors with zero eigenvalue. Let $\left\{\mathbf{u}_{i} \mid i=1, \ldots, m\right\}$ be a set of $m$ orthonormal eigenvectors with unit eigenvalue and $\left\{\mathbf{u}_{i} \mid i=m+1, \ldots, n\right\}$ be a set of $n-m$ orthonormal eigenvectors with zero eigenvalue, so that

$$
\begin{equation*}
\mathbf{P}=\sum_{i=1}^{m} \mathbf{u}_{i} \mathbf{u}_{i}^{\mathrm{T}} \tag{5.103}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbf{1}_{n \times n}-\mathbf{P}=\sum_{i=m+1}^{n} \mathbf{u}_{i} \mathbf{u}_{i}^{\mathrm{T}} \tag{5.104}
\end{equation*}
$$

are projection operators onto the two orthogonal subspaces.
Now, let $\boldsymbol{\lambda}_{\text {true }}$ be the column vector of true, unknown parameters. Then

$$
\begin{equation*}
\mathbf{Z}=\boldsymbol{\sigma}^{-1}\left[\mathbf{X}-\boldsymbol{\mu}\left(\boldsymbol{\lambda}_{\text {true }}\right)\right] \tag{5.105}
\end{equation*}
$$

is a vector of $n$ independent standard normal random variables. Since $\left\{\mathbf{u}_{i} \mid i=1, \ldots, n\right\}$ is an orthonormal basis, we can also construct independent standard normal random variables $\left\{\mathcal{Z}_{i} \mid i=1, \ldots, n\right\}$ where

$$
\begin{equation*}
\mathcal{Z}_{i}=\mathbf{u}_{i}^{\mathrm{T}} \mathbf{Z}=\mathbf{u}_{i}^{\mathrm{T}} \boldsymbol{\sigma}^{-1}\left(\mathbf{X}-\boldsymbol{\mu}\left(\boldsymbol{\lambda}_{\text {true }}\right)\right) \tag{5.106}
\end{equation*}
$$

If we take the sum of the squares of the last $n-m$ of these random variables, that will obey a chi-squared distribution with
$n-m$ degrees of freedom:

$$
\begin{align*}
W_{n-m} & =\sum_{i=m+1}^{n}\left[\mathcal{Z}_{i}\right]^{2}=\left[\mathbf{X}-\boldsymbol{\mu}\left(\boldsymbol{\lambda}_{\text {true }}\right)\right]^{\mathrm{T}} \boldsymbol{\sigma}^{-1}\left(\sum_{i=m+1}^{n} \mathbf{u}_{i} \mathbf{u}_{i}^{\mathrm{T}}\right) \boldsymbol{\sigma}^{-1}\left[\mathbf{X}-\boldsymbol{\mu}\left(\boldsymbol{f}_{\text {true }}\right)\right] \\
& =\left[\mathbf{X}-\boldsymbol{\mu}\left(\boldsymbol{\lambda}_{\text {true }}\right)\right]^{\mathrm{T}} \boldsymbol{\sigma}^{-1}\left(\mathbf{1}_{n \times n}-\mathbf{P}\right) \boldsymbol{\sigma}^{-1}\left[\mathbf{X}-\boldsymbol{\mu}\left(\boldsymbol{\lambda}_{\text {true }}\right)\right] \tag{5.107}
\end{align*}
$$

But by (5.100) and 5.98,

$$
\begin{align*}
\left(\mathbf{1}_{n \times n}-\mathbf{P}\right) \boldsymbol{\sigma}^{-1}\left[\mathbf{X}-\boldsymbol{\mu}\left(\boldsymbol{\lambda}_{\text {true }}\right)\right] & =\left(\mathbf{1}_{n \times n}-\mathbf{P}\right) \boldsymbol{\sigma}^{-1} \mathbf{X}=\boldsymbol{\sigma}^{-1} \mathbf{X}-\mathbf{P} \boldsymbol{\sigma}^{-1} \mathbf{X} \\
& =\boldsymbol{\sigma}^{-1}[\mathbf{X}-\mu(\hat{\boldsymbol{\lambda}}(\mathbf{X}))] \tag{5.108}
\end{align*}
$$

