

Chapter 2

Avian Bone Growth: A Case Study

In this chapter we use bone growth data to illustrate modeling methodology. Many important ideas and techniques are introduced in this case study.

2.1 Scientific problem

Growth is a fundamental process of life. Growth of multicellular organisms involves the multiplication and differentiation of cells. Cells can be arranged in different ways to create different body shapes, just like bricks can be arranged in different ways to create different buildings.

Size constrains an organism's life. You would not be surprised to see a spider walking up a vertical wall, but you would be surprised to see a human doing this. Spider locomotion is constrained by electrostatic forces, whereas human locomotion is constrained by gravitational forces. Size makes the difference.

Shape also plays a defining role in life. A giraffe feeds on leaves high in trees. By contrast, a hippo, shaped much differently, eats different things and would find it impossible to feed like the giraffe. As anatomists say, "structure follows function".

Through *development*, multicellular organisms grow from single cells into species-typical sizes and shapes. Organs within organisms do the same thing. As organs grow, they also change shape. And as they change shape, they change (or gain or lose) function.

Bones are organs of internal support in vertebrates. As a vertebrate grows, the size and shape of its bones change as well. Different bones change in different ways. By comparing the growth of various bones, we can learn about where the animal is funneling its energy and what is happening in its life.

Glaucous-winged gulls (*Larus glaucescens*) breed in large colonies in America's Pacific Northwest. For many years, we have studied these birds at

Protection Island National Wildlife Refuge, Strait of Juan de Fuca, Washington. We noticed that newly-hatched young behave differently than older juveniles, and that older juveniles behave differently than adults. So we asked a simple question: How does development of a particular behavior depend on the development of a particular bone?

2.1.1 Data

Gull chicks hatch in late June and early July on Protection Island. Many chicks die from predation, overheating, and dehydration or are killed by neighboring adults. If they are lucky enough to survive, they will fledge at about 44 days old and reach maturity at four years old.

During the hatching period, we banded 373 newly-hatched chicks. Over the next several weeks, we collected 80 banded chicks that had died. Because we knew when each chick had been banded, we could tell how old it was when it died. The dead chicks ranged from 0 to 42 days old. We also collected 13 dead adults on the colony.

We shipped the dead chicks and adults back to Michigan where we prepared their skeletons—a tedious, time-consuming, and smelly task! We decided to focus on three wing bones (humerus, ulna, and carpometacarpus) and three leg bones (femur, tibiotarsus, and tarsometatarsus), because these bones facilitate the easily-observed behaviors of flying and walking.

Using calipers, we measured the diaphyseal length and midshaft diameter of each of six bones from the right side of each skeleton. We graphed the lengths and diameters of each of the six bones against age and noted that the bones exhibited different growth patterns.

Gulls often lay three eggs per nest. The first two eggs, A and B, are generally bigger, are first to hatch, and produce larger, healthier chicks than the C egg. Our sample consisted of 52 (65% of the sample) C chicks, probably because they are least likely to survive. Thus, our data set is skewed in favor of smaller birds, something that should be kept in mind when interpreting the data.

It is also important to note that we used different bones from different birds for each age level—in other words, ours was a *cross-sectional* study. Had we followed the growth of individual bones in individual birds, we would have done a *longitudinal* study. Cross-sectional studies are limited by the fact that they tell us nothing about individual rates of growth, but only provide estimates of mean rates of growth for a population. Although it would be more informative, carrying out a longitudinal study on bone growth in wild gulls would be a logistical nightmare!

2.2 Translation into mathematics

How do gull bones grow? A bit of thought will convince you that this question cannot be translated into the language of mathematics, because it is

too vague. What does “grow” mean? We might define an object to be “growing” if and only if its size is changing over time. (This definition for growth includes shrinking as well as expanding.) But what does “size” mean? Are we interested in length, diameter, volume, or what? And which kind of bone are we talking about? Humerus? Ulna?

You can see that the very first step in the modeling process, translation into mathematics, typically requires careful thought. This step can be quite fruitful in and of itself, even if you never go any further in the modeling cycle, because the act of translating forces you to clarify concepts and sharpen questions. Translation into mathematics can help you ask whether your scientific question makes sense and can be expected to have a solution. This is important, because some apparently meaningful questions are actually nonsensical (exercise 1). Indeed, some of the burning scientific questions of history have simply “gone away” because they were discovered to be meaningless.

