11

Logistic Regression

11.1 Modeling Conditional Probabilities

So far, we either looked at estimating the conditional expectations of continuous variables (as in regression), or at estimating distributions. There are many situations where however we are interested in input-output relationships, as in regression, but the output variable is discrete rather than continuous. In particular there are many situations where we have binary outcomes (it snows in Pittsburgh on a given day, or it doesn't; this squirrel carries plague, or it doesn't; this loan will be paid back, or it won't; this person will get heart disease in the next five years, or they won't). In addition to the binary outcome, we have some input variables, which may or may not be continuous. How could we model and analyze such data?

We could try to come up with a rule which guesses the binary output from the input variables. This is called **classification**, and is an important topic in statistics and machine learning. However, guessing "yes" or "no" is pretty crude — especially if there is no perfect rule. (Why should there be a perfect rule?) Something which takes noise into account, and doesn't just give a binary answer, will often be useful. In short, we want probabilities — which means we need to fit a stochastic model.

What would be nice, in fact, would be to have conditional distribution of the response Y, given the input variables, $\Pr(Y|X)$. This would tell us about how precise our predictions should be. If our model says that there's a 51% chance of snow and it doesn't snow, that's better than if it had said there was a 99% chance of snow (though even a 99% chance is not a sure thing). We will see, in Chapter 14, general approaches to estimating conditional probabilities non-parametrically, which can use the kernels for discrete variables from Chapter 4. While there are a lot of merits to this approach, it does involve coming up with a model for the joint distribution of outputs Y and inputs X, which can be quite time-consuming.

Let's pick one of the classes and call it "1" and the other "0". (It doesn't matter which is which.) Then Y becomes an **indicator variable**, and you can convince yourself that $\Pr(Y = 1) = \mathbb{E}[Y]$. Similarly, $\Pr(Y = 1|X = x) = \mathbb{E}[Y|X = x]$. (In a phrase, "conditional probability is the conditional expectation of the indicator".) This helps us because by this point we know all about estimating conditional expectations. The most straightforward thing for us to do at this point

257

would be to pick out our favorite smoother and estimate the regression function for the indicator variable; this will be an estimate of the conditional probability function.

There are two reasons not to just plunge ahead with that idea. One is that probabilities must be between 0 and 1, but our smoothers will not necessarily respect that, even if all the observed y_i they get are either 0 or 1. The other is that we might be better off making more use of the fact that we are trying to estimate probabilities, by more explicitly modeling the probability.

Assume that $\Pr(Y = 1|X = x) = p(x;\theta)$, for some function p parameterized by θ. parameterized function θ, *and* further assume that observations are independent of each other. The the (conditional) likelihood function is

$$\prod_{i=1}^{n} \Pr(Y = y_i|X = x_i) = \prod_{i=1}^{n} p(x_i;\theta)^{y_i}(1 - p(x_i;\theta))^{1-y_i} \qquad (11.1)$$

Recall that in a sequence of Bernoulli trials $y_1, \ldots y_n$, where there is a constant probability of success p, the likelihood is

$$\prod_{i=1}^{n} p^{y_i}(1 - p)^{1-y_i} \qquad (11.2)$$

As you learned in basic statistics, this likelihood is maximized when $p = \hat{p} = n^{-1}\sum_{i=1}^{n} y_i$. If each trial had its own success probability p_i, this likelihood becomes

$$\prod_{i=1}^{n} p_i^{y_i}(1 - p_i)^{1-y_i} \qquad (11.3)$$

Without some constraints, estimating the "inhomogeneous Bernoulli" model by maximum likelihood doesn't work; we'd get $\hat{p}_i = 1$ when $y_i = 1$, $\hat{p}_i = 0$ when $y_i = 0$, and learn nothing. If on the other hand we assume that the p_i aren't just arbitrary numbers but are linked together, if we *model* the probabilities, those constraints give non-trivial parameter estimates, and let us generalize. In the kind of model we are talking about, the constraint, $p_i = p(x_i;\theta)$, tells us that p_i must be the same whenever x_i is the same, and if p is a continuous function, then similar values of x_i must lead to similar values of p_i. Assuming p is known (up to parameters), the likelihood is a function of θ, and we can estimate θ by maximizing the likelihood. This chapter will be about this approach.

11.2 Logistic Regression

To sum up: we have a binary output variable Y, and we want to model the conditional probability $\Pr(Y = 1|X = x)$ as a function of x; any unknown parameters in the function are to be estimated by maximum likelihood. By now, it will not surprise you to learn that statisticians have approached this problem by asking themselves "how can we use linear regression to solve this?"

1. The most obvious idea is to let $p(x)$ be a linear function of x. Every increment of a component of x would add or subtract so much to the probability. This is called a "linear probability model". The conceptual problem here is that p must be between 0 and 1, and linear functions are unbounded. Moreover, in many situations we empirically see "diminishing returns" — changing p by the same amount requires a bigger change in x when p is already large (or small) than when p is close to $1/2$. Linear models can't do this.
2. The next most obvious idea is to let $\log p(x)$ be a linear function of x, so that changing an input variable *multiplies* the probability by a fixed amount. The problem is that logarithms of probabilities are unbounded in only one direction, and linear functions are not.
3. Finally, the easiest modification of $\log p$ which has an unbounded range is the **logistic transformation** (or **logit**) , $\log \frac{p}{1-p}$. We can make *this* a linear function of x without fear of nonsensical results. (Of course the results could still happen to be *wrong*, but they're not *guaranteed* to be wrong.)

This last alternative is **logistic regression**.

