> Part Three
> Data-Driven Models

The first two parts of this book have covered techniques to explore the behavior of a model but have assumed that the model is known. Frequently, finding the model in the first place is the most difficult, interesting, and important question. Models can come from introspection or observation (or both); here we turn to the problem of inferring a model from measured data. The model may be used to characterize and classify the data, to generalize from the measurements in order to make predictions about new observations, or most ambitiously to learn something about the rules underlying the observed behavior.

In any data-driven modeling effort the two central tasks are always choosing the functional form of the model, and using the data to determine the adjustable parameters of the model. These are closely connected but can lead to different kinds of errors. Model mismatch errors are those that arise from a model that is unable to represent the data, and model estimation errors come from using incorrect values for the model parameters. Decreasing one kind of error is likely to increases the other kind. This is called a bias/variance tradeoff - if you want less bias in the estimate of a model parameter, it usually costs you more variance. A more flexible model that can better represent the data may also be more easily led astray by noise in the data. Each of the coming chapters covers some kind of optimization to minimize these errors. Since one person's data may be another's noise, it will be important throughout to keep an eye on which principle is being used to optimize what quantity with respect to which error measure, and on how the "best" solution is defined. After all, everything is optimal with respect to something, or conversely no one thing is optimal with respect to everything (the No Free Lunch theorem [Macready \& Wolpert, 1996; Wolpert \& Macready, 1997; Wolpert \& Macready, 2005]).

Chapter 12 introduces the canon of function fitting. Chapters $13-17$ revisit these standard choices in order to uncover the power of less-familiar alternatives, and Chapters 19-21 look at the essential role of time in modeling.

## 12 Function Fitting

The goal of function fitting is to choose values for the parameters in a function to best describe a set of data. There are many possible reasons to do this. If a specific meaningful form for the function with a small number of free parameters is known in advance, this is called parametric fitting, and finding the parameter values themselves may be the goal. For example, an exponential decay constant might be sought to determine a reaction rate. If the form of the function is not known, and so a very flexible function with many free parameters is used, this becomes nonparametric fitting (although this distinction is often vague). One reason to do nonparametric fitting is to try to find and describe underlying trends in noisy data, in which case the fit must build in some prior beliefs and posterior observations about what defines the difference between the signal and the noise. In function approximation there is no noise; the goal is to take known (and perhaps laboriously calculated) values of a function and find a new function that is easier to evaluate and that can interpolate between, or extrapolate beyond, the known values.

It's useful to view function fitting in a context such as the Minimum Description Length principle (MDL) [Rissanen, 1986], or the related Algorithmic Information Theory [Chaitin, 1990]. One extreme is to report the observed data itself as your model. This is not hard to do, but the "model" is very large and has no ability to generalize. Another extreme is report the smallest amount of information possible needed to describe the data, but this may require a great deal of supporting documentation about how to use the model. A tidy computer program is of no use to a Martian unless it comes with a complete description of a computer that can run it. The best model typically lies between these extremes: there is some amount of information about the data, and some about the model architecture. According to MDL, the sum of these two kinds of information taken together should be as small as possible. While this principle cannot be applied directly (like so many other attractive ideas, it includes a solution to the halting problem of deciding if an arbitrary program will terminate, which is known to be impossible [Turing, 1936; Chaitin, 1994]), it is a useful guiding principle that can be made explicit given specific assumptions about a problem.

In this chapter we will look at the basic features of function fitting: the general principles by which data can be used to constrain a model, the (often overlooked) connection with the choice of an error measure, how to fit a model with linear and nonlinear parameters, and the limits on what we can expect to learn from fitting. We will not be particularly concerned with the functional form used; coming chapters will look in much more detail at the representation of data, functions, and optimization strategies.

### 12.1 MODEL ESTIMATION

The general fitting problem has three ingredients: a model architecture (which we'll call $m$ ) that has a set of adjustable parameters $\varphi$, and measured data $d$. The goal is to find values of the parameters that lead to the best agreement (in some sense) between the predictions of the model and the data. An example of $m$ might be a polynomial of a given order, where the $\varphi$ are the coefficients.

A reasonable way to go about finding the best coefficients is to ask for the $\varphi$ that are most likely given the choice of the model and the measured data. This means that we want to find the $\varphi$ that maximizes $p(\varphi \mid d, m)$. Using Bayes' rule (equation 6.11), our job is then to find

$$
\begin{align*}
\max _{\varphi} p(\varphi \mid d, m) & =\max _{\varphi} \frac{p(\varphi, d, m)}{p(d, m)} \\
& =\max _{\varphi} \frac{p(d \mid \varphi, m) p(\varphi \mid m) p(m)}{p(d \mid m) p(m)} \\
& =\max _{\varphi} \frac{p(d \mid \varphi, m) p(\varphi \mid m)}{p(d \mid m)} \\
& =\max _{\varphi} \frac{p(d \mid \varphi, m) p(\varphi \mid m)}{\int_{\varphi} p(d, \varphi \mid m) d \varphi} \\
& =\max _{\varphi} \frac{p(d \mid \varphi, m) p(\varphi \mid m)}{\int_{\varphi} p(d \mid \varphi, m) p(\varphi \mid m) d \varphi} \\
& =\max _{\varphi} \frac{\text { likelihood } \times \operatorname{prior}}{\text { evidence }} \tag{12.1}
\end{align*}
$$

The probability has been factored into three terms. The likelihood measures the match between the data and the predictions of the model with the coefficents, based on an error model. The prior introduces advance beliefs about which values of the coefficients are reasonable and which are not. And the evidence measures how well the model can describe the data.

If you solve equation (12.1) then you are an official card-carrying Bayesian [Bernardo \& Smith, 1994; Kruschke, 2010]. The reason that there are not too many of them around is that solving equation (12.1) represents a lot of work. First of all, it's necessary to explicitly put priors on every parameter that is used. Then, the integration for the evidence is over all values of all the parameters, which can be an enormous computational task for a large model. Although efficient techniques have been developed for these kinds of integrals using Monte-Carlo sampling techniques that replace exact integration with a probabilistic approximation [Besag et al., 1995], they are still computationally intensive. Finally, the maximization over parameters is just an inner loop; the best description is given by a maximization over model architectures as well, or, even better, over the combined outputs from multiple models [Burnham \& Anderson, 2002].