Let’s pose our problem precisely. How does the length of the humerus change in time over the life of a gull? That is, how does the length of the humerus change as a function of age? Consider a single “average” gull. Let

$$\begin{aligned}x &= \text{age in days} \\f(x) &= \text{length of humerus in cm}\end{aligned}$$

Mathematically, the question becomes: how does $f(x)$ depend on x ?

2.2.1 Simplifying assumptions

Whenever we relate the variables x and $f(x)$ through an equation, we are making implicit biological assumptions. As far as possible, these assumptions should be stated explicitly in the language of biology. Ideally, the assumptions should address two kinds of questions. First, what deterministic mechanisms are most important in driving the system? Second, what kind of process error is in the system, and what kind of measurement error is in the data?

Deterministic mechanisms

We pose two alternative hypotheses as the deterministic modeling assumption. They will give rise to two competing deterministic models which we will test against one another.

- A1a) The length of the humerus increases linearly with age until the chick reaches some critical age, at which point the bone abruptly ceases to grow.
- A1b) The length of the humerus grows first at an increasing rate and then at a decreasing rate, asymptotically leveling off (*saturating*) toward some maximal length. (This is called *sigmoidal growth* because the length vs. age curve is s-shaped.)

Stochastic mechanisms

We also make assumptions about process and measurement error in the system:

- A2) (Process error) Demographic stochasticity: The chicks represented typical sizes for their ages when they died. For a given age x , chick bone lengths are normally distributed with standard deviation σ_p . The standard deviation σ_p is independent of age. There is no environmental stochasticity.
- A3) (Measurement error) There is no measurement error in the age x . For a given bone, the measurement errors in the length are normally distributed with mean zero and standard deviation σ_m . The standard deviation σ_m does not depend on the bone length or the age of the chick.

2.2.2 The deterministic model

Assumption A1a says that bone length $f(x)$, when graphed against age x , is a line with positive slope from age 0 (hatching) to some critical age, at which point the graph of $f(x)$ becomes a horizontal line. Such a graph can be completely determined by three parameters. We will certainly want parameters for the critical age and the maximal bone length. The third parameter could be either the y -intercept (bone length at hatching) or the slope. Let

$$\begin{aligned} b &= \text{age at which bone growth stops} \\ K &= \text{maximal bone length} \\ a &= \text{slope} \end{aligned}$$

Then assumption A1a can be translated into mathematics as the model

$$f(x) = \begin{cases} a(x - b) + K & x < b \\ K & x \geq b \end{cases} \quad (2.1)$$

(exercise 5). Here x is age in days, $f(x)$ is the length of the humerus in cm at age x , and $a, b, K > 0$ are parameters.

The alternative assumption A1b is not specific enough to yield a unique equation, because there are several classes of functions that produce sigmoidal curves. In this example we will assume a modified *Holling Type III* growth model

$$f(x) = \frac{(K - b)x^2}{a^2 + x^2} + b, \quad (2.2)$$

where $a, b, K > 0$ are constants. It is important to note that in this model the parameters a and b have biological meanings different from those in model (2.1). The parameter K , however, has the same interpretation in both models (exercise 6).

The two deterministic models (2.1) and (2.2) are competing hypotheses that we are setting forth to explain the data in Table 1.

2.2.3 The stochastic model

We wish to explicitly model the stochasticity in the system, based on our assumptions A2 and A3 about the source of the noise. Let $F(x)$ be a random variable denoting the measurement of the humerus length in a chick of age x . We can write

$$F(x) = f(x) + \text{process error} + \text{measurement error}.$$

According to assumptions A2 and A3, the process errors and measurement errors are normally distributed with mean zero and standard deviations σ_p and σ_m . Thus, we can write the stochastic model as

$$F(x) = f(x) + \sigma_p \varepsilon_p + \sigma_m \varepsilon_m,$$

where ε_p and ε_m are *standard normal* random variables (mean zero and standard deviation one). Each realization of this stochastic model will first compute the deterministic part $f(x)$, and then add on two random *perturbations*, or error terms. The perturbations, which are realizations of the random variables $\sigma_p \varepsilon_p$ and $\sigma_m \varepsilon_m$, are drawn from the two respective normal distributions.