Formally, the logistic regression model is that

$$\log \frac{p(x)}{1 - p(x)} = \beta_0 + x \cdot \beta \tag{11.4}$$

Solving for p, this gives

$$p(x; \beta) = \frac{e^{\beta_0 + x \cdot \beta}}{1 + e^{\beta_0 + x \cdot \beta}} = \frac{1}{1 + e^{-(\beta_0 + x \cdot \beta)}} \tag{11.5}$$

Notice that the overall specification is a lot easier to grasp in terms of the transformed probability that in terms of the untransformed probability.[1]

To minimize the mis-classification rate, we should predict $Y = 1$ when $p \geq 0.5$ and $Y = 0$ when $p < 0.5$ (Exercise 11.1). This means guessing 1 whenever $\beta_0 + x \cdot \beta$ is non-negative, and 0 otherwise. So logistic regression gives us a **linear classifier**. The **decision boundary** separating the two predicted classes is the solution of $\beta_0 + x \cdot \beta = 0$, which is a point if x is one dimensional, a line if it is two dimensional, etc. One can show (exercise!) that the distance from the decision boundary is $\beta_0 / \|\beta\| + x \cdot \beta / \|\beta\|$. Logistic regression not only says where the boundary between the classes is, but also says (via Eq. 11.5) that the class probabilities depend on distance from the boundary, in a particular way, and that they go towards the extremes (0 and 1) more rapidly when $\|\beta\|$ is larger. It's these statements about probabilities which make logistic regression more than just a classifier. It makes stronger, more detailed predictions, and can be fit in a different way; but those strong predictions could be wrong.

Using logistic regression to predict class probabilities is a *modeling choice*, just like it's a modeling choice to predict quantitative variables with linear regression. In neither case is the appropriateness of the model guaranteed by the gods, nature,

---

[1] Unless you've taken thermodynamics or physical chemistry, in which case you recognize that this is the Boltzmann distribution for a system with two states, which differ in energy by $\beta_0 + x \cdot \beta$.
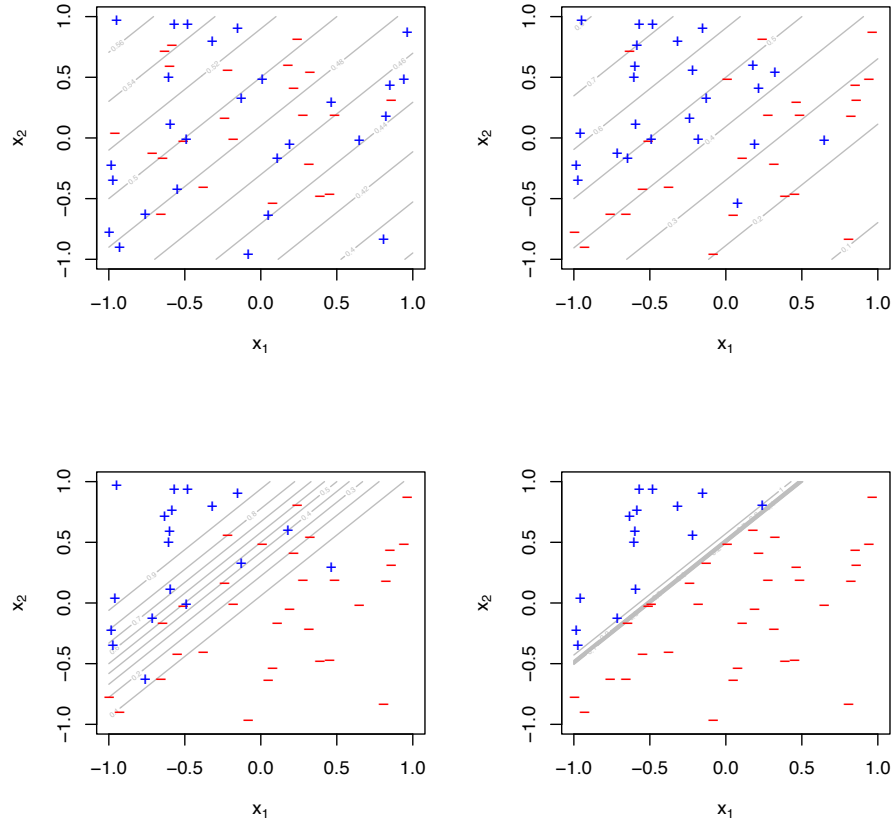
```
x <- matrix(runif(n = 50 * 2, min = -1, max = 1), ncol = 2)
par(mfrow = c(2, 2))
plot.logistic.sim(x, beta.0 = -0.1, beta = c(-0.2, 0.2))
y.1 <- plot.logistic.sim(x, beta.0 = -0.5, beta = c(-1, 1))
plot.logistic.sim(x, beta.0 = -2.5, beta = c(-5, 5))
plot.logistic.sim(x, beta.0 = -250, beta = c(-500, 500))
```

Figure 11.1 Effects of scaling logistic regression parameters. Values of x_1 and x_2 are the same in all plots ($\sim \text{Unif}(-1, 1)$ for both coordinates), but labels were generated randomly from logistic regressions with $\beta_0 = -0.1$, $\beta = (-0.2, 0.2)$ (top left); from $\beta_0 = -0.5$, $\beta = (-1, 1)$ (top right); from $\beta_0 = -2.5$, $\beta = (-5, 5)$ (bottom left); and from $\beta_0 = 2.5 \times 10^2$, $\beta = (-5 \times 10^2, 5 \times 10^2)$. Notice how as the parameters get increased in constant ratio to each other, we approach a deterministic relation between Y and x, with a linear boundary between the classes. (We save one set of the random binary responses for use later, as the imaginatively-named y.1.)

```
sim.logistic <- function(x, beta.0, beta, bind = FALSE) {
    require(faraway)
    linear.parts <- beta.0 + (x %*% beta)
    y <- rbinom(nrow(x), size = 1, prob = ilogit(linear.parts))
    if (bind) {
        return(cbind(x, y))
    }
    else {
        return(y)
    }
}

plot.logistic.sim <- function(x, beta.0, beta, n.grid = 50, labcex = 0.3, col = "grey",
    ...) {
    grid.seq <- seq(from = -1, to = 1, length.out = n.grid)
    plot.grid <- as.matrix(expand.grid(grid.seq, grid.seq))
    require(faraway)
    p <- matrix(ilogit(beta.0 + (plot.grid %*% beta)), nrow = n.grid)
    contour(x = grid.seq, y = grid.seq, z = p, xlab = expression(x[1]), ylab = expression(x[2]),
        main = "", labcex = labcex, col = col)
    y <- sim.logistic(x, beta.0, beta, bind = FALSE)
    points(x[, 1], x[, 2], pch = ifelse(y == 1, "+", "-"), col = ifelse(y == 1, "blue",
        "red"))
    invisible(y)
}
```

CODE EXAMPLE 26: *Code to simulate binary responses from a logistic regression model, and to plot a 2D logistic regression's probability contours and simulated binary values. (How would you modify this to take the responses from a data frame?*

mathematical necessity, etc. We begin by positing the model, to get something to work with, and we end (if we know what we're doing) by checking whether it really does match the data, or whether it has systematic flaws.