Much of the work in Bayesian model estimation goes into the integration for the evidence term. But this does not affect a single maximization over $\varphi$; it comes in making comparisons among competing model architectures. If we decide in advance that we are going to stick with one architecture then equation (12.1) can be simplified by dropping
the conditioning on the model:

$$
\begin{equation*}
\max _{\varphi} p(\varphi \mid d)=\max _{\varphi} \frac{p(d \mid \varphi) p(\varphi)}{p(d)} \tag{12.2}
\end{equation*}
$$

Now the evidence term has become a simple prior on the likelihood of the data set. Even this can usually be dropped; it's relevant only in combining multiple data sets of varying pedigrees. Finding parameters with (12.2) is called Maximum A Posteriori estimation (MAP).

MAP still requires putting a prior on the parameters. This is a very powerful idea, to be explored in the next chapter, but if we make the simplest choice of a uniform prior $p(\varphi)=p(d)=1$ then we're left with

$$
\begin{equation*}
\max _{\varphi} p(\varphi \mid d)=\max _{\varphi} p(d \mid \varphi) \tag{12.3}
\end{equation*}
$$

This is the easiest kind of model estimation of all, called Maximum Likelihood (ML). That is what we will now apply.

### 12.2 LEAST SQUARES

Let's assume that we are given a set of $N$ noisy measurements of a quantity $y_{n}$ as a function of a variable $x_{n}$, and we seek to find values for coefficients $\varphi$ in a function $y_{n}=y\left(x_{n}, \varphi\right)$ that describes their relationship (the generalization to vector variables will be straightforward). In Section 6.1.2 we learned that in the absence of any other information the Central Limit Theorem tells us that the most reasonable choice for the distribution of a random variable is Gaussian, and so we will make that choice for the distribution of errors in $y_{n}$. Problem 11.3 will use an entropy argument to reach the same conclusion. In practice, many systems choose to ignore this insight and have nonGaussian distributions; the real reason why Gaussianity is so commonly (and frequently implictly) assumed is that it leads to a particularly simple and useful error model: least squares.

If the errors do have a Gaussian distribution around the true value $y\left(x_{n}, \varphi\right)$ then the probability to observe a value between $y$ and $y+d y$ is given by a Gaussian centered on the correct value

$$
\begin{equation*}
p(y) d y=\frac{1}{\sqrt{2 \pi \sigma_{n}^{2}}} e^{-\left[y-y\left(x_{n}, \varphi\right)\right]^{2} /\left(2 \sigma_{n}^{2}\right)} d y \tag{12.4}
\end{equation*}
$$

The variance $\sigma_{n}^{2}$ might depend on quantities such as the noise in a photodetector or the number of samples that are measured.

We will further assume that the errors between samples are independent as well as identically distributed (iid). This means that the probability to see the entire data set is given by the product of the probabilities to see each point,

$$
\begin{equation*}
p(\text { data } \mid \text { model })=\prod_{n=1}^{N} \frac{1}{\sqrt{2 \pi \sigma_{n}^{2}}} e^{-\left[y_{n}-y\left(x_{n}, \varphi\right)\right]^{2} /\left(2 \sigma_{n}^{2}\right)} \tag{12.5}
\end{equation*}
$$

We seek the $\varphi$ that maximizes this probability. If $p$ is maximal then so is its logarithm
(the log-likelihood), and since the log of a product is equal to the sum of the logs, this becomes

$$
\begin{equation*}
-\log p(\text { data } \mid \text { model })=\sum_{n=1}^{N} \frac{\left[y_{n}-y\left(x_{n}, \varphi\right)\right]^{2}}{2 \sigma_{n}^{2}}+\frac{1}{2} \log \left(2 \pi \sigma_{n}^{2}\right) \tag{12.6}
\end{equation*}
$$

Because we've moved the minus sign to the left hand side we now want to find the $\varphi$ that minimizes the right hand side. The first term measures the distance between the data and the model, and the second one catches us if we try to cheat and make a model with a huge variance that explains everything equally well (or poorly). We can drop the second term since it does not depend on the parameters $\varphi$ that we are adjusting, and so we want to find the values that satisfy

$$
\begin{equation*}
\min _{\varphi} \sum_{n=1}^{N} \frac{\left[y_{n}-y\left(x_{n}, \varphi\right)\right]^{2}}{2 \sigma_{n}^{2}} \tag{12.7}
\end{equation*}
$$

If the variances $\sigma_{n}^{2}$ are constant (in particular, if we set $\sigma_{n}^{2}=1$ when we have no idea at all what it should be) this reduces to

$$
\begin{equation*}
\min _{\varphi} \sum_{n=1}^{N}\left[y_{n}-y\left(x_{n}, \varphi\right)\right]^{2} \tag{12.8}
\end{equation*}
$$

This is the familiar least squares error measure. It is the maximum likelihood estimator for data with normally distributed errors, but it is used much more broadly because it is simple, convenient, and frequently not too far off from an optimal choice for a particular problem. An example of where least squares might be a bad choice for an error measure is a bi-modal data set that has two peaks. The least squares error is minimized by a point between the peaks, but such a point has very little probability of actually occurring in the data set.

Instead of the square of the deviation between the model and the data, other powers can be used as an error measure. The first power (the magnitude of the difference) is the maximum likelihood estimate if the errors are distributed exponentially, and higher powers place more emphasis on outliers.

### 12.3 LINEAR LEAST SQUARES

Once we've chosen our error measure we need to find the parameters for the distribution that minimizes it. Perhaps the most important example of such a technique is linear least squares, because it is straightforward to implement and broadly applicable.