It can be shown that the sum of two random normal variables is a random normal variable whose variance is the sum of the two variances and whose mean is the sum of the two means. Therefore, we can combine the two sources of stochasticity, and write

$$F(x) = f(x) + \sigma \varepsilon, \tag{2.3}$$

where $\sigma^2 = \sigma_p^2 + \sigma_m^2$ and ε is a standard normal random variable. That is, we are assuming that any particular measurement of the length of a bone is the predicted length $f(x)$ plus a random normal perturbation. We can think of the stochastic model (2.3) as a *deterministic skeleton* $f(x)$ “clothed” with noise. The deterministic skeleton of a stochastic model is the part of the model that would remain if all the noise could be tuned to zero. In model (2.3), we can tune the noise to zero by letting $\sigma \rightarrow 0$; the skeleton that remains is $f(x)$.

It is fairly straightforward to show that if X and Y are random variables and c is a constant, then the following two properties hold:

$$\begin{aligned} E[X + Y] &= E[X] + E[Y] \\ E[cX] &= cE[X] \end{aligned}$$

(exercise 7). The expected value of $F(x)$ is therefore

$$\begin{aligned} E[F(x)] &= E[f(x) + \sigma \varepsilon] \\ &= E[f(x)] + \sigma E[\varepsilon] \\ &= f(x) + 0 \\ &= f(x). \end{aligned}$$

Let d_x be an actual data point (a measurement of the length of a bone in a particular chick of age x). Since we are thinking of model (2.3) as a surrogate

for the system, we can think of d_x as a realization of the random variable $F(x)$. The *residual error* of the deterministic prediction is the actual measurement minus the predicted value:

$$\text{residual} = d_x - f(x).$$

Note that the residual itself is a realization of the random variable $\sigma\varepsilon$ in the stochastic model (2.3). Thus, assumptions A2 and A3 boil down to the assumption that *the residuals are normally distributed with mean zero and standard deviation σ* .

2.3 Model parameterization

There are three parameters in our model: a , b , and K . We do not have any explicit information on the values of a and b . We do have some explicit data on the maximal bone length K , however. The lengths of 13 adult humeri are listed at the bottom of Table 1. We cannot include these data in estimating a and b , because the exact ages of the adult birds were unknown. However, we can use this information to estimate directly a value for K , and thus reduce the number of model parameters by one.

Our assumptions A2 and A3 imply that the adult bone length data are drawn from a normal distribution whose mean is the “true value” of K . Thus, we estimate the value of K by finding the mean of the adult data in Table 1:

$$\hat{K} = 11.96923077.$$

We will save quite a few decimals for the computations that follow, even though they are not all significant. Only the results of final computations should be rounded.

Our next goal is to use the juvenile data to estimate a and b .

2.3.1 Dividing the data set

Which data are to be used for parameter estimation? This often requires careful thought; the choice depends on the scientific question being asked, as well as how much and what kind of data exist. In general, if the data set is fairly large and robust, it is best to set aside some of the data for purposes of independent model validation.

We will divide the juvenile data set in Table 1 into two parts, one for model fitting (parameterization), and one for model evaluation (validation). We will call the two data sets the *estimation data set* and the *validation data set*. We want to do this randomly, but at the same time we would like to have representative data points from a variety of ages in each data set. We therefore use a technique called *stratified random sampling* in the following steps.

1. Chick ages range from $x = 0$ to $x = 42$. *Bin* the juvenile data into the age intervals $[0, 4]$, $[5, 9]$, $[10, 14]$, \dots , $[35, 39]$, $[40, 44]$ (Table 2). Our goal is to randomly select half of the data from each bin.

2. Number the data points in each bin (Table 2).
3. Consider the first bin. It has 18 data points numbered 1 through 18. Use a random number generator to flip an “18-sided coin” 50 times and write down the sequence of random numbers generated. Repeat this step for each bin. Such a procedure yielded the sequences of numbers in Table 3.
4. Consider the random sequence for the first bin: 17, 12, 16, 2, 15, 8, 16, 13, 12, 6, 3, 3, \dots . Put data point number 17 into data set 1, number 12 into data set 2, number 16 into data set 1, number 2 into data set 2, number 15 into data set 1, number 8 into data set 2, number 13 into data set 1 (skip over the 2nd occurrence of number 16), number 6 into data set 2, etc. Do this for each bin. The data are now divided into two sets, data set 1 and data set 2. Notice that if a bin had an odd number of data points, then data set 1 will always have one more data point than data set 2.
5. Flip a (2-sided) coin for bin 1. Heads means data set 1 for that bin goes into the estimation data; tails means it goes into the validation data. Do this for each bin. Such a procedure yielded the data sets in Table 4.

