Data analysis recipes:
Probability calculus for inference[*]

David W. Hogg

Center for Cosmology and Particle Physics, Department of Physics, New York University
and *Max-Planck-Institut für Astronomie, Heidelberg*

Abstract

In this pedagogical text aimed at those wanting to start thinking about or brush up on probabilistic inference, I review the rules by which probability distribution functions can (and cannot) be combined. I connect these rules to the operations performed in probabilistic data analysis. Dimensional analysis is emphasized as a valuable tool for helping to construct non-wrong probabilistic statements. The applications of probability calculus in constructing likelihoods, marginalized likelihoods, posterior probabilities, and posterior predictions are all discussed.

When constructing the plan or basis for a probabilistic inference—a data analysis making use of likelihoods and also prior probabilities and posterior probabilities and marginalization of nuisance parameters—the options are *incredibly strongly constrained* by the simple rules of probability calculus. That is, there are only certain ways that probability distribution functions ("pdfs" in what follows) *can* be combined to make new pdfs, compute expectation values, or make other kinds of non-wrong statements. For this reason, it behooves the data analyst to have good familiarity and facility with these rules.

Formally, probability calculus is extremely simple. However, it is not always part of a physicist's (or biologist's or chemist's or economist's) education. That motivates this document.

For space and specificity—and because it is useful for most problems I encounter—I will focus on continuous variables (think "model parameters" for the purposes of inference) rather than binary, integer, or discrete parameters, although I might say one or two things here or there. I also won't always be specific about the domain of variables (for example whether a variable a is defined on $0 < a < 1$ or $0 < a < \infty$ or $-\infty < a < \infty$); the limits of integrals—all of which are definite—will usually be implicit.[3] Along the way,

---

[*]The notes begin on page 15, including the license[1] and the acknowledgements[2]

a key idea I want to convey is—and this shows that I come from physics—
that it is very helpful to think about the units or dimensions of probability
quantities.

I learned the material here by using it; probability calculus is so simple
and constraining, you can truly learn it as you go. However if you want more
information or discussion about any of the subjects below, Sivia & Skilling
(2006) provide a book-length, practical introduction and the first few chap-
ters of Jaynes (2003) make a great and exceedingly idiosyncratic read.

1   Generalities

A probability distribution function has units or dimensions. Don't ignore
them. For example, if you have a continuous parameter a, and a pdf $p(a)$ for
a, it must obey the normalization condition

$$1 \;=\; \int p(a)\,\mathrm{d}a \quad, \tag{1}$$

where the limits of the integral should be thought of as going over the entire
domain of a. This (along with, perhaps, $p(a) \geq 0$ everywhere) is almost
the *definition* of a pdf, from my (pragmatic, informal) point of view. This
normalization condition shows that $p(a)$ has units of a^{-1}. Nothing else would
integrate properly to a dimensionless result. Even if a is a multi-dimensional
vector or list or tensor or field or even point in function space, the pdf must
have units of a^{-1}.

In the multi-dimensional case, the units of a^{-1} are found by taking the
product of all the units of all the dimensions. So, for example, if a is a six-
dimensional phase-space position in three-dimensional space (three cartesian
position components measured in m and three cartesian momentum compo-
nents measured in $\mathrm{kg\,m\,s^{-1}}$), the units of $p(a)$ would be $\mathrm{kg^{-3}\,m^{-6}\,s^3}$.

Most problems we will encounter will have multiple parameters; even if
we *condition* $p(a)$ on some particular value of another parameter b, that is,
ask for the pdf for a *given* that b has a particular, known value to make
$p(a\,|\,b)$ (read "the pdf for a given b"), it must obey the same normalization

$$1 \;=\; \int p(a\,|\,b)\,\mathrm{d}a \quad, \tag{2}$$

but you can *absolutely never do* the integral

$$\textbf{wrong:} \quad \int p(a \,|\, b) \,\mathrm{d}b \tag{3}$$

because that integral would have units of $a^{-1} b$, which is (for our purposes) absurd.[4]