Logistic regression is one of the most commonly used tools for applied statistics and discrete data analysis. There are basically four reasons for this.

1. Tradition.
2. In addition to the heuristic approach above, the quantity $\log p/(1-p)$ plays an important role in the analysis of contingency tables (the "log odds"). Classification is a bit like having a contingency table with two columns (classes) and infinitely many rows (values of x). With a finite contingency table, we can estimate the log-odds for each row empirically, by just taking counts in the table. With infinitely many rows, we need some sort of interpolation scheme; logistic regression is linear interpolation for the log-odds.
3. It's closely related to "exponential family" distributions, where the probability of some vector v is proportional to $\exp\left\{\beta_0 + \sum_{j=1}^{m} f_j(v)\beta_j\right\}$. If one of the components of v is binary, and the functions f_j are all the identity function, then we get a logistic regression. Exponential families arise in many contexts in statistical theory (and in physics!), so there are lots of problems which can be turned into logistic regression.

4. It often works surprisingly well as a classifier. But, many simple techniques often work surprisingly well as classifiers, and this doesn't really testify to logistic regression getting the probabilities right.

11.2.1 Likelihood Function for Logistic Regression

[[TODO: Standardize notation here for likelihood function compared to theory appendix]]

Because logistic regression predicts probabilities, rather than just classes, we can fit it using likelihood. For each training data-point, we have a vector of features, x_i, and an observed class, y_i. The probability of that class was either p, if $y_i = 1$, or $1 - p$, if $y_i = 0$. The likelihood is then

$$L(\beta_0, \beta) = \prod_{i=1}^{n} p(x_i)^{y_i} (1 - p(x_i))^{1-y_i} \tag{11.6}$$

(I could substitute in the actual equation for p, but things will be clearer in a moment if I don't.) The log-likelihood turns products into sums:

$$\ell(\beta_0, \beta) = \sum_{i=1}^{n} y_i \log p(x_i) + (1 - y_i) \log (1 - p(x_i)) \tag{11.7}$$

$$= \sum_{i=1}^{n} \log (1 - p(x_i)) + \sum_{i=1}^{n} y_i \log \frac{p(x_i)}{1 - p(x_i)} \tag{11.8}$$

$$= \sum_{i=1}^{n} \log (1 - p(x_i)) + \sum_{i=1}^{n} y_i(\beta_0 + x_i \cdot \beta) \tag{11.9}$$

$$= \sum_{i=1}^{n} -\log \left(1 + e^{\beta_0 + x_i \cdot \beta}\right) + \sum_{i=1}^{n} y_i(\beta_0 + x_i \cdot \beta) \tag{11.10}$$

where in the next-to-last step we finally use equation 11.4.

Typically, to find the maximum likelihood estimates we'd differentiate the log likelihood with respect to the parameters, set the derivatives equal to zero, and solve. To start that, take the derivative with respect to one component of β, say β_j.

$$\frac{\partial \ell}{\partial \beta_j} = -\sum_{i=1}^{n} \frac{1}{1 + e^{\beta_0 + x_i \cdot \beta}} e^{\beta_0 + x_i \cdot \beta} x_{ij} + \sum_{i=1}^{n} y_i x_{ij} \tag{11.11}$$

$$= \sum_{i=1}^{n} (y_i - p(x_i; \beta_0, \beta)) \, x_{ij} \tag{11.12}$$

We are not going to be able to set this to zero and solve exactly. (That's a transcendental equation, and there is no closed-form solution.) We can however approximately solve it numerically.

11.3 Numerical Optimization of the Likelihood

While our likelihood isn't nice enough that we have an explicit expression for the maximum (the way we do in OLS or WLS), it is a pretty well-behaved function, and one which is amenable to lots of the usual numerical methods for optimization. In particular, like most log-likelihood functions, it's suitable for an application of Newton's method. Briefly (see Appendix D.2 for details), Newton's method starts with an initial guess about the optimal parameters, and then calculates the gradient of the log-likelihood with respect to those parameters. It then adds an amount proportional to the gradient to the parameters, moving up the surface of the log-likelihood function. The size of the step in the gradient direction is dictated by the second derivatives — it takes bigger steps when the second derivatives are small (so the gradient is a good guide to what the function looks like), and small steps when the curvature is large.

11.3.1 Iteratively Re-Weighted Least Squares

Remarkably enough, in the case of logistic regression, each step of Newton's method ends up looking like a good, old-fashioned linear regression problem.

Fundamentally, this is because logistic regression is a linear model for a transformation of the probability. Let's call this transformation g:

$$g(p) \equiv \log \frac{p}{1-p} \tag{11.13}$$

So the model is

$$g(p) = \beta_0 + x \cdot \beta \tag{11.14}$$

and $Y|X = x \sim \text{Binom}(1, g^{-1}(\beta_0 + x \cdot \beta))$. It seems that what we should want to do is take $g(y)$ and regress it linearly on x. Of course, the variance of Y, according to the model, is going to change depending on x — it will be $(g^{-1}(\beta_0 + x \cdot \beta))(1 - g^{-1}(\beta_0 + x \cdot \beta))$ — so we really ought to do a weighted linear regression, with weights inversely proportional to that variance. (We learned about weighted linear regression in Chapter 10.) Since writing $g^{-1}(\beta_0 + x \cdot \beta)$ is getting annoying, let's abbreviate it by $p(x)$ or just p, and let's abbreviate that variance as $V(p)$.