The data are now divided. The estimation and validation data are graphed together in Fig. 2.1. Make sure you can reconstruct Table 4 from the original data set, given the sequences of random numbers in Table 3 (exercise 8).

Our next goal is to use the estimation data set to estimate the parameters a and b . In what follows, we discuss two methods for doing this.

2.3.2 Maximum likelihood method (ML)

Our model assumes the residual errors $d_x - f(x)$ are normally distributed with mean zero and standard deviation σ . Thus, the *likelihood* that a given data point (and hence a given residual) occurs is

$$\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{d_x - f(x)}{\sigma}\right)^2}. \quad (2.4)$$

The likelihood that the entire data set occurs (as a realization of our stochastic model) is the product of the likelihoods over the whole data set:

$$L = \prod_{data} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{d_x - f(x)}{\sigma}\right)^2}, \quad (2.5)$$

where the symbol \prod denotes the product operator. If the data set has size q (that is, if there are q data points), equation (2.5) is equivalent to

$$L = \left(\frac{1}{\sigma\sqrt{2\pi}}\right)^q \exp\left(-\frac{1}{2} \sum_{data} \left(\frac{d_x - f(x)}{\sigma}\right)^2\right),$$

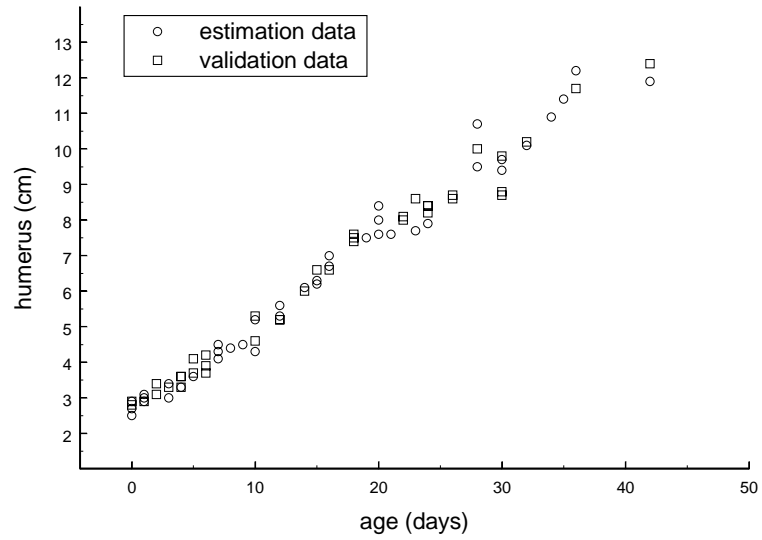


Figure 2.1: Estimation and validation data vs. age.

where $\exp y$ means e^y (exercise 9).

We want to maximize the likelihood L that the data set occurs as a realization of our stochastic model. The only thing that can be adjusted on the right hand side of equation (2.5) is the deterministic model prediction $f(x)$, which can be adjusted by changing the values of the parameters a and b in the model. That is, L is a function of a and b :

$$L(a, b) = \left(\frac{1}{\sigma\sqrt{2\pi}} \right)^q \exp \left(-\frac{1}{2\sigma^2} \sum_{data} (d_x - f(x))^2 \right). \quad (2.6)$$

This is called the *likelihood function*. We want to maximize $L(a, b)$ as a function of a and b . Actually, we want to find the *maximizer*, that is, the pair of parameters (\hat{a}, \hat{b}) that maximizes the function $L(a, b)$. These maximizer parameters are called the *maximum likelihood* (ML) parameters.

Note that a maximizer for L will also be a maximizer for $\ln L$, and vice versa (exercise 10). It is easy to show that the log-likelihood is

$$\ln L = -q \ln \sigma - \frac{q}{2} \ln(2\pi) - \frac{1}{2\sigma^2} \sum_{data} (d_x - f(x))^2 \quad (2.7)$$

(exercise 11).