If you have a probability distribution for two things ("the pdf for a and b"), you can always factorize it into two distributions, one for a, and one for b given a or the other way around:

$$p(a, b) \;=\; p(a) \, p(b \,|\, a) \tag{4}$$
$$p(a, b) \;=\; p(a \,|\, b) \, p(b) \quad, \tag{5}$$

where the units of both sides of both equations are $a^{-1} b^{-1}$. These two factorizations taken together lead to what is sometimes called "Bayes's theorem", or

$$p(a \,|\, b) \;=\; \frac{p(b \,|\, a) \, p(a)}{p(b)} \quad, \tag{6}$$

where the units are just a^{-1} (the b^{-1} units cancel out top and bottom), and the "divide by $p(b)$" aspect of that gives many philosophers and mathematicians the chills (though certainly not me).[5] Conditional probabilities factor just the same as unconditional ones (and many will tell you that there is no such thing as an unconditional probability[6]); they factor like this:

$$p(a, b \,|\, c) \;=\; p(a \,|\, c) \, p(b \,|\, a, c) \tag{7}$$
$$p(a, b \,|\, c) \;=\; p(a \,|\, b, c) \, p(b \,|\, c) \tag{8}$$
$$p(a \,|\, b, c) \;=\; \frac{p(b \,|\, a, c) \, p(a \,|\, c)}{p(b \,|\, c)} \quad, \tag{9}$$

where the condition c must be carried through all the terms; the whole right-hand side must be conditioned on c if the left-hand side is. Again, there was Bayes's theorem, and you can see its role in conversions of one kind of conditional probability into another. For technical reasons,[7] I usually write Bayes's theorem like this:

$$p(a \,|\, b, c) \;=\; \frac{1}{Z} \, p(b \,|\, a, c) \, p(a \,|\, c) \tag{10}$$

$$Z \;\equiv\; \int p(b \,|\, a, c) \, p(a \,|\, c) \,\mathrm{d}a \quad. \tag{11}$$

Here are things you *can't* do:

$$\textbf{wrong:} \quad p(a \,|\, b, c)\, p(b \,|\, a, c) \tag{12}$$

$$\textbf{wrong:} \quad p(a \,|\, b, c)\, p(a \,|\, c) \quad ; \tag{13}$$

the first over-conditions (it is not a factorization of anything possible) and the second has units of a^{-2}, which is absurd (for our purposes). Know these and *don't do them.*

One important and confusing point about all this: The terminology used throughout this document *enormously overloads* the symbol $p(\cdot)$. That is, we are using, in each line of this discussion, the function $p(\cdot)$ to mean something different; it's meaning is set by the letters used in its arguments. That is a nomenclatural abomination.[8] I apologize, and encourage my readers to do things that aren't so ambiguous (like maybe add informative subscripts), but it is so standard in our business that I won't change (for now).

The theory of continuous pdfs is measure theory; measure theory (for me, anyway[9]) is the theory of things in which you can do integrals. You can *integrate out* or *marginalize away* variables you want to get rid of (or, in what follows, *not* infer) by integrals that look like

$$p(a \,|\, c) \;=\; \int p(a, b \,|\, c)\, \mathrm{d}b \tag{14}$$

$$p(a \,|\, c) \;=\; \int p(a \,|\, b, c)\, p(b \,|\, c)\, \mathrm{d}b \quad , \tag{15}$$

where the second is a factorized version of the first. Once again the integrals go over the entire domain of b in each case, and again if the left-hand side is conditioned on c, then everything on the right-hand side must be also. This equation is a natural consequence of the things written above and dimensional analysis. Recall that because b is some kind of arbitrary, possibly very high-dimensional mathematical object, these integrals can be extremely daunting in practice (see below). Sometimes equations like (15) can be written

$$p(a \,|\, c) \;=\; \int p(a \,|\, b)\, p(b \,|\, c)\, \mathrm{d}b \quad , \tag{16}$$

where the dependence of $p(a \,|\, b)$ on c has been dropped. This is only permitted if it *happens to be the case* that $p(a \,|\, b, c)$ doesn't, in practice, depend on c. The dependence on c is really there (in some sense), it just might be trivial or null.