To do a least squares fit we will start by expanding our unknown function as a linear sum of $M$ known basis functions $f_{m}$

$$
\begin{equation*}
y(x)=\sum_{m=1}^{M} a_{m} f_{m}(x) \tag{12.9}
\end{equation*}
$$

We want to find the coefficients $a_{m}$ that minimize the sum of the squared errors between this model and a set of $N$ given observations $y_{n}\left(x_{n}\right)$. The basis functions $f_{m}$ need not be orthogonal, but they must not be linear (otherwise the sum would be trivial); it is the
coefficients $a_{m}$ that enter linearly. For example, the $f_{m}$ could be polynomial terms, with the $a_{m}$ as the coefficients of the polynomial.

A least squares fit can be written as a matrix problem

$$
\left(\begin{array}{cccc}
f_{1}\left(x_{1}\right) & f_{2}\left(x_{1}\right) & \cdots & f_{M}\left(x_{1}\right)  \tag{12.10}\\
f_{1}\left(x_{2}\right) & f_{2}\left(x_{2}\right) & \cdots & f_{M}\left(x_{2}\right) \\
f_{1}\left(x_{3}\right) & f_{2}\left(x_{3}\right) & \cdots & f_{M}\left(x_{3}\right) \\
\vdots & \vdots & \vdots & \vdots \\
f_{1}\left(x_{N-1}\right) & f_{2}\left(x_{N-1}\right) & \cdots & f_{M}\left(x_{N-1}\right) \\
f_{1}\left(x_{N}\right) & f_{2}\left(x_{N}\right) & \cdots & f_{M}\left(x_{N}\right)
\end{array}\right)\left(\begin{array}{c}
a_{1} \\
a_{2} \\
\vdots \\
a_{M}
\end{array}\right)=\left(\begin{array}{c}
y_{1} \\
y_{2} \\
y_{3} \\
\vdots \\
y_{N-1} \\
y_{N}
\end{array}\right)
$$

If we have the same number of free parameters as data points then the matrix will be square, and so the coefficients can be found by inverting the matrix and multiplying it by the observations. As long as the matrix is not singular (which would happen if our basis functions were linearly dependent), this inversion could be done exactly and our fit would pass through all of the data points. If our data are noisy this is a bad idea; we'd like to have many more observations than we have model parameters. We can do this if we use the pseudo-inverse of the matrix to minimize the least squared error. The Singular Value Decomposition (SVD) is a powerful technique that solves this (as well as many other problems).

### 12.3.1 Singular Value Decomposition

For a general linear equation $\mathbf{A} \cdot \vec{v}=\vec{b}$, the space of all possible vectors $\vec{b}$ for which the equation is solvable is the range of $\mathbf{A}$. The dimension of the range (i.e., the number of vectors needed to form a basis of the range) is called the rank of $\mathbf{A}$. There is an associated homogeneous problem $\mathbf{A} \cdot \vec{v}=0$; the vectors $\vec{v}$ that satisfy the homogeneous equation lie in the nullspace of $\mathbf{A}$. If there is no nullspace then the matrix is of full rank (which is equal to the number of columns of $\mathbf{A}$ ).

If A is an arbitrary $N \times M$ matrix, an important result from linear algebra is that it can always be written in the form [Golub \& Loan, 1996]

$$
(\mathbf{A})=(\mathbf{U}) \underbrace{\left(\begin{array}{llll}
w_{1} & & & \mathbf{0}  \tag{12.11}\\
& w_{2} & & \\
& & \ddots & \\
\mathbf{0} & & w_{M}
\end{array}\right)}_{\mathbf{W}}\left(\begin{array}{l}
\mathbf{V}^{T} \\
\end{array}\right)
$$

where $\mathbf{U}$ and $\mathbf{V}$ are orthogonal matrices whose inverse is equal to their transpose $\mathbf{U}^{T} \cdot \mathbf{U}=$ $\mathbf{V} \cdot \mathbf{V}^{T}=\mathbf{I}$ (where $\mathbf{I}$ is the $M \times M$ identity matrix). This is called the Singular Value Decomposition (SVD) of the matrix, and the elements of the $M \times M$ diagonal matrix $\mathbf{W}$ are the singular values $w_{i}$.

The reason that the SVD is so important is that the columns of $\mathbf{U}$ associated with nonzero singular values $\left(w_{i} \neq 0\right)$ form an orthonormal basis for the range of $\mathbf{A}$, and the columns of $\mathbf{V}$ associated with $w_{i}=0$ form an orthonormal basis for the nullspace of
A. The singular values $w_{i}$ give the lengths of the principal axes of the hyper-ellipsoid defined by $\mathbf{A} \cdot \vec{x}$, where $\vec{x}$ lies on a hyper-sphere $|\vec{x}|^{2}=1$.

In terms of the SVD, the solution to $\mathbf{A} \cdot \vec{v}=\vec{b}$ is

$$
\begin{equation*}
\vec{v}=\mathbf{V} \cdot \mathbf{W}^{-1} \cdot \mathbf{U}^{T} \cdot \vec{b} \tag{12.12}
\end{equation*}
$$

(the errors in this are discussed in Section 12.5). Since $\mathbf{W}$ is diagonal, its inverse $\mathbf{W}^{-1}$ is also diagonal and is found by replacing each diagonal element by its inverse. This sounds like a recipe for disaster since we just saw that $w_{i}=0$ for elements of the nullspace. The odd solution to this problem is simply to declare that $1 / 0=0$, and set to zero the diagonal elements of $\mathbf{W}^{-1}$ corresponding to $w_{i}=0$. To see how this apparent nonsense works, let's look for another solution to $\mathbf{A} \cdot \vec{v}=\vec{b}$. Assume that $\vec{v}$ is found according to the prescription for zeroing singular values in equation (12.12). We can add an arbitrary vector $\vec{v}^{\prime}$ from the nullspace $\left(\mathbf{A} \cdot \vec{v}^{\prime}=0\right)$ and still have a solution. The magnitude of the sum of these vectors is

$$
\begin{align*}
\left|\vec{v}+\vec{v}^{\prime}\right| & =\left|\mathbf{V} \cdot \mathbf{W}^{-1} \cdot \mathbf{U}^{T} \cdot \vec{b}+\vec{v}^{\prime}\right| \\
& =\left|\mathbf{V} \cdot\left(\mathbf{W}^{-1} \cdot \mathbf{U}^{T} \cdot \vec{b}+\mathbf{V}^{T} \cdot \vec{v}^{\prime}\right)\right| \\
& =\left|\mathbf{W}^{-1} \cdot \mathbf{U}^{T} \cdot \vec{b}+\mathbf{V}^{T} \cdot \vec{v}^{\prime}\right| . \tag{12.13}
\end{align*}
$$