The problem is that y is either 0 or 1, so $g(y)$ is either $-\infty$ or $+\infty$. We will evade this by using first-order Taylor expansion (App. B).

$$g(y) \approx g(p) + (y - p)g'(p) \equiv z \tag{11.15}$$

The right hand side, z will be our *effective* response variable, which we will regress on x. To see why this should give us the right coefficients, substitute for $g(p)$ in the definition of z,

$$z = \beta_0 + x \cdot \beta + (y - p)g'(p) \tag{11.16}$$

and notice that, if we've got the coefficients right, $\mathbb{E}[Y|X = x] = p$, so $(y - p)$ should be mean-zero noise. In other words, when we have the right coefficients, z *is* a linear function of x plus mean-zero noise. (This is our excuse for throwing

away the rest of the Taylor expansion, even though we know the discarded terms are infinitely large!) That noise doesn't have constant variance, but we can work it out,

$$\mathbb{V}[Z|X = x] = \mathbb{V}[(Y - p)g'(p)|X = x] = (g'(p))^2 V(p) , \qquad (11.17)$$

and so use that variance in weighted least squares to recover β.

Notice that z and the weights both involve the parameters of our logistic regression, through $p(x)$. So having done this once, we should really use the new parameters to update z and the weights, and do it again. Eventually, we come to a fixed point, where the parameter estimates no longer change. This loop — start with a guess about the parameters, use it to calculate the z_i and their weights, regress on the x_i to get new parameters, and repeat — is known as **iterative reweighted least squares** (IRLS or IRWLS), **iterative weighted least squares** (IWLS), etc.

The treatment above is rather heuristic[2], but it turns out to be equivalent to using Newton's method, only with the expected second derivative of the log likelihood, instead of its actual value. This takes a reasonable amount of algebra to show, so we'll skip it (but see Exercise 11.3)[3]. Since, with a large number of observations, the observed second derivative should be close to the expected second derivative, this is only a small approximation.

Perfect Classification

One caution about using maximum likelihood to fit logistic regression is that it can seem to work badly when the training data *can* be linearly separated. The reason is that, to make the likelihood large, $p(x_i)$ should be large when $y_i = 1$, and p should be small when $y_i = 0$. If β_0, β_0 is a set of parameters which perfectly classifies the training data, then $c\beta_0, c\beta$ is too, for any $c > 1$, but in a logistic regression the second set of parameters will have more extreme probabilities, and so a higher likelihood. For linearly separable data, then, there is no parameter vector which *maximizes* the likelihood, since ℓ can always be increased by making the vector larger but keeping it pointed in the same direction.

You should, of course, be so lucky as to have this problem.

---

[2] That is, mathematically incorrect.

[3] The two key points are as follows. First, the gradient of the log-likelihood turns out to be the sum of the $z_i x_i$. (Cf. Eq. 11.12.) Second, take a single Bernoulli observation with success probability p. The log-likelihood is $Y \log p + (1 - Y) \log 1 - p$. The first derivative with respect to p is
$Y/p - (1 - Y)/(1 - p)$, and the second derivative is $-Y/p^2 - (1 - Y)/(1 - p)^2$. Taking expectations of the second derivative gives $-1/p - 1/(1 - p) = -1/p(1 - p)$. In other words, $V(p) = -1/\mathbb{E}[\ell'']$. Using weights inversely proportional to the variance thus turns out to be equivalent to dividing by the expected second derivative. But gradient divided by second derivative is the increment we use in Newton's method, QED.

11.4 Generalized Linear and Generalized Additive Models

Logistic regression is part of a broader family of **generalized linear models** (GLMs), where the conditional distribution of the response falls in some parametric family, and the parameters are set by the linear predictor. Ordinary, least-squares regression is the case where response is Gaussian, with mean equal to the linear predictor, and constant variance. Logistic regression is the case where the response is binomial, with n equal to the number of data-points with the given x (usually but not always 1), and p is given by Equation 11.5. Changing the relationship between the parameters and the linear predictor is called changing the **link function**. For computational reasons, the link function is actually the function you apply to the mean response to get back the linear predictor, rather than the other way around — (11.4) rather than (11.5). There are thus other forms of binomial regression besides logistic regression.[4] There is also Poisson regression (appropriate when the data are counts without any upper limit), gamma regression, etc.; we will say more about these in Chapter 12.

In R, any standard GLM can be fit using the (base) `glm` function, whose syntax is very similar to that of `lm`. The major wrinkle is that, of course, you need to specify the family of probability distributions to use, by the `family` option — `family=binomial` defaults to logistic regression. (See `help(glm)` for the gory details on how to do, say, probit regression.) All of these are fit by the same sort of numerical likelihood maximization.

11.4.1 Generalized Additive Models

A natural step beyond generalized linear models is **generalized additive models** (GAMs), where instead of making the transformed mean response a *linear* function of the inputs, we make it an *additive* function of the inputs. This means combining a function for fitting additive models with likelihood maximization. This is actually done in R with the same `gam` function we used for additive models (hence the name). We will look at how this works in some detail in Chapter 12. For now, the basic idea is that the iteratively re-weighted least squares procedure of §11.3.1 doesn't really require the model for the log odds to be linear. We get a GAM when we fit an additive model to the z_i; we could even fit an arbitrary non-parametric model, like a kernel regression, though that's not often done.

GAMs can be used to check GLMs in much the same way that smoothers can be used to check parametric regressions: fit a GAM and a GLM to the same data, then simulate from the GLM, and re-fit both models to the simulated data. Repeated many times, this gives a distribution for how much better the GAM will seem to fit than the GLM does, *even when the GLM is true*. You can then read a p-value off of this distribution. This is illustrated in §11.6 below.

----

[4] My experience is that these tend to give similar error rates as classifiers, but have rather different guesses about the underlying probabilities.

11.5 Model Checking

The validity of the logistic regression model is no more a fact of mathematics or nature than is the validity of the linear regression model. Both are sometimes convenient assumptions, but neither is guaranteed to be correct, nor even some sort of generally-correct default. In either case, if we want to use the model, the proper scientific (and statistical) procedure is to *check* the validity of the modeling assumptions.