In rare cases, you can get factorizations that look like this:

$$p(a, b \,|\, c) \;\; = \;\; p(a \,|\, c)\, p(b \,|\, c) \;\; ; \tag{17}$$

this factorization doesn't have the pdf for a depend on b or vice versa. When this happens—and it is rare—it says that a and b are "independent" (at least conditional on c).[10] In many cases of data analysis, in which we want probabilistic models of data sets, we often have some number N of data a_n (indexed by n) each of which is independent in this sense. If for each datum a_n you can write down a probability distribution $p(a_n \,|\, c)$, then the probability of the full data set—when the data are independent—is simply the product of the individual data-point probabilities:

$$p(\{a_n\}_{n=1}^{N} \,|\, c) \;\; = \;\; \prod_{n=1}^{N} p(a_n \,|\, c) \;\; . \tag{18}$$

This is the definition of "independent data". If all the functions $p(a_n \,|\, c)$ are identical—that is, if the functions don't depend on n—we say the data are "iid" or "independent and identically distributed". However, this will not be true in general in data analysis. Real data are heteroscedastic at the very least.[11]

I am writing here mainly about continuous variables, but one thing that comes up frequently in data analysis is the idea of a "mixture model" in which data are produced by two (or more) qualitatively different processes (some data are good, and some are bad, for example) that have different relative probabilities. When a variable (b, say) is discrete, the marginalization integral corresponding to equation (15) becomes a sum

$$p(a \,|\, c) \;\; = \;\; \sum_{b} p(a \,|\, b, c)\, p(b \,|\, c) \;\; , \tag{19}$$

and the normalization of $p(b \,|\, c)$ becomes

$$1 \;\; = \;\; \sum_{b} p(b \,|\, c) \;\; ; \tag{20}$$

in both sums, the sum is implicitly over the (countable number of) possible states of b.[12]

If you have a conditional pdf for a, for example $p(a \,|\, c)$, and you want to know the expectation value $E(a \,|\, c)$ of a under this pdf (which would be, for

example, something like the mean value of a you would get if you drew many draws from the conditional pdf), you just integrate

$$E(a \mid c) \;=\; \int a \, p(a \mid c) \, \mathrm{d}a \quad . \tag{21}$$

This generalizes to any function $f(a)$ of a:

$$E(f \mid c) \;=\; \int f(a) \, p(a \mid c) \, \mathrm{d}a \quad . \tag{22}$$

You can see the marginalization integral (15) that converts $p(a \mid b, c)$ into $p(a \mid c)$ as providing the *expectation value* of $p(a \mid b, c)$ under the conditional pdf $p(b \mid c)$. That's deep and relevant for what follows.

Exercise 1: You have conditional pdfs $p(a \mid d)$, $p(b \mid a, d)$, and $p(c \mid a, b, d)$. Write expressions for $p(a, b \mid d)$, $p(b \mid d)$, and $p(a \mid c, d)$.

Exercise 2: You have conditional pdfs $p(a \mid b, c)$ and $p(a \mid c)$ expressed or computable for any values of a, b, and c. You are not permitted to multiply these together, of course. But can you use them to construct the conditional pdf $p(b \mid a, c)$ or $p(b \mid c)$? Did you have to make any assumptions?

Exercise 3: You have conditional pdfs $p(a \mid c)$ and $p(b \mid c)$ expressed or computable for any values of a, b, and c. Can you use them to construct the conditional pdf $p(a \mid b, c)$?

Exercise 4: You have a function $g(b)$ that is a function only of b. You have conditional pdfs $p(a \mid c)$ and $p(b \mid a, c)$. What is the expectation value $E(g \mid c)$ for g conditional on c but *not* conditional on a?

Exercise 5: Take the integral on the right-hand side of equation (15) and replace the "$\mathrm{d}b$" with a "$\mathrm{d}a$". Is it permissible to do this integral? Why or why not? If it *is* permissible, what do you get?