Multiplication by $\mathbf{V}$ was eliminated in the last line because $\mathbf{V}$ is an orthogonal matrix and hence does not change the magnitude of the answer (equation 13.7). The magnitude of $\left|\vec{v}+\vec{v}^{\prime}\right|$ is made up of the sum of two vectors. The first one will have its $i$ th element equal to 0 for every vanishing singular value $w_{i}=0$ because of the rule for zeroing these elements of $\mathbf{W}^{-1}$. On the other hand, since $\vec{v}^{\prime}$ is in the nullspace,

$$
\begin{align*}
\mathbf{A} \cdot \vec{v}^{\prime} & =\overrightarrow{0} \\
\mathbf{U} \cdot \mathbf{W} \cdot \mathbf{V}^{T} \cdot \vec{v}^{\prime} & =\overrightarrow{0} \\
\mathbf{W} \cdot \mathbf{V}^{T} \cdot \vec{v}^{\prime} & =\mathbf{U}^{T} \cdot \overrightarrow{0} \\
\mathbf{W} \cdot \mathbf{V}^{T} \cdot \vec{v}^{\prime} & =\overrightarrow{0} \tag{12.14}
\end{align*}
$$

This means that the $i$ th component of the vector $\mathbf{V}^{T} \cdot \vec{v}^{\prime}$ must equal 0 for every $w_{i} \neq 0$. Returning to the last line of equation (12.13), the left hand term can be nonzero only if $w_{i} \neq 0$, and the right hand term can be nonzero only if $w_{i}=0$ : these two vectors are orthogonal. Therefore, the magnitude of their sum is a minimum if $\vec{v}^{\prime}=\overrightarrow{0}$. Adding any component from the nullspace to $\vec{x}$ increases its magnitude. This means that for an underdetermined problem (i.e., one in which there is a nullspace), the SVD along with the rule for zeroing elements in $\mathbf{W}^{-1}$ chooses the answer with the smallest magnitude (i.e., no component in the nullspace).

Now let's look at what the SVD does for an overdetermined problem. This is the case that we care about for fitting data, where we have more measurements than free parameters. We can no longer hope for an exact solution, but we can look for one that minimizes the residual

$$
\begin{equation*}
|\mathbf{A} \cdot \vec{v}-\vec{b}|=\sqrt{(\mathbf{A} \cdot \vec{v}-\vec{b})^{2}} \tag{12.15}
\end{equation*}
$$

Let's once again choose $\vec{v}$ by zeroing singular values in equation (12.12), and see what happens to the residual if we add an arbitrary vector $\vec{v}^{\prime}$ to $i t$. This adds an error of
$\overrightarrow{b^{\prime}}=\mathbf{A} \cdot \vec{v}^{\prime}:$

$$
\begin{align*}
\left|\mathbf{A} \cdot\left(\vec{v}+\vec{v}^{\prime}\right)-\vec{b}\right| & =\left|\mathbf{A} \cdot \vec{v}+\vec{b}^{\prime}-\vec{b}\right| \\
& =\left|\left(\mathbf{U} \cdot \mathbf{W} \cdot \mathbf{V}^{T}\right) \cdot\left(\mathbf{V} \cdot \mathbf{W}^{-1} \cdot \mathbf{U}^{T} \cdot \vec{b}\right)+\overrightarrow{b^{\prime}}-\vec{b}\right| \\
& =\left|\left(\mathbf{U} \cdot \mathbf{W} \cdot \mathbf{W}^{-1} \cdot \mathbf{U}^{T}-\mathbf{I}\right) \cdot \vec{b}+\vec{b}^{\prime}\right| \\
& =\left|\mathbf{U} \cdot\left[\left(\mathbf{W} \cdot \mathbf{W}^{-1}-\mathbf{I}\right) \cdot \mathbf{U}^{T} \cdot \vec{b}+\mathbf{U}^{T} \cdot \vec{b}^{\prime}\right]\right| \\
& =\left|\left(\mathbf{W} \cdot \mathbf{W}^{-1}-\mathbf{I}\right) \cdot \mathbf{U}^{T} \cdot \vec{b}+\mathbf{U}^{T} \cdot \vec{b}^{\prime}\right| . \tag{12.16}
\end{align*}
$$

Here again the magnitude is the sum of two vectors. ( $\mathbf{W} \cdot \mathbf{W}^{-1}-\mathbf{I}$ ) is a diagonal matrix, with nonzero entries for $w_{i}=0$, and so the elements of the left term can be nonzero only where $w_{i}=0$. The right hand term can be rewritten as follows:

$$
\begin{align*}
\mathbf{A} \cdot \vec{v}^{\prime} & =\overrightarrow{b^{\prime}} \\
\mathbf{U} \cdot \mathbf{W} \cdot \mathbf{V}^{T} \cdot \vec{v}^{\prime} & =\vec{b}^{\prime} \\
\mathbf{W} \cdot \mathbf{V}^{T} \cdot \vec{v}^{\prime} & =\mathbf{U}^{T} \cdot \vec{b}^{\prime} \tag{12.17}
\end{align*}
$$

The $i$ th component can be nonzero only where $w_{i} \neq 0$. Once again, we have the sum of two vectors, one of which is nonzero only where $w_{i}=0$, and the other where $w_{i} \neq 0$, and so these vectors are orthogonal. Therefore, the magnitude of the residual is a minimum if $\vec{v}^{\prime}=0$, which is the choice that SVD makes. Thus, the SVD finds the vector that minimizes the least squares residual for an overdetermined problem.