2.3.3 Nonlinear least squares method (LS)

In our example, σ is assumed to be constant. Therefore the log-likelihood (2.7) is maximized when the *residual sum of squares* (RSS)

$$RSS(a, b) = \sum_{data} (d_x - f(x))^2 \quad (2.8)$$

is minimized as a function of the parameters (exercise 12). The minimizing parameters for RSS are called the *least squares* (LS) parameters. In this example, we will find the LS parameters.

Note: For univariate models (models with only one dependent variable), ML parameters and LS parameters are the same whenever the residuals are normally distributed with constant variance.

Once a and b are estimated, the variance of the residuals is estimated by

$$\hat{\sigma}^2 = \frac{RSS}{q}.$$

2.3.4 Downhill minimization routine: Nelder-Mead algorithm

Typically, one cannot minimize the function (2.8) analytically; it must be done numerically on a computer. Think of $RSS(a, b)$ as a surface suspended over

parameter space. In this case, parameter space is the horizontal plane spanned by the a -axis and b -axis. The RSS -axis rises vertically out of the plane. Each point on the plane corresponds to a parameter pair (a, b) , and the value of $RSS(a, b)$ is plotted above each point on the plane, generating a surface. We want to locate the *minimizer* point(s) (\hat{a}, \hat{b}) on the plane at which the surface attains a minimum value.

The Nelder-Mead algorithm begins with an initial “guess” (a_0, b_0) in parameter space, and then systematically checks around nearby in parameter space to see if there is a lower point on the surface. In this way, the routine “walks downhill” along the surface and converges on the minimum/minimizer. It is important to remember that a surface can have more than one local minimum. There is no fool-proof way of making sure you have found a global minimum. It pays to restart the Nelder-Mead routine from several different initial places to see if it always converges to the same minimum. Your teacher will provide you with a code for the Nelder-Mead routine.

2.3.5 Results of parameterization

The LS parameters for models (2.1) and (2.2), estimated from the estimation data set under the assumption that $\hat{K} = 11.96923077$, are, to four significant figures (exercise 14):

model	\hat{a}	\hat{b}	RSS	q	$\hat{\sigma}^2$	
A1a	0.2498	37.69	5.892	40	0.1473	(2.9)
A1b	17.49	2.789	14.10	40	0.3524	

The first question we should ask is: do these parameters make biological sense (exercise 15)? If one or more parameters persist in wandering off to zero or infinity during the parameterization process no matter what starting parameter values you choose, then the model is bad and should be discarded.

Note: A model that cannot be parameterized with feasible parameter values is not a good model.

The predictions of both models at their LS parameter values are shown with the estimation data in Fig. 2.2a,b. Which of the two alternative models best fits the data? We turn to this question in the next section.

2.4 Model selection

Which of the two alternative models, (2.1) or (2.2), best describes the data? The minimized RSS is smaller, and the estimated variance of errors $\hat{\sigma}^2$ is smaller, for model A1a. It would seem that model A1a therefore provides the best fit to the data. This conclusion is valid in this particular example, because both alternative models have the same number of parameters. If, however, the two models had contained different numbers of parameters, then the model with more parameters should be penalized. This is because an overparameterized

model can fit just about any data but has little explanatory power. When alternative models have different numbers of parameters, one can use *theoretic information criteria* to select the best model. These criteria will be explored in a later chapter.

In this example we can discard model A1b in favor of model A1a. Our deterministic model is therefore

$$\begin{aligned} f(x) &= \begin{cases} a(x-b) + K & x < b \\ K & x \geq b \end{cases} & (2.10) \\ a &= 0.2498 \\ b &= 37.69 \\ K &= 11.97, \end{aligned}$$

and our stochastic model is

$$\begin{aligned} F(x) &= \begin{cases} a(x-b) + K & x < b \\ K & x \geq b \end{cases} + \sigma\varepsilon & (2.11) \\ a &= 0.2498 \\ b &= 37.69 \\ K &= 11.97 \\ \sigma &= \sqrt{0.1473}. \end{aligned}$$

2.5 Model validation

To validate model (2.10), we must evaluate how well it explains the validation data, *without reparameterizing*. We will compute a number called the *goodness-of-fit* on both the parameterization data set and the validation data set. If the model fits the validation data as well as the parameterization data, then we will consider the model validated.