2   Likelihoods

Imagine you have N data points or measurements D_n of some kind, possibly times or temperatures or brightnesses. I will say that you have a "generative model" of data point n if you can write down or calculate a pdf $p(D_n \mid \theta, I)$ for the measurement D_n, conditional on a vector or list θ of parameters and a (possibly large) number of other things I ("prior information") on which the D_n pdf depends, such as assumptions, or approximations, or knowledge about the noise process, or so on. If all the data points are independently drawn (that would be one of the assumptions in I), then the pdf for the full data set $\{D_n\}_{n=1}^N$ is just the product

$$p(\{D_n\}_{n=1}^N \mid \theta, I) \;=\; \prod_{n=1}^{N} p(D_n \mid \theta, I) \quad . \tag{23}$$

(This requires the data to be independent, but not necessarily iid.) When the pdf of the data is thought of as being a *function* of the parameters at *fixed* data, the pdf of the data given the parameters is called the "likelihood" (for historical reasons I don't care about). In general, in contexts in which the data are thought of as being fixed and the parameters are thought of as variable, any kind of conditional pdf for the data—conditional on parameters— is called a likelihood "for the parameters" even though it is a pdf "for the data".[13]

Now imagine that the parameters divide into two groups. One group θ are parameters of great interest, and another group α are of no interest. The α parameters are nuisance parameters. In this situation, the likelihood can be written

$$p(\{D_n\}_{n=1}^N \mid \theta, \alpha, I) \;=\; \prod_{n=1}^{N} p(D_n \mid \theta, \alpha, I) \quad . \tag{24}$$

If you want to make likelihood statements about the important parameters θ without committing to anything regarding the nuisance parameters α, you can marginalize rather than infer them. You might be tempted to do

$$\textbf{wrong:} \quad \int p(\{D_n\}_{n=1}^N \mid \theta, \alpha, I) \, \mathrm{d}\alpha \quad , \tag{25}$$

but that's not allowed for dimensional arguments given in the previous Section.[14] In order to integrate over the nuisances α, something with units of

α^{-1} needs to be multiplied in—a pdf for the α of course:

$$p(\{D_n\}_{n=1}^N \,|\, \theta, I) \;=\; \int p(\{D_n\}_{n=1}^N \,|\, \theta, \alpha, I) \, p(\alpha \,|\, \theta, I) \, \mathrm{d}\alpha \quad , \qquad (26)$$

where $p(\alpha \,|\, \theta, I)$ is called the "prior pdf" for the α and it *can* depend (but doesn't *need to* depend) on the other parameters θ and the more general prior information I. This marginalization is incredibly useful but note that it comes at a substantial cost: It required specifying a prior pdf over parameters that, by assertion, you don't care about![15] Equation (26) could be called a "partially marginalized likelihood" because it is a likelihood (a pdf for the data) but it is conditional on fewer parameters than the original, rawest, likelihood.

Sometimes I have heard concern that when you perform the marginalization (26), you are allowing the nuisance parameters to "have any values they like, whatsoever" as if you are somehow not constraining them. It is true that you are not *inferring* the nuisance parameters, but you certainly are using the data to limit their range, in the sense that the integral in (26) only gets significant weight where the likelihood is large. That is, if the data strongly constrain the α to a narrow range of good values, then only those good values are making any (significant) contribution to the marginalization integral. That's important!

Because the data were independent by assumption (the full-data likelihood is a product of individual-datum likelihoods), you might be tempted to do things like

$$\textbf{wrong:} \quad \prod_{n=1}^N \left[\int p(D_n \,|\, \theta, \alpha, I) \, p(\alpha \,|\, \theta, I) \, \mathrm{d}\alpha \right] \quad . \qquad (27)$$

This is wrong because if you do the integral *inside* the product, you end up doing the integral N times over. Or another way to put it, although you don't want to infer the α parameters, you want the support in the marginalization integral to be consistently set by *all* the data taken together, not inconsistently set by each individual datum separately.

One thing that is often done with likelihoods (and one thing that many audience members think, instinctively, when you mention the word "likelihood") is "maximum likelihood". If all you want to write down is your likelihood, and you want the "best" parameters θ given your data, you can find the parameters that maximize the full-data likelihood. The only really

probabilistically responsible use of the maximum-likelihood parameter value is—when coupled with a likelihood width estimate—to make an approximate description of a likelihood function.[16] Of course a maximum-likelihood value could in principle be useful on its own when the data are *so decisive* that there is (for the investigator's purposes) no significant remaining uncertainty in the parameters. I have never seen that happen. The key idea is that it is the *likelihood function* that is useful for inference, not some parameter value suggested by that function.