The computational cost of finding the SVD of an $N \times M$ matrix is $\mathcal{O}\left(N M^{2}+M^{3}\right)$. This is comparable to the ordinary inversion of an $M \times M$ matrix, which is $\mathcal{O}\left(M^{3}\right)$, but the prefactor is larger. Because of its great practical significance, good SVD implementations are available in most mathematical packages. We can now see that it is ideal for solving equation (12.10). If the basis functions are chosen to be polynomials, the matrix to be inverted is called a Vandermonde matrix. For example, let's say that we want to fit a 2D bilinear model $z=a_{0}+a_{1} x+a_{2} y+a_{3} x y$. Then we must invert

$$
\left(\begin{array}{cccc}
1 & x_{1} & y_{1} & x_{1} y_{1}  \tag{12.18}\\
1 & x_{2} & y_{2} & x_{2} y_{2} \\
1 & x_{3} & y_{3} & x_{3} y_{3} \\
\vdots & \vdots & \cdots & \vdots \\
1 & x_{N-1} & y_{N-1} & x_{N-1} y_{N-1} \\
1 & x_{N} & y_{N} & x_{N} y_{N}
\end{array}\right)\left(\begin{array}{c}
a_{0} \\
a_{1} \\
a_{2} \\
a_{3}
\end{array}\right)=\left(\begin{array}{c}
z_{1} \\
z_{2} \\
z_{3} \\
\vdots \\
z_{N-1} \\
z_{N}
\end{array}\right)
$$

If the matrix is square $(M=N)$, the solution can go through all of the data points. These are interpolating polynomials, like those we used for finite elements. A rectangular matrix $(M<N)$ is the relevant case for fitting data. If there are any singular values near zero (given noise and the finite numerical precision they won't be exactly zero) it means that some of our basis functions are nearly linearly dependent and should be removed from the fit. If they are left in, SVD will find the set of coefficients with the smallest overall magnitude, but it is best for numerical accuracy (and convenience in using the fit) to remove the terms with small singular values. The SVD inverse will then provide the coefficients that give the best least squares fit. Choosing where to cut off the spectrum of singular values depends on the context; a reasonable choice is the largest singular value
weighted by the computer's numerical precision, or by the fraction of noise in the data [Golub \& Loan, 1996].

In addition to removing small singular values to eliminate terms which are weakly determined by the data, another good idea is to scale the expansion terms so that the magnitudes of the coefficients are comparable, or even better to rescale the data to have unit variance and zero mean (almost always a good idea in fitting). If 100 is a typical value for $x$ and $y$, then $10^{4}$ will be a typical value for their product. This means that the coefficient of the $x y$ term must be $\sim 100$ times smaller than the coefficient of the $x$ or $y$ terms. For higher powers this problem will be even worse, ranging from an inconvenience in examining the output from the fit, to a serious loss of numerical precision as a result of multiplying very large numbers by very small numbers.

### 12.4 NONLINEAR LEAST SQUARES

Using linear least squares we were able to find the best set of coefficients $\vec{a}$ in a single step (the SVD inversion). The price for this convenience is that the coefficients cannot appear inside the basis functions. For example, we could use Gaussians as our bases, but we would be able to vary only their amplitude and not their location or variance. It would be much more general if we could write

$$
\begin{equation*}
y(x)=\sum_{m=1}^{M} f_{m}\left(x, \vec{a}_{m}\right) \tag{12.19}
\end{equation*}
$$

where the coefficients are now inside the nonlinear basis functions. We can still seek to minimize the error

$$
\begin{equation*}
\chi^{2}(\vec{a})=\sum_{n=1}^{N}\left(\frac{y_{n}-y\left(x_{n}, \vec{a}\right)}{\sigma_{n}}\right)^{2} \tag{12.20}
\end{equation*}
$$

but we will now need to do an iterative search to find the best solution, and we are no longer guaranteed to find it (see Section 14.5).

The basic techniques for nonlinear fitting that we'll cover in this chapter are based on the insight that we may not be able to invert a matrix to find the best solution, but we can evaluate the error locally and then move in a direction that improves it. The gradient of the error is

$$
\begin{equation*}
\left(\nabla \chi^{2}\right)_{k}=\frac{\partial \chi^{2}}{\partial a_{k}}=-2 \sum_{n=1}^{N} \frac{y_{n}-y\left(x_{n}, \vec{a}\right)}{\sigma_{i}^{2}} \frac{\partial y\left(x_{n}, \vec{a}\right)}{\partial a_{k}} \tag{12.21}
\end{equation*}
$$

and its second derivative (the Hessian) is

$$
\begin{align*}
\mathbf{H}_{k l} & =\frac{\partial^{2} \chi^{2}}{\partial a_{k} \partial a_{l}}  \tag{12.22}\\
& =2 \sum_{n=1}^{N} \frac{1}{\sigma_{i}^{2}}\left[\frac{\partial y\left(x_{n}, \vec{a}\right)}{\partial a_{k}} \frac{\partial y\left(x_{n}, \vec{a}\right)}{\partial a_{l}}-\left[y_{n}-y\left(x_{n}, \vec{a}\right)\right] \frac{\partial^{2} y\left(x_{n}, \vec{a}\right)}{\partial a_{l} \partial a_{k}}\right]
\end{align*}
$$

Since the second term in the Hessian depends on the sum of terms proportional to the
residual between the model and the data, which should be small and can change sign, it is customary to drop this term in nonlinear fitting.

From a starting guess for $\vec{a}$, we can update the estimate by the method of steepest descent or gradient descent, taking a step in the direction in which the error is decreasing most rapidly

$$
\begin{equation*}
\vec{a}_{\text {new }}=\vec{a}_{\text {old }}-\alpha \nabla \chi^{2}\left(\vec{a}_{\text {old }}\right) \tag{12.23}
\end{equation*}
$$

where $\alpha$ determines how big a step we make. On the other hand, $\chi^{2}$ can be expanded around a point $\vec{a}_{0}$ to second order as

$$
\begin{equation*}
\chi^{2}(\vec{a})=\chi^{2}\left(\vec{a}_{0}\right)+\left[\nabla \chi^{2}\left(\vec{a}_{0}\right)\right] \cdot\left(\vec{a}-\vec{a}_{0}\right)+\frac{1}{2}\left(\vec{a}-\vec{a}_{0}\right) \cdot \mathbf{H} \cdot\left(\vec{a}-\vec{a}_{0}\right) \tag{12.24}
\end{equation*}
$$

which has a gradient

$$
\begin{equation*}
\nabla \chi^{2}(\vec{a})=\nabla \chi^{2}\left(\vec{a}_{0}\right)+\mathbf{H} \cdot\left(\vec{a}-\vec{a}_{0}\right) \tag{12.25}
\end{equation*}
$$

The minimum $\left(\nabla \chi^{2}(\vec{a})=0\right)$ can therefore be found by iterating

$$
\begin{equation*}
\vec{a}_{\text {new }}=\vec{a}_{\text {old }}-\mathbf{H}^{-1} \cdot \nabla \chi^{2}\left(\vec{a}_{\text {old }}\right) \tag{12.26}
\end{equation*}
$$

(this is Newton's method). Either of these techniques lets us start with an initial guess for $\vec{a}$ and then successively refine it.