One measure of goodness-of-fit is the *R-squared* value, denoted R^2 . Consider a data set with members d_x . Let \bar{d}_x be the sample mean of the data set. Then the R-squared value is defined to be

$$R^2 = 1 - \frac{\sum_{data} (d_x - f(x))^2}{\sum_{data} (d_x - \bar{d}_x)^2}.$$

The residual sum of squares $\sum_{data} (d_x - f(x))^2$ measures the variability that is unexplained by the model, whereas $\sum_{data} (d_x - \bar{d}_x)^2$ measures the variability around the mean of the data. Thus, the quotient of these two numbers measures the fraction of variability not explained by the model, relative to use of the mean as a predictor. One minus the quotient measures the fraction of variability that *is* explained by the model (as opposed to using the mean as a predictor). The higher the R^2 value, the better the model fit. Although R^2 is usually between zero and one, it can also be negative (exercise 18).

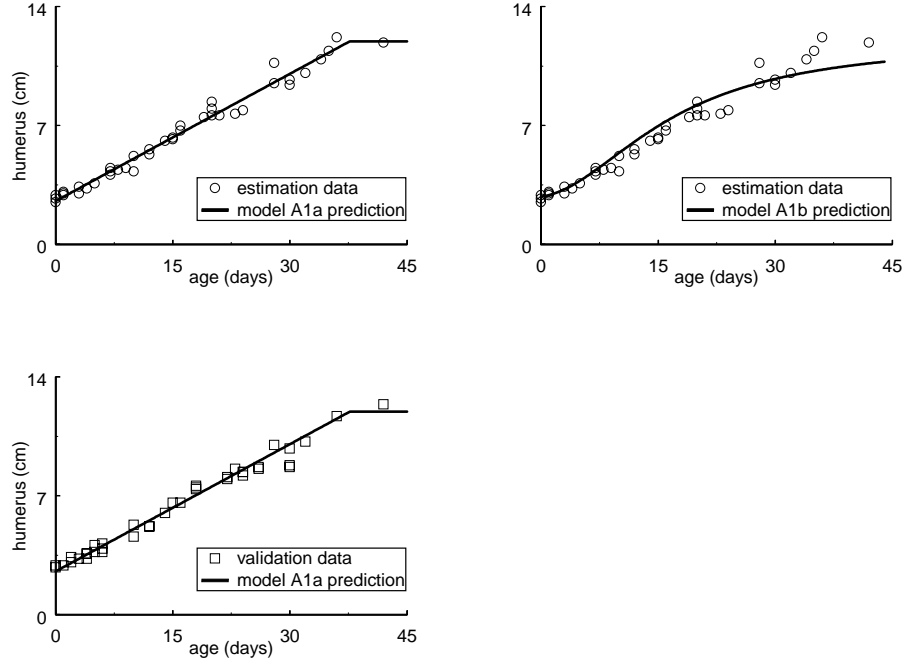


Figure 2.2: Model predictions. a) model A1a as fitted to the estimation data. b) model A1b as fitted to the estimation data. c) model A1a as compared to the validation data without refitting.

The R^2 values for model (2.10) on the two data sets are

$$R^2 \begin{array}{cc} \text{estimation data} & \text{validation data} \end{array} \quad (2.12)$$

$$R^2 \begin{array}{cc} 0.9821 & 0.9770 \end{array}$$

(exercise 19). Successful validation is supported because the R^2 values are about the same. The prediction of model (2.10) is shown in Fig. 2.2a,c with both the estimation and the validation data. Visually, we note that the model appears to fit the validation data about as well as it fits the estimation data set upon which it was parameterized.

2.6 Exercises

1. Some questions that seem to make sense are actually nonsensical. Explain why each of the following questions is nonsensical.
 - (a) What is the uptake rate (feeding rate) of a jackalope?
 - (b) What proportion of the mass of iron is due to flogiston?
 - (c) How many hours are there in a mile [1]?
 - (d) Is yellow square or round [1]?
 - (e) Let A be the set of all sets which do not have themselves as members; that is, $A = \{B \mid B \notin B\}$. Is $A \in A$?
2. What kinds of different physiological mechanisms could give rise to the situations in assumptions A1a and A1b?
3. Criticize assumption A2.
4. Criticize assumption A3. Design an experiment to test it.
5. Derive model (2.1) from assumption A1a. Explain why the model has to have at least three parameters.
6. The Holling Type III sigmoidal curve has the form

$$y = \frac{Mx^2}{a^2 + x^2}; \quad x \in [0, \infty)$$

Use calculus to do the following problems.