To make all of the above concrete, we can consider a simple example, where measurements D_n are made at a set of "horizontal" postions x_n. Each measurement is of a "vertical" position y_n but the measurement is noisy, so the generative model looks like:

$$y_n = a\,x_n + b \tag{28}$$

$$D_n = y_n + e_n \tag{29}$$

$$p(e_n) = N(e_n\,|\,0,\sigma_n^2) \tag{30}$$

$$p(D_n\,|\,\theta,I) = N(D_n\,|\,a\,x_n + b,\sigma_n^2) \tag{31}$$

$$p(\{D_n\}_{n=1}^N\,|\,\theta,I) = \prod_{n=1}^N p(D_n\,|\,\theta,I) \tag{32}$$

$$\theta \equiv [a,b] \tag{33}$$

$$I \equiv [\{x_n,\sigma_n^2\}_{n=1}^N,\text{and so on...}]\quad, \tag{34}$$

where the y_n are the "true values" for the heights, which lie on a straight line of slope a and intercept b, the e_n are noise contributions, which are drawn independently from Gaussians $N(\cdot\,|\,\cdot)$ with zero means and variances σ_n^2, the likelihood is just a Gaussian for each data point, there are two parameters, and the x_n and σ_n^2 values are considered prior information.[17]

Exercise 6: Show that the likelihood for the model given in equations (28) through (34) can be written in the form $Q\exp(-\chi^2/2)$, where χ^2 is the standard statistic for weighted least-squares problems. On what does Q depend, and what are its dimensions?

Exercise 7: The likelihood in equation (32) is a product of Gaussians in D_n. At fixed data and b, what shape will it have in the a direction? That is, what functional form will it have when thought of as being a function

of a? You will have to use the properties of Gaussians (and products of Gaussians).

3 Posterior probabilities

A large fraction of the inference that is done in the quantitative sciences can be left in the form of likelihoods and marginalized likelihoods, and probably should be.[18] However, there are many scientific questions the answers to which require going beyond the likelihood—which is a pdf for the data, conditional on the parameters—to a pdf for the parameters, conditional on the data.

To illustrate, imagine that you want to make a probabilistic prediction, given your data analysis. For one example, you might want to predict what θ values you will find in a subsequent experiment. Or, for another, say some quantity t of great interest (like, say, the age of the Universe) is a function $t(\theta)$ of the parameters and you want to predict the outcome you expect in some future independent measurement of that same parameter (some other experiment that measures—differently—the age of the Universe). The distributions or expectation values for these predictions, conditioned on your data, will require a pdf for the parameters or functions of the parameters; that is, if you want the expectation $E(t\,|\,D)$, where for notational convenience I have defined

$$D \;\equiv\; \{D_n\}_{n=1}^{N} \quad , \tag{35}$$

you need to do the integral

$$E(t\,|\,D) \;=\; \int t(\theta)\,p(\theta\,|\,D,I)\,\mathrm{d}\theta \quad , \tag{36}$$

this in turn requires the pdf $p(\theta\,|\,D,I)$ for the parameters θ given the data. This is called the "posterior pdf" because it is the pdf you get *after* digesting the data.

The posterior pdf is obtained by Bayes rule (10)

$$p(\theta\,|\,D,I) \;=\; \frac{1}{Z}\,p(D\,|\,\theta,I)\,p(\theta\,|\,I) \tag{37}$$

$$Z \;\equiv\; \int p(D\,|\,\theta,I)\,p(\theta\,|\,I)\,\mathrm{d}\theta \quad , \tag{38}$$

where we had to introduce the prior pdf $p(\theta \mid I)$ for the parameters. The prior can be thought of as the pdf for the parameters *before* you took the data D. The prior pdf brings in new assumptions but also new capabilities, because posterior expectation values as in (36) and other kinds of probabilistic predictions become possible with its use.