### 12.4.1 Levenberg-Marquardt Method

Far from a minimum Newton's method is completely unreliable: the local slope may shoot the new point further from the minimum than the old one was. On the other hand, near a minimum Newton's method converges very quickly and gradient descent slows to a crawl since the gradient being descended is disappearing. A natural strategy is to use gradient descent far away, and then switch to Newton's method close to a minimum. But how do we decide when to switch between them, and how large should the gradient descent steps be? The Levenberg-Marquardt method [Marquardt, 1963] is a clever solution to these questions, and is the most common method used for nonlinear least squares fitting.

We can use the Hessian to measure the curvature of the error surface, taking small gradient descent steps if the surface is curving quickly. Using the diagonal elements alone in the Hessian to measure the curvature is suggested by the observation that this gives the correct units for the scale factor $\alpha$ in equation (12.23):

$$
\begin{equation*}
\delta a_{i}=-\frac{1}{\lambda \mathbf{H}_{i i}} \frac{\partial \chi^{2}}{\partial a_{i}} \tag{12.27}
\end{equation*}
$$

where $\lambda$ is a new dimensionless scale factor. If we use this weighting for gradient descent, we can then combine it with Newton's method by defining a new matrix

$$
\begin{align*}
& M_{i i}=\frac{1}{2} \frac{\partial^{2} \chi^{2}}{\partial a_{i}^{2}}(1+\lambda) \\
& M_{i j}=\frac{1}{2} \frac{\partial^{2} \chi^{2}}{\partial a_{i} \partial a_{j}} \quad(i \neq j) \tag{12.28}
\end{align*}
$$

If we use this to take steps given by

$$
\begin{equation*}
\mathbf{M} \cdot \delta \vec{a}=-\nabla \chi^{2} \tag{12.29}
\end{equation*}
$$

or

$$
\begin{equation*}
\delta \vec{a}=-\mathbf{M}^{-1} \cdot \nabla \chi^{2} \tag{12.30}
\end{equation*}
$$

when $\lambda=0$ this just reduces to Newton's method. On the other hand, if $\lambda$ is very large then the diagonal terms will dominate, which is just gradient descent (equation 12.23). $\lambda$ controls an interpolation between steepest descent and Newton's method. To use the Levenberg-Marquardt method, $\lambda$ starts off moderately large. If the step improves the error, $\lambda$ is decreased (Newton's method is best near a minimum), and if the step increases the error then $\lambda$ is increased (gradient descent is better).
$\chi^{2}$ can easily have many minima, but the Levenberg-Marquardt method will find only the local minimum closest to the starting condition. For this reason, a crucial sanity check is to plot the fitting function with the starting parameters and compare it with the data. If it isn't even close, it is unlikely that Levenberg-Marquardt will converge to a useful answer. It is possible to improve its performance in these cases by adding some kind of randomness that lets it climb out of small minima. If a function is hard to fit because it has very many local minima, or the parameter space is so large that it is hard to find sane starting values, then a technique that is better at global searching is called for. These extensions will be covered in Chapter 15.

### 12.5 ERRORS

The preceeding fitting procedures find the best values for adjustable parameters, but there's no guarantee that the best is particularly good. As important as finding the values is estimating their errors. These come in two flavors: statistical errors from the experimental sampling procedure, and systematic errors due to biases in the measuring and modeling process. There are a number of good techniques for estimating the former; bounding the latter can be more elusive and requires a detailed analysis of how the data are acquired and handled.

It's frequently the case that there is an error model for the measurements known in advance, such as counting statistics or the presence of Johnson noise in an amplifier [Gershenfeld, 2000]. For a linear fitting procedure it is then straightforward to propagate this through the calculation. If there is an error $\vec{\zeta}$ in a measurement, it causes an error $\vec{\eta}$ in an SVD fit:

$$
\begin{equation*}
\vec{v}+\vec{\eta}=\mathbf{V} \cdot \mathbf{W}^{-1} \cdot \mathbf{U}^{T} \cdot(\vec{b}+\vec{\zeta}) \tag{12.31}
\end{equation*}
$$

Since this is a linear equation, these random variables are related by

$$
\begin{equation*}
\vec{\eta}=\mathbf{V} \cdot \mathbf{W}^{-1} \cdot \mathbf{U}^{T} \cdot \vec{\zeta} \tag{12.32}
\end{equation*}
$$

If the components of $\vec{\zeta}$ are zero-mean with a variance $\sigma_{\zeta}^{2}$ then the variance in the $i$ th component of $\vec{\eta}$ is

$$
\begin{aligned}
\sigma_{\eta, i}^{2} & =\left\langle\eta_{i} \eta_{i}\right\rangle \\
& =\left\langle\sum_{j} V_{i j} \frac{1}{w_{j}} \sum_{k} U_{k j} \zeta_{k} \sum_{l} V_{i l} \frac{1}{w_{l}} \sum_{m} U_{m l} \zeta_{m}\right\rangle
\end{aligned}
$$

$$
\begin{align*}
& =\sum_{j} \sum_{l} V_{i j} V_{i l} \frac{1}{w_{j}} \frac{1}{w_{l}} \sum_{k} \sum_{m} U_{k j} U_{m l} \underbrace{\left\langle\zeta_{k} \zeta_{m}\right\rangle}_{\sigma_{\zeta}^{2} \delta_{k m}} \\
& =\sum_{j} \sum_{l} V_{i j} V_{i l} \frac{1}{w_{j}} \frac{1}{w_{l}} \underbrace{\sum_{k} U_{k j} U_{k l} \sigma_{\zeta}^{2}}_{\delta_{j l}} \\
& =\sigma_{\zeta}^{2} \sum_{j} \frac{V_{i j}^{2}}{w_{j}^{2}} \tag{12.33}
\end{align*}
$$