11.5.1 Residuals

In your linear models course, you learned a lot of checks based on the residuals of the model (see Chapter 2). Many of these ideas translates to logistic regression, but we need to re-define residuals. Sometimes people work with the "response" residuals,

$$y_i - p(x_i) \qquad (11.18)$$

which should have mean zero (why?), but are heteroskedastic even when the model is true (why?). Others work with standardized or **Pearson** residuals,

$$\frac{y_i - p(x_i)}{\sqrt{V(p(x_i))}} \qquad (11.19)$$

and there are yet other notions of residuals for logistic models. Still, both the response and the Pearson residuals should be unpredictable from the covariates, and the latter should have constant variance.

11.5.2 Non-parametric Alternatives

Chapter 9 discussed how non-parametric regression models can be used to check whether parametric regressions are well-specified. The same ideas apply to logistic regressions, with the minor modification that in place of the difference in MSEs, one should use the difference in log-likelihoods, or (what comes to the same thing, up to a factor of 2) the difference in deviances. The use of generalized additive models (§11.4.1) as the alternative model class is illustrated in §11.6 below.

11.5.3 Calibration

Because logistic regression predicts actual *probabilities*, we can check its predictions in a more stringent way than an ordinary regression, which just tells us the mean value of Y, but is otherwise silent about its distribution. If we've got a model which tells us that the probability of rain on a certain class of days is 50%, it had better rain on half of those days, or there model is just *wrong* about the probability of rain. More generally, we'll say that the model is **calibrated** (or **well-calibrated**) when

$$\Pr\left(Y = 1 | \hat{p}(X) = p\right) = p \tag{11.20}$$

That is, the actual probabilities should match the predicted probabilities. If we have a large sample, by the law of large numbers, observed relative frequencies will converge on true probabilities. Thus, the observed relative frequencies should be close to the predicted probabilities, or else the model is making systematic mistakes.

In practice, each case often has its own unique predicted probability p, so we can't really accumulate many cases with the same p and check the relative frequency among those cases. When that happens, one option is to look at all the cases where the predicted probability is in some small range $[p, p + \epsilon)$; the observed relative frequency had then better be in that range too. §11.7 below illustrates some of the relevant calculations.

A second option is to use what is called a **proper scoring rule**, which is a function of the outcome variables and the predicted probabilities that attains its minimum when, and only when, the predicted are calibrated. For binary outcomes, one proper scoring rule (historically the oldest) is the **Brier score**,

$$n^{-1} \sum_{i=1}^{n} (y_i - p_i)^2 \tag{11.21}$$

Another however is simply the (normalized) negative log-likelihood,

$$-n^{-1} \sum_{i=1}^{n} y_i \log p_i + (1 - y_i) \log (1 - p_i) \tag{11.22}$$

Of course, proper scoring rules are better evaluated out-of-sample, or, failing that, through cross-validation, than in-sample. Even an in-sample evaluation is better than nothing, however, which is too often what happens.

11.6 A Toy Example

Here's a worked R example, using the data from the upper right panel of Figure 11.1. The 50×2 matrix `x` holds the input variables (the coordinates are independently and uniformly distributed on $[-1, 1]$), and `y.1` the corresponding class labels, themselves generated from a logistic regression with $\beta_0 = -0.5$, $\beta = (-1, 1)$.

```
df <- data.frame(y = y.1, x1 = x[, 1], x2 = x[, 2])
logr <- glm(y ~ x1 + x2, data = df, family = "binomial")
```

The **deviance** of a model fitted by maximum likelihood is twice the difference between its log likelihood and the maximum log likelihood for a **saturated** model, i.e., a model with one parameter per observation. Hopefully, the saturated model can give a perfect fit.[5] Here the saturated model would assign probability 1 to the observed outcomes[6], and the logarithm of 1 is zero, so $D = 2\ell(\widehat{\beta_0}, \widehat{\beta})$. The null deviance is what's achievable by using just a constant bias β_0 and setting the rest of β to 0. The fitted model definitely improves on that.[7]

If we're interested in inferential statistics on the estimated model, we can see those with `summary`, as with `lm`:

```
summary(logr, digits = 2, signif.stars = FALSE)
##
## Call:
## glm(formula = y ~ x1 + x2, family = "binomial", data = df)
##
## Deviance Residuals:
##      Min        1Q    Median        3Q       Max
## -2.34521  -0.82798   0.01499   0.83880   2.07197
##
## Coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept)  -0.5091     0.3793  -1.342  0.17957
## x1           -2.2365     0.7293  -3.067  0.00216 **
## x2            2.4894     0.8556   2.909  0.00362 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##     Null deviance: 69.315  on 49  degrees of freedom
## Residual deviance: 48.929  on 47  degrees of freedom
## AIC: 54.929
##
## Number of Fisher Scoring iterations: 5
```

---

[5] The factor of two is so that the deviance will have a χ^2 distribution. Specifically, if the model with p parameters is right, as $n \to \infty$ the deviance will approach a χ^2 distribution with $n - p$ degrees of freedom.

[6] This is not possible when there are multiple observations with the same input features, but different classes.

[7] AIC is of course the Akaike information criterion, $-2\ell + 2p$, with p being the number of parameters (here, $p = 3$). (Some people divide this through by n.) See §D.5.5.5 for more on AIC, and why I mostly ignore it in this book.

```
simulate.from.logr <- function(df, mdl) {
    probs <- predict(mdl, newdata = df, type = "response")
    df$y <- rbinom(n = nrow(df), size = 1, prob = probs)
    return(df)
}
```

CODE EXAMPLE 27: *Code for simulating from an* estimated *logistic regression model. By default* `(type="link")`, `predict` *for logistic regressions returns predictions for the log odds; changing the* `type` *to* `"response"` *returns a probability.*

The fitted values of the logistic regression are the class probabilities; this next line gives us the (in-sample) mis-classification rate.

```
mean(ifelse(fitted(logr) < 0.5, 0, 1) != df$y)
## [1] 0.16
```

An error rate of 16% may sound bad, but notice from the contour lines in Figure 11.1 that lots of the probabilities are near 0.5, meaning that the classes are just genuinely hard to predict.