Some notes about all this: *(a)* Computation of expectation values is not the only—or even the primary—use of posterior pdfs; I was just using the expectation value as an example of *why* you might want the posterior. You can use posterior pdfs to make predictions that are themselves pdfs; this is in general the only way to propagate the full posterior uncertainty remaining after your experiment. *(b)* The normalization Z in equation (37) is a marginalization of a likelihood; in fact it is a fully marginalized likelihood, and could be written $p(D \mid I)$. It has many uses in model evaluation and model averaging, to be discussed in subsequent documents; it is sometimes called "the evidence" though that is a very unspecific term I don't like so much. *(c)* You can think of the posterior expression (37) as being a "belief updating" operation, in which you start with the prior pdf, multiply in the likelihood (which probably makes it narrower, at least if your data are useful) and re-normalize to make a posterior pdf. There is a "subjective" attitude to take towards all this that makes the prior and the posterior pdfs specific to the individual inferrer, while the likelihood is (at least slightly) more objective.[19] *(d)* Many committed Bayesians take the view that you *always* want a posterior pdf—that is, you are never satisfied with a likelihood—and that you therefore *must* always have a prior, even if you don't think you do. That view is false, but it contains an element of truth: If you eschew prior pdfs, then you are relegated to only ever asking and answering questions about the probability of the data. You can answer questions like "what regions of parameter space are consistent with the data?" but within the set of consistent models, you can't answer questions like "is this parameter neighborhood more plausible than that one?"[20] You also can't marginalize out nuisance parameters.

Although you might use the posterior pdf to report some kind of mean prediction as discussed above, it almost never makes sense to just *optimize* the posterior pdf. The posterior-optimal parameters are called the "maximum *a posteriori*" (MAP) parameters.[21] Like the maximum likelihood parameters, these only make sense to compute if they are being used to provide an approximation to the posterior pdf (in the form, say, of a mode and a width).[22]

The key idea is that the results of responsible data analysis is not *an answer* but a *distribution over answers*. Data are inherently noisy and incomplete; they never answer your question precisely. So no single number—no maximum-likelihood or MAP value—will adequately represent the result of a data analysis. Results are always pdfs (or full likelihood functions); we must embrace that.

Continuing our example of the model given in equations (28) through (34), if we want to learn the posterior pdf over the parameters $\theta \equiv [a, b]$, we need to choose a prior pdf $p(\theta \,|\, I)$. In principle this should represent our actual prior or exterior knowledge. In practice, investigators often want to "assume nothing" and put a very or infinitely broad prior on the parameters; of course putting a broad prior is *not* equivalent to assuming nothing, it is just as severe an assumption as any other prior. For example, even if you go with a very broad prior on the parameter a, that is a *different* assumption than the same form of very broad prior on a^2 or on $\arctan(a)$. The prior doesn't just set the ranges of parameters, it places a *measure on parameter space.* That's why it is so important.[23]

If you choose an infinitely broad prior pdf, it can become *improper*, in the sense that it can become impossible to satisfy the normalization condition

$$1 \;=\; \int p(\theta \,|\, I)\, \mathrm{d}\theta \quad . \tag{39}$$

The crazy thing is that—although it is not advised—even with improper priors you can still often do inference, because an infinitely broad Gaussian (for example) is a well defined limit of a wide but finite Gaussian, and the posterior pdf can be well behaved in the limit. That is, posterior pdfs can be proper even when prior pdfs are not.

Sometimes your prior knowledge can be very odd. For example, you might be willing to take any slope a over a wide range, but require that the line $y = a\,x + b$ pass through a specific point $(x, y) = (X_0, Y_0)$. Then your prior might look like

$$p(\theta \,|\, I) \;=\; p(a \,|\, I)\, p(b \,|\, a, I) \tag{40}$$

$$p(b \,|\, a, I) \;=\; \delta(b + a\, X_0 - Y_0) \quad , \tag{41}$$

where $p(a \,|\, I)$ is some broad function but $\delta(\cdot)$ is the Dirac delta function, and, implicitly, X_0 and Y_0 are part of the prior information I. These examples all go to show that you have an enormous range of options when you start to

write prior pdfs. In general, you should include all the things you know to be true when you write your priors. With great power comes great responsibility.