In the second line we've assumed that the measurement errors are an uncorrelated random variable with a fixed variance $\sigma_{b}^{2}$; if this is not the case then their covariance must be carried through the calculation. In the third line we've used the orthonormality of $\mathbf{U}$. Therefore, the error in the fit relative to that in the data is

$$
\begin{equation*}
\frac{\sigma_{\eta, i}^{2}}{\sigma_{\zeta}^{2}}=\sum_{j} \frac{V_{i j}^{2}}{w_{j}^{2}} \tag{12.34}
\end{equation*}
$$

Absent such insight into an error model for the measurements it's necessary to use the observations themselves to estimate the fitting errors. If you are in the fortunate position of being able to generate unlimited amounts of data this can be done following its definition by analyzing an ensemble of independent data sets and then reporting the distribution of fitting results. But in the much more likely circumstance of limited data it's not possible to produce an ensemble of measurements.

Or is it? Bootstrap resampling is based on the apparently-circular reasoning that the data set is drawn from the distribution that describes it and hence drawing from the data approximates that distribution. This is done by sampling with replacement. A random number generator is used to choose elements of the data set, with an element remaining in the original data set after it is chosen so that it can reappear in the derived one multiple times. Random selection continues until a new data set has been produced of the same size as the original one. This one uses the same elements, but it is an independent sampling of them. These data can then be fit, and the resampling done again as many times as desired. Problem 11.1 looks at an example of this.

Bootstrap error estimation was originally statistically suspect because it appears to violate the deeply-held belief that data should not be reused in analysis. But, over time, both theoretical and experimental work has shown that, while it is not as reliable as using truly independent data sets, bootstrap can provide useful information about errors not available in a single fit [Efron \& Tibshirani, 1994]. Its typical performance can in fact be quite good, although the worst-case errors in the errors can be quite bad.

Beyond bootstrap, in fitting functions that will be used for making forecasts the most important error is how well the model generalizes on data not seen in the fitting procedure. This can be evaluated through cross-validation, discussed in Section 14.4.

### 12.6 ESTIMATION, FISHER INFORMATION, AND THE CRAMER-RAO INEQUALITY

We've seen increasingly powerful techniques to extract functions and errors from data. Is there no limit to this cleverness? Unfortunately, and not surprisingly, there is indeed a limit on how much information about unknown parameters can be extracted from a set of measurements. This chapter closes with a view of the information in a probability distribution that sets a limit on the accuracy of measurements.

Let $p_{\alpha}(x)$ be a probability distribution that depends on a parameter $\alpha$ (such as the variance of a Gaussian); the goal is to estimate the value of $\alpha$ from a series of measurements of $x$. Let $f\left(x_{1}, x_{2}, \ldots, x_{N}\right)$ be the estimator of $\alpha$. It is biased if $\left\langle f\left(x_{1}, x_{2}, \ldots, x_{N}\right)\right\rangle \neq \alpha$, and it is consistent if $\lim _{N \rightarrow \infty} f\left(x_{1}, x_{2}, \ldots, x_{N}\right)=\alpha$. An estimator $f_{1}$ dominates $f_{2}$ if $\left\langle\left(f_{1}\left(x_{1}, x_{2}, \ldots, x_{N}\right)-\alpha\right)^{2}\right\rangle \leq\left\langle\left(f_{2}\left(x_{1}, x_{2}, \ldots, x_{N}\right)-\alpha\right)^{2}\right\rangle$. This raises the question of what is the minimum variance possible for an unbiased estimator of $\alpha$ ? The answer is given by the Cramér-Rao bound.

Start by defining the score:

$$
\begin{equation*}
V=\frac{\partial}{\partial \alpha} \log p_{\alpha}(x)=\frac{\partial_{\alpha} p_{\alpha}(x)}{p_{\alpha}(x)} \tag{12.35}
\end{equation*}
$$

The expected value of the score is

$$
\begin{align*}
\langle V\rangle & =\int_{-\infty}^{\infty} p_{\alpha}(x) \frac{\partial_{\alpha} p_{\alpha}(x)}{p_{\alpha}(x)} d x \\
& =\int_{-\infty}^{\infty} \partial_{\alpha} p_{\alpha}(x) d x \\
& =\partial_{\alpha} \int_{-\infty}^{\infty} p_{\alpha}(x) d x \\
& =\partial_{\alpha} 1 \\
& =0 \tag{12.36}
\end{align*}
$$

This means that $\sigma^{2}(V)=\left\langle V^{2}\right\rangle$. The variance of the score is called the Fisher information:

$$
\begin{equation*}
J(\alpha)=\left\langle\left[\partial_{\alpha} \log p_{\alpha}(x)\right]^{2}\right\rangle \tag{12.37}
\end{equation*}
$$

The score for a set of independent, identically distributed variables is the sum of the individual scores

$$
\begin{align*}
V\left(x_{1}, x_{2}, \ldots, x_{N}\right) & =\partial_{\alpha} \log p_{\alpha}\left(x_{1}, x_{2}, \ldots, x_{N}\right) \\
& =\sum_{n=1}^{N} \partial_{\alpha} \log p_{\alpha}\left(x_{n}\right) \\
& =\sum_{n=1}^{N} V\left(x_{n}\right) \tag{12.38}
\end{align*}
$$

and so the Fisher information for the set is

$$
\begin{aligned}
J_{N}(\alpha) & =\left\langle\left[\partial_{\alpha} \log p_{\alpha}\left(x_{1}, x_{2}, \ldots, x_{N}\right)\right]^{2}\right\rangle \\
& =\left\langle V^{2}\left(x_{1}, x_{2}, \ldots, x_{N}\right)\right\rangle
\end{aligned}
$$

$$
\begin{align*}
& =\left\langle\left(\sum_{n=1}^{N} V\left(x_{n}\right)\right)^{2}\right\rangle \\
& =\sum_{n=1}^{N}\left\langle V^{2}\left(x_{n}\right)\right\rangle \\
& =N J(\alpha) \tag{12.39}
\end{align*}
$$