To see how well the logistic regression assumption holds up, let's compare this to a GAM. We'll use the same package for estimating the GAM, `mgcv`, that we used to fit the additive models in Chapter 8.

```
library(mgcv)
(gam.1 <- gam(y ~ s(x1) + s(x2), data = df, family = "binomial"))
##
## Family: binomial
## Link function: logit
##
## Formula:
## y ~ s(x1) + s(x2)
##
## Estimated degrees of freedom:
## 1.00 2.09   total = 4.09
##
## UBRE score: 0.03774858
```

This fits a GAM to the same data, using spline smoothing of both input variables. (Figure 11.2 shows the partial response functions.) The (in-sample) deviance is

```
signif(gam.1$deviance, 3)
## [1] 43.7
```

which is lower than the logistic regression, so the GAM gives the data higher likelihood. We expect this; the question is whether the difference is significant, or within the range of what we should expect when logistic regression is valid. To test this, we need to simulate from the logistic regression model.

Now we simulate from our fitted model, and re-fit both the logistic regression and the GAM.

```
delta.deviance.sim <- function(df, mdl) {
    sim.df <- simulate.from.logr(df, mdl)
```
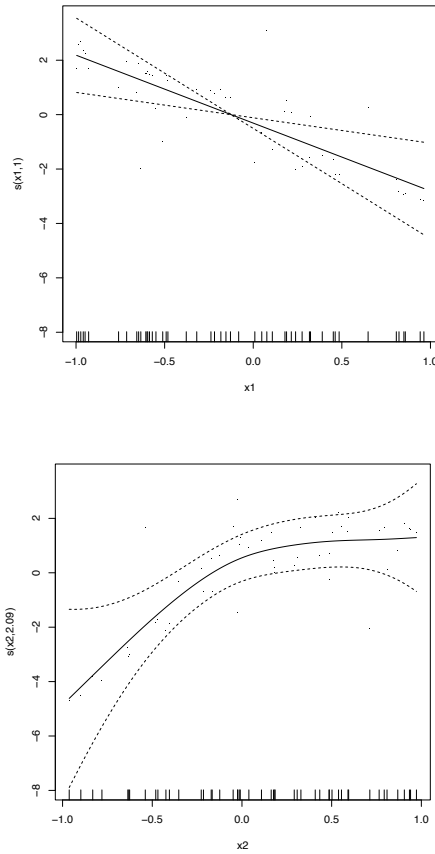
```
plot(gam.1, residuals = TRUE, pages = 0)
```

Figure 11.2 Partial response functions estimated when we fit a GAM to the data simulated from a logistic regression. Notice that the vertical axes are on the logit scale.

```
GLM.dev <- glm(y ~ x1 + x2, data = sim.df, family = "binomial")$deviance
GAM.dev <- gam(y ~ s(x1) + s(x2), data = sim.df, family = "binomial")$deviance
return(GLM.dev - GAM.dev)
}
```

Notice that in this simulation we are not generating new $\vec{X}$ values. The logistic regression and the GAM are both models for the response *conditional* on the inputs, and are agnostic about how the inputs are distributed, or even whether it's meaningful to talk about their distribution.

Finally, we repeat the simulation a bunch of times, and see where the observed difference in deviances falls in the sampling distribution.

```
(delta.dev.observed <- logr$deviance - gam.1$deviance)
## [1] 5.212441
delta.dev <- replicate(100, delta.deviance.sim(df, logr))
mean(delta.dev.observed <= delta.dev)
## [1] 0.41
```

In other words, the amount by which a GAM fits the data better than logistic regression is pretty near the middle of the null distribution. Since the example data really *did* come from a logistic regression, this is a relief.
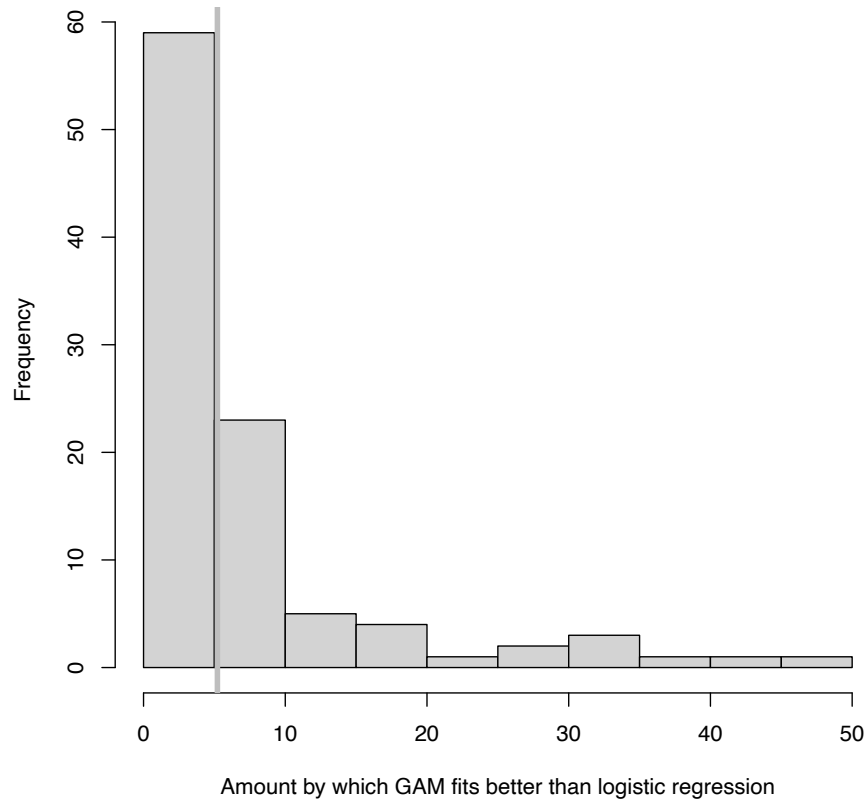
```
hist(delta.dev, main = "", xlab = "Amount by which GAM fits better than logistic regression")
abline(v = delta.dev.observed, col = "grey", lwd = 4)
```