In other documents in this series (for example Hogg *et al.* 2010a), more advanced topics will be discussed. One example is the (surprisingly common) situation in which you have far more parameters than data. This sounds impossible, but the rules of probability calculus don't prohibit it, and once you marginalize out most of them you can be left with extremely strong constraints on the parameters you care about. Another example is that in many cases you care about the *prior* on your parameters, not the parameters themselves. Imagine, for example, that you don't know what prior pdf to put on your parameters θ but you want to take it from a family that itself has parameters β. Then you can marginalize out the θ—yes, marginalize out all the parameters we once thought we cared about—using the prior $p(\theta \mid \beta, I)$ and be left with a likelihood for the parameters β of the prior, like so:

$$p(D \mid \beta, I) \;=\; \int p(D \mid \theta, I) \, p(\theta \mid \beta, I) \, \mathrm{d}\theta \quad . \tag{42}$$

The parameters β of the prior pdf are usually called "hyperparameters"; this kind of inference is called "hierarchical".[24]

Exercise 8: Show that if you take the model in equations (28) through (34) and put a Gaussian prior pdf on a and an independent Gaussian prior pdf on b that your posterior pdf for a and b will be a two-dimensional Gaussian. Feel free to use informal or even hand-waving arguments; there are no mathematicians present.

Exercise 9: Take the limit of your answer to Exercise 8 as the width of the a and b prior pdfs go to infinity and show that you still get a well-defined Gaussian posterior pdf. Feel free to use informal or even hand-waving arguments.

Exercise 10: Show that the posterior pdf you get in Exercise 9 is just a rescaling of the likelihood function by a scalar. Two questions: Why *must* there be a re-scaling, from a dimensional point of view? What is the scaling, specifically? You might have to do some linear algebra, for which I won't apologize; it's awesome.

Exercise 11:   Equation (41) implies that the delta function $\delta(q)$ has what dimensions?

Notes

[3]In my world, *all integrals are definite integrals*. The integral is never just the anti-derivative, it is always a definite, finite, area or volume or measure. That has to do, somehow, with the fact that the rules of probability calculus are an application of the rules of measure theory.

[4]To see that the units in equation (3) are $a^{-1} b$, you have to see that the *integration operator* $\int db$ itself has units of b (just as the derivative operator d/db has units of b^{-1}). Why are the units $a^{-1} b$ absurd? Well, they might not be absurd units in general, but they aren't what we want when we integrate a probability distribution.

[5]Division by zero is a huge danger—in principle—when applying Bayes's theorem. In practice, if there is support for the model in your data, or support for the data in your model, you don't hit any zeros. This sounds a little crazy, but if you have data with unmodeled outliers—for example if your noise model can't handle enormous data excursions caused by rare events—you can easily get real data sets that have vanishing probability in a naïve model.

[6]There are no unconditional probabilities! This is because whenever *in practice* you calculate a probability or a pdf, you are always making strong assumptions. Your probabilities are all conditioned on these assumptions.

[7]The "technical reason" here that I treat the denominator in equation (10) as a renormalization constant Z is that when we perform Markov-Chain Monte Carlo sampling methods to obtain samples from $p(a \,|\, b, c)$, we will not need to know Z at all; it is usually (in practice) hard to compute and often unnecessary.

[8]Serious, non-ambiguous mathematicians often distinguish between the

name of the variable and the name of a draw from the probability distribution for the variable, and then instead of $p(a \mid b, c)$ they can write things like $p_{A,B=b,C=c}(a)$, which are unambiguous (or far less ambiguous). This permits, for example, another thing q to be drawn from the same distribution as a by a notation like $p_{A,B=b,C=c}(q)$. In our (very bad) notation $p(q \mid b, c)$ would in general be different from $p(a \mid b, c)$ because we take the meaning of the function from the names of the argument variables. See why that is very bad? Some explicit subscripting policy seems like much better practice and we should probably all adopt something like it, though I won't here.

Another good idea is to make the probability equations be about statements, so instead of writing $p(a \mid b, c)$ you write $p(A = a \mid B = b, C = c)$. This is unambiguous—you can write useful terms like $p(A = q \mid B = b, C = c)$ to deal with the a, q problem—but it makes the equations big and long.

[9]Have you noticed that I am *not* a mathematician?

[10]The word "independent" has many meanings in different contexts of mathematics and probability; I will avoid it in what follows, except in context of independently drawn data points in a generative model of data. I prefer the word "separable" for this situation, because I think it is less ambiguous.

[11]I love the word "heteroscedastic". It means "having heterogeneous noise properties", or that you can't treat every data point as being drawn from the same noise model. Heteroscedasticity is a property of every data set I have ever used.