- (a) Show that $y(0) = 0$ and $\lim_{x \rightarrow \infty} y(x) = M$.
 - (b) Show that $y(x)$ is always increasing.
 - (c) Where is y concave up? Concave down? Find all inflection points.
 - (d) Graph y vs x .
 - (e) Find the value of x at which the curve reaches one half of its asymptotic limit. This is called the *half-saturation constant*.
 - (f) The Holling Type III curve cannot describe the length of the humerus, because the humerus has a positive length b upon hatching; that is, $f(0) = b$. We want to translate the Holling Type III curve vertically so that the initial value is b instead of zero. Also, we want the saturation level to be K . Derive model (2.2).
7. Prove that the expectation is linear. That is, prove that If X and Y are random variables and c is a constant, then

$$E[cX + Y] = cE[X] + E[Y]$$

8. Make sure you can reconstruct Table 4 from the original data set in Table 1, given the sequences of random numbers.
9. Show that equation (2.5) can be written as

$$L = \left(\frac{1}{\sigma\sqrt{2\pi}} \right)^q e^{-\frac{1}{2} \sum_{data} \left(\frac{d_x - f(x)}{\sigma} \right)^2}$$

10. Let $g(z)$ be any positive-valued function. Prove that a maximizer for $g(z)$ will also be a maximizer for $\ln g(z)$, and vice versa.
11. Derive the log-likelihood function (2.7) from the likelihood function (2.6).
12. Explain why the log-likelihood is maximized when the residual sum of squares is minimized.
13. Use the Nelder-Mead routine to solve the transcendental equation $x = e^{-x}$. (Hint: Let $g(a) = (a - e^{-a})^2$. Find the minimizer for g .) Attach your programs.
14. Reproduce the LS parameter estimates in (2.9). Attach your programs, input files, and output.
15. Are the LS parameter estimates (2.9) for the humerus model A1a biologically reasonable? Explain.
16. Write a program to produce a simulated data set of humerus lengths using the stochastic model (2.11). For each of the ages 0, 1, 2, ..., 45 generate 10 simulated data points. Present the output as a scatter plot of length against age. Attach your program.
17. What does $R^2 = 1$ mean? What does $R^2 = 0$ mean? Can R^2 be greater than one? Explain.
18. Consider the number of seals hauled out on a beach at hour t . Suppose a wildlife refuge biologist models the number of seals $N(t)$ on the beach at hour t with the equation

$$N(t) = 3t - 7$$

The biologist then collects 10 hours of data:

time t	seals	model prediction $N(t)$
5	10	
6	9	
7	11	
8	10	
9	10	
10	9	
11	8	
12	10	
13	11	
14	12	

- (a) Fill in the model predictions for each hour.
 - (b) Compute R^2 .
 - (c) Under what mathematical conditions is R^2 negative?
 - (d) What does a negative R^2 mean to a scientist?
19. Reproduce the R^2 values in (2.12). Attach your programs, input files, and output.
20. Model the growth of the ulna in the glaucous-winged gull. That is, repeat the comprehensive analysis in this chapter for the ulna length data in Table 1. Take your two alternative deterministic models to be

$$f(x) = \begin{cases} a(x - b) + K & x < b \\ K & x \geq b \end{cases}$$

and

$$f(x) = \frac{K}{1 + ae^{-bx}}$$

Use the same binning procedure as in Table 2 and the same sequences of random numbers as in Table 3, but work through all the details. Present your work in a complete, precise, and organized fashion.

21. We might suspect that the humerus in C chicks grows more slowly than in A and B chicks. Given the work already done in this chapter, we can now test this hypothesis. The idea is to fit the model to A and B chick data only, and then see if it also fits the C chick data.
- (a) Divide the humerus data in Table 1 into two sets: data set AB (consisting of all the A and B chick data) and data set C (consisting of all the C chick data).
 - (b) Estimate the LS parameters for model (2.1) on the AB data set. Record the parameter values and the RSS . Compute $\hat{\sigma}^2$ and R^2 .
 - (c) Without re-estimating the parameters, compute the RSS and R^2 on the independent data set C.
 - (d) Present your work in tables and graphs.
 - (e) What is your conclusion? Does the humerus grow differently in C chicks?

2.7 References

1. C. S. Lewis, quoted by A. N. Wilson, C. S. Lewis: A Biography (New York: Norton, 1990) p. 285.