(remember that the individual scores are uncorrelated).
The Cramér-Rao inequality states that the mean square error of an unbiased estimator $f$ of $\alpha$ is lower bounded by the reciprocal of the Fisher information:

$$
\begin{equation*}
\sigma^{2}(f) \geq \frac{1}{J(\alpha)} \tag{12.40}
\end{equation*}
$$

To prove this, start with the Cauchy-Schwarz inequality

$$
\begin{align*}
\langle(V-\langle V\rangle)(f-\langle f\rangle)\rangle^{2} & \leq\left\langle(V-\langle V\rangle)^{2}\right\rangle\left\langle(f-\langle f\rangle)^{2}\right\rangle \\
\langle V f-\langle V\rangle f-V\langle f\rangle+\langle V\rangle\langle f\rangle\rangle^{2} & \leq\left\langle V^{2}-2 V\langle V\rangle+\langle V\rangle^{2}\right\rangle\left\langle(f-\langle f\rangle)^{2}\right\rangle \\
\langle V f\rangle^{2} & \leq\left\langle V^{2}\right\rangle\left\langle(f-\langle f\rangle)^{2}\right\rangle \\
\langle V f\rangle^{2} & \leq J(\alpha) \sigma^{2}(f) \tag{12.41}
\end{align*}
$$

(remember $\langle V\rangle=0$ ). The lefthand side equals one:

$$
\begin{align*}
\langle V f\rangle & =\int_{-\infty}^{\infty} \frac{\partial_{\alpha} p_{\alpha}(x)}{p_{\alpha}(x)} f(x) p_{\alpha}(x) d x \\
& =\int_{-\infty}^{\infty} \partial_{\alpha} p_{\alpha}(x) f(x) d x \\
& =\partial_{\alpha} \int_{-\infty}^{\infty} p_{\alpha}(x) f(x) d x \\
& =\partial_{\alpha}\langle f(x)\rangle \\
& =\partial_{\alpha} \alpha \\
& =1 \tag{12.42}
\end{align*}
$$

thus proving the Cramér-Rao inequality.
Just like the information theoretic channel capacity, this sets a lower limit on what is possible but does not provide any guidance in finding the minimum variance unbiased estimator. The inequality measures how much information the distribution provides about a parameter. Not surprisingly, the Fisher information can be related to the entropy of the distribution; this is done by de Bruijn's identity [Cover \& Thomas, 2006]. Roughly, the entropy measures the volume and the Fisher information measures the surface of the distribution.

One final caution about the Cramér-Rao bound. Not only may it not be reachable in practice, but it may be misleading because it is a bound on unbiased estimators. Unbiased does not necessarily mean better: it is possible for a biased estimator to dominate an unbiased one (as well as have other desirable characteristics). However, just like channel capacity, although it should not be taken too literally it does provide a good rough estimate of what is plausible and what is not.

### 12.7 SELECTED REFERENCES

[Press et al., 2007] Press, William H., Teukolsky, Saul A., Vetterling, William T., \& Flannery, Brian P. (2007). Numerical Recipes in C: The Art of Scientific Computing. 3nd edn. Cambridge: Cambridge University Press.
Numerical Recipes is particularly strong for function fitting.
[Cover \& Thomas, 2006] Cover, Thomas M., \& Thomas, Joy A. (2006). Elements of Information Theory. 2nd edn. New York: Wiley-Interscience.
Good coverage of the many connections between information theory and statistics.

### 12.8 PROBLEMS

(11.1) Generate 100 points $x$ uniformly distributed between 0 and 1 , and let $y=2+3 x+\zeta$, where $\zeta$ is a Gaussian random variable with a standard deviation of 0.5 . Use an SVD to fit $y=a+b x$ to this data set, finding $a$ and $b$. Evaluate the errors in $a$ and $b$
(a) With equation (12.34)
(b) By bootstrap sampling to generate 100 data sets
(c) From fitting an ensemble of 100 independent data sets
(11.2) Generate 100 points $x$ uniformly distributed between 0 and 1 , and let $y=\sin (2+$ $3 x)+\zeta$, where $\zeta$ is a Gaussian random variable with a standard deviation of 0.1 . Write a Levenberg-Marquardt routine to fit $y=\sin (a+b x)$ to this data set starting from $a=b=1$ (remembering that the second-derivative term can be dropped in the Hessian), and investigate the convergence for both fixed and adaptively adjusted $\lambda$ values.
(11.3) An alternative way to choose among models is to select the one that makes the weakest assumptions about the data; this is the purpose of maximum entropy methods. Assume that what is measured is a set of expectation values for functions $f_{i}$ of a random variable $x$,

$$
\begin{equation*}
\left\langle f_{i}(x)\right\rangle=\int_{-\infty}^{\infty} p(x) f_{i}(x) d x \tag{12.43}
\end{equation*}
$$

(a) Given these measurements, find the compatible normalized probability distribution $p(x)$ that maximizes the differential entropy

$$
\begin{equation*}
S=-\int_{-\infty}^{\infty} p(x) \log p(x) d x \tag{12.44}
\end{equation*}
$$

(b) What is the maximum entropy distribution if we know only the second moment

$$
\begin{equation*}
\sigma^{2}=\int_{-\infty}^{\infty} p(x) x^{2} d x \quad ? \tag{12.45}
\end{equation*}
$$

(11.4) Now consider the reverse situation. Let's say that we know that a data set $\left\{x_{n}\right\}_{n=1}^{N}$ was drawn from a Gaussian distribution with variance $\sigma^{2}$ and unknown mean $\mu$. Try to find an optimal estimator of the mean (one that is unbiased and has the smallest possible error in the estimate).