[12]In many sources, when a variable b is made discrete, because the pdf $p(b \mid a)$ becomes just a set of probabilities, the symbol "p" is often changed to "P". That's not crazy; anything that makes this ambiguous nomenclature less ambiguous is good.

[13]I hate the terminology "likelihood" and "likelihood for parameters" but again, it is so entrenched, it would be a disservice to the reader not to use it. If you really want to go down the rabbit hole, Jaynes (2003) calls the pdf for the data a "frequency distribution" not a "probability distribution" because he *always* sees the data as being *fixed* and reserves the word "probability" for the things that are unknown (that is, being inferred).

[14]The pseudo-marginalization operation in wrong equation (25) has occasionally been done in the literature, and not led to disastrous results. Why

not? It is because that operation, though not permitted, is very very similar to the operation you *in fact* perform when you do the correct marginalization integration given in equation (26) but with a prior pdf $p(\alpha \,|\, \theta, I)$ that is flat in α. Implicitly, *any* marginalization or marginalization-like integral necessarily involves the choice of a prior pdf.

[15]I will say a lot more about this in some subsequent document in this series; in general in real-world problems you need to put a lot of care and attention into the parts of the problem that you *don't* care about; it is a consequence of the need to make precise measurements in the parts of the problem that you *do* care about.

[16]A probabilistic reasoner returns probabilistic information, not just some "best" value. That applies to frequentists and Bayesians alike.

[17]In general, the investigator has a lot of freedom in deciding what to treat as given, prior information, and what to treat as free parameters of the model. If you don't trust your uncertainty variances σ_n^2, make them model parameters! Same with the horizontal positions x_n.

[18]I will say more about what you both gain and lose by going beyond the likelihood in a subsequent document in this series. It relates to the battles between frequentists and Bayesians, which tend towards the boring and unproductive.

[19]Nothing in inference is ever truly *objective* since every generative model involves making choices and assumptions and approximations. In particular, it involves choosing a (necessarily very limited) model space for inference and testing. This informs my view that it is inconsistent—as a probabilistic data analyzer—to be a *realist* (Hogg, 2009). All that said, the likelihood function is necessarily *more* objective than the posterior pdf.

[20]It is hard to live as a frequentist but it is possible. Some of the aforementioned battles between frequentists and Bayesians are set off by overinterpretation of a likelihood as a posterior pdf; in these cases the investigator is indeed unconsciously multiplying in a prior. Other battles are set off by loose language in which a proper frequentist analysis is spoiled by overstatement of the results as making some parameters "more probable" than others.

[21]The MAP parameters include prior information, which provide a measure

on the parameters, so the MAP parameters change as the parameter space is transformed (as coordinate transformations are applied to the parameters). This is not true of the maximum likelihood parameters, because the likelihood framework never makes use of any measure information.

[22]I have seen many cases in which investigators report the MAP parameters instead of the maximum-likelihood (ML) parameters and describe the inference as "Bayesian". That is pretty misleading, because neither the MAP nor the ML is a probabilistic output, and the MAP has additional assumptions in the form of a prior pdf. That said, the MAP value is like a "regularized" version of the ML, and can therefore be valuable in engineering and decision systems; more about this elsewhere.

[23]The insertion of the prior pdf absolutely and necessarily increases the number of assumptions; there is no avoiding that. Bayesians sometimes like to say that frequentists make just as many assumptions as Bayesians do. It isn't true: A principled frequentist—an investigator who only uses the likelihood function and nothing else—genuinely makes fewer assumptions than any Bayesian.

[24]An example of simple hierarchical inference from my own work is Hogg *et al.* (2010b).

References

Hogg, D. W., 2009, "Is cosmology just a plausibility argument?", arXiv:0910.3374

Hogg, D. W., Bovy, J., & Lang, D., 2010a, "Data analysis recipes: Fitting a model to data", arXiv:1008.4686

Hogg, D. W., Myers, A. D., & Bovy, J., 2010b, "Inferring the eccentricity distribution", arXiv:1008.4146

Jaynes, E. T., 2003, *Probability theory: The logic of science* (Cambridge University Press)

Sivia, D. S. & Skilling, J., 2006, *Data analysis: A Bayesian tutorial*, 2ed (Oxford University Press)