SO

$$
\begin{equation*}
W_{n-m}=[\mathbf{X}-\mu(\hat{\boldsymbol{\lambda}}(\mathbf{X}))]^{\mathrm{T}} \sigma^{-2}[\mathbf{X}-\mu(\hat{\boldsymbol{\lambda}}(\mathbf{X}))]=W_{\mathrm{red}} \tag{5.109}
\end{equation*}
$$

which means this quantity we've constructed, which is $\chi^{2}$ distributed with $n-m$ degrees of freedom (in the case where the (??) modelled expectation values are linear in the parameters, i.e., $\boldsymbol{\mu}=\mathbf{A} \boldsymbol{\lambda}+\mathbf{b})$, is indeed the residual chi-squared $w(\mathbf{X} ; \widehat{\boldsymbol{\lambda}}(\mathbf{X}))$.


[^0]:    *Copyright 2017, John T. Whelan, and all that

[^1]:    ${ }^{1}$ Note that the denominator $P(D \mid I)$ has cancelled out, which is a good thing since to calculate it we'd need to consider every possible hypothesis.
    ${ }^{2}$ In the strict frequentist interpretation, we're also not allowed to use $H$ or $I$ in a conditional probability so it's usually written something like $P(D ; H)$ or $P_{H}(D)$, but let's stick with the Bayesian notation since we know what it means.

[^2]:    ${ }^{3}$ The characteristic function for the Cauchy distribution is $\Phi(\xi)=e^{-|\xi|}$, which is not analytic at $\xi=0$, which is why we cannot use analytic continuation to define the mgf.

[^3]:    ${ }^{5}$ For more review and practical discussion of the binomial coëfficient, see section 1.1 of http://ccrg.rit.edu/~whelan/courses/2013_3fa_STAT_ 405/notes03.pdf

[^4]:    ${ }^{6}$ It's only approximate because the total number of events is fixed to be $M \mu$. To make the construction exact, we would need to generate one Poisson random value using a Poisson distribution with mean $M \mu$, and then use that number

[^5]:    ${ }^{7}$ Note that some sources define the Gamma distribution in terms of a rate parameter which they call $\beta$, but is one over our scale parameter $\beta$. (Sources using this other convention include Wikipedia and the notes for previous versions of this class!) When in doubt, you can fall back on dimensional analysis. If $\beta$ is a scale parameter, it will have the same units as $x$, and $x / \beta$ will appear in the exponential. If it were a rate parameter, it would have units of $1 / x$. Note that Gregory uses the scale parameter definition, but calls the parameter $\theta$.

[^6]:    ${ }^{8}$ We haven't demonstrated this one, but it is easily done either by applying a change of variables to the pdf, or by using the mgf to show that $M_{a X}(t)=E\left[e^{t(a X)}\right]=E\left[e^{(t a) X}\right]=M_{X}(a t)=(1-\beta a t)^{-\alpha}$ which is the $m g f$ of a $\operatorname{Gamma}(\alpha, a \beta)$

[^7]:    ${ }^{9} \mathrm{An}$ alternative, slick shortcut is to write

    $$
    f_{T}(t)=\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \delta\left(t-\left[x_{1}-x_{2}\right]\right) f\left(x_{1}, x_{2}\right) d x_{1} d x_{2}
    $$

[^8]:    ${ }^{11}$ Remember that $\mathbf{A}$ is $(n+1) \times n, \boldsymbol{\sigma}^{2}$ is $n \times n$, and $\mathbf{A}^{\mathrm{T}}$ is $n \times(n-1)$, so $\mathbf{A} \boldsymbol{\sigma}^{2} \mathbf{A}^{\mathrm{T}}$ is $(n+1) \times(n+1)$.