Figure 11.3 Sampling distribution for the difference in deviance between a GAM and a logistic regression, on data generated from a logistic regression. The observed difference in deviances is shown by the grey vertical line.

```
snoqualmie <- scan("http://www.stat.washington.edu/peter/book.data/set1", skip = 1)
snoq <- data.frame(tomorrow = c(tail(snoqualmie, -1), NA), today = snoqualmie)
years <- 1948:1983
days.per.year <- rep(c(366, 365, 365, 365), length.out = length(years))
snoq$year <- rep(years, times = days.per.year)
snoq$day <- rep(c(1:366, 1:365, 1:365, 1:365), times = length(years)/4)
snoq <- snoq[-nrow(snoq), ]
```

CODE EXAMPLE 28: *Read in and re-shape the Snoqualmie data set. Prof. Guttorp, who has kindly provided the data, formatted it so that each year was a different row, which is rather inconvenient for R.*

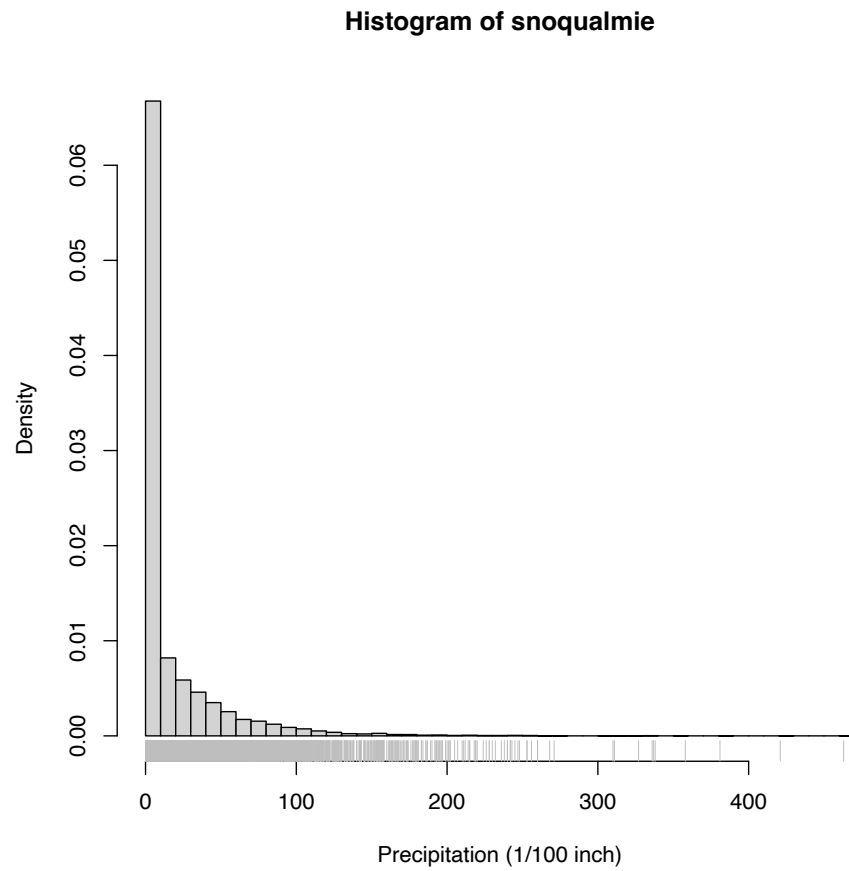### 11.7 Weather Forecasting in Snoqualmie Falls

For our worked data example, we are going to build a simple weather forecaster. Our data consist of daily records, from the start of 1948 to the end of 1983, of precipitation at Snoqualmie Falls, Washington (Figure 11.4)[8]. Each row of the data file is a different year; each column records, for that day of the year, the day's precipitation (rain or snow), in units of $\frac{1}{100}$ inch. Because of leap-days, there are 366 columns, with the last column having an NA value for three out of four years.

What we want to do is predict tomorrow's weather from today's. This would be of interest if we lived in Snoqualmie Falls, or if we operated one of the local hydroelectric power plants, or the tourist attraction of the Falls themselves. Examining the distribution of the data (Figures 11.5 and 11.6) shows that there is a big spike in the distribution at zero precipitation, and that days of no precipitation can follow days of any amount of precipitation but seem to be less common after heavy precipitation.

---

[8] I learned of this data set from Guttorp (1995); the data file is available from http://www.stat.washington.edu/peter/stoch.mod.data.html. See Code Example 28 for the commands used to read it in, and to reshape it into a form more convenient for R.

Figure 11.4 Snoqualmie Falls, Washington, on a low-precipitation day.
Photo by Jeannine Hall Gailey, from http://myblog.webbish6.com/2011/
07/17-years-and-hoping-for-another-17.html [[TODO: Get
permission for photo use!]]

Histogram of snoqualmie


```
hist(snoqualmie, n = 50, probability = TRUE, xlab = "Precipitation (1/100 inch)")
rug(snoqualmie, col = "grey")
```

Figure 11.5 Histogram of the amount of daily precipitation at Snoqualmie
Falls

```
plot(tomorrow ~ today, data = snoq, xlab = "Precipitation today (1/100 inch)", ylab = "Precipitation
    cex = 0.1)
rug(snoq$today, side = 1, col = "grey")
rug(snoq$tomorrow, side = 2, col = "grey")
```

Figure 11.6 Scatterplot showing relationship between amount of
precipitation on successive days. Notice that days of no precipitation can
follow days of any amount of precipitation, but seem to be more common
when there is little or no precipitation to start with.

These facts suggest that "no precipitation" is a special sort of event which would be worth predicting in its own right (as opposed to just being when the precipitation happens to be zero), so we will attempt to do so with logistic regression. Specifically, the input variable X_i will be the amount of precipitation on the i^{th} day, and the response Y_i will be the indicator variable for whether there was any precipitation on day $i + 1$ — that is, $Y_i = 1$ if $X_{i+1} > 0$, an $Y_i = 0$ if $X_{i+1} = 0$. We expect from Figure 11.6, as well as common experience, that the coefficient on X should be positive.[9]

The estimation is straightforward:

```
snoq.logistic <- glm((tomorrow > 0) ~ today, data = snoq, family = binomial)
```

To see what came from the fitting, run `summary`:

```
print(summary(snoq.logistic), digits = 3, signif.stars = FALSE)
##
## Call:
## glm(formula = (tomorrow > 0) ~ today, family = binomial, data = snoq)
##
## Deviance Residuals:
##    Min     1Q  Median     3Q     Max
## -4.525  -0.999   0.167   1.170   1.367
##
## Coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) -0.43520    0.02163   -20.1   <2e-16
## today        0.04523    0.00131    34.6   <2e-16
##
## (Dispersion parameter for binomial family taken to be 1)
##
##     Null deviance: 18191  on 13147  degrees of freedom
## Residual deviance: 15896  on 13146  degrees of freedom
## AIC: 15900
##
## Number of Fisher Scoring iterations: 5
```

The coefficient on the amount of precipitation today is indeed positive, and (if we can trust R's assumptions) highly significant. There is also an intercept term, which is slightly positive. We can see what the intercept term means by considering what happens on days of no precipitation. The linear predictor is then just the intercept, -0.435, and the predicted probability of precipitation is 0.393. That is, even when there is no precipitation today, it's almost as likely as not that there will be some precipitation tomorrow.[10]

We can get a more global view of what the model is doing by plotting the data and the predictions (Figure 11.7). This shows a steady increase in the probability of precipitation tomorrow as the precipitation today increases, though with the leveling off characteristic of logistic regression. The (approximate) 95% confidence limits for the predicted probability are (on close inspection) asymmetric.

---

[9] This does not attempt to model *how much* precipitation there will be tomorrow, if there is any. We could make that a separate model, if we can get this part right.

[10] For western Washington State, this is plausible — but see below.

```
plot((tomorrow > 0) ~ today, data = snoq, xlab = "Precipitation today (1/100 inch)",
    ylab = "Positive precipitation tomorrow?")
rug(snoq$today, side = 1, col = "grey")
data.plot <- data.frame(today = (0:500))
pred.bands <- function(mdl, data, col = "black", mult = 1.96) {
    preds <- predict(mdl, newdata = data, se.fit = TRUE)
    lines(data[, 1], ilogit(preds$fit), col = col)
    lines(data[, 1], ilogit(preds$fit + mult * preds$se.fit), col = col, lty = "dashed")
    lines(data[, 1], ilogit(preds$fit - mult * preds$se.fit), col = col, lty = "dashed")
}
pred.bands(snoq.logistic, data.plot)
```

Figure 11.7 Data (dots), plus predicted probabilities (solid line) and
approximate 95% confidence intervals from the logistic regression model
(dashed lines). Note that calculating standard errors for predictions on the
logit scale, and then transforming, is better practice than getting standard
errors directly on the probability scale.

How well does this work? We can get a first sense of this by comparing it to a simple nonparametric smoothing of the data. Remembering that when Y is binary, $\Pr(Y = 1|X = x) = \mathbb{E}[Y|X = x]$, we can use a smoothing spline to estimate $\mathbb{E}[Y|X = x]$ (Figure 11.8). This would not be so great as a model — it ignores the fact that the response is a binary event and we're trying to estimate a probability, the fact that the variance of Y therefore depends on its mean, etc. — but it's at least suggestive.

The result starts out notably above the logistic regression, then levels out and climbs much more slowly. It also has a bunch of dubious-looking wiggles, despite the cross-validation.

We can try to do better by fitting a generalized additive model. In this case, with only one predictor variable, this means using non-parametric smoothing to estimate the log odds — we're still using the logistic transformation, but only requiring that the log odds change smoothly with X, not that they be linear in X. The result (Figure 11.9) is initially similar to the spline, but has some more exaggerated undulations, and has confidence intervals. At the largest values of X, the latter span nearly the whole range from 0 to 1, which is not unreasonable considering the sheer lack of data there.

Visually, the logistic regression curve is hardly ever within the confidence limits of the non-parametric predictor. What can we say about the difference between the two models more quantiatively?

Numerically, the deviance is 1.59×10^4 for the logistic regression, and 1.51×10^4 for the GAM. We can go through the testing procedure outlined in §11.6. We need a simulator (which presumes that the logistic regression model is true), and we need to calculate the difference in deviance on simulated data many times.

```
snoq.sim <- function(model = snoq.logistic) {
    fitted.probs = fitted(model)
    return(rbinom(n = length(fitted.probs), size = 1, prob = fitted.probs))
}
```

A quick check of the simulator against the observed values:

```
summary(ifelse(snoq[, 1] > 0, 1, 0))
##    Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##  0.0000  0.0000  1.0000  0.5262  1.0000  1.0000
summary(snoq.sim())
##    Min. 1st Qu.  Median    Mean 3rd Qu.    Max.
##  0.0000  0.0000  1.0000  0.5241  1.0000  1.0000
```

This suggests that the simulator is not acting crazily.
Now for the difference in deviances:

```
diff.dev <- function(model = snoq.logistic, x = snoq[, "today"]) {
    y.new <- snoq.sim(model)
    GLM.dev <- glm(y.new ~ x, family = binomial)$deviance
    GAM.dev <- gam(y.new ~ s(x), family = binomial)$deviance
    return(GLM.dev - GAM.dev)
}
```

```
plot((tomorrow > 0) ~ today, data = snoq, xlab = "Precipitation today (1/100 inch)",
    ylab = "Positive precipitation tomorrow?")
rug(snoq$today, side = 1, col = "grey")
data.plot <- data.frame(today = (0:500))
pred.bands(snoq.logistic, data.plot)
snoq.spline <- smooth.spline(x = snoq$today, y = (snoq$tomorrow > 0))
lines(snoq.spline, col = "red")
```

Figure 11.8 As Figure 11.7, plus a smoothing spline (red).

A single run of this takes about 0.6 seconds on my computer.

Finally, we calculate the distribution of difference in deviances under the null (that the logistic regression is properly specified), and the corresponding *p*-value:

```
diff.dev.obs <- snoq.logistic$deviance - snoq.gam$deviance
null.dist.of.diff.dev <- replicate(100, diff.dev())
p.value <- (1 + sum(null.dist.of.diff.dev > diff.dev.obs))/(1 + length(null.dist.of.diff.dev))
```

Using a thousand replicates takes about 67 seconds, or a bit over a minute; it gives a p-value of $< 1/101$. (A longer run of 1000 replicates, not shown, gives a p-values of $< 10^{-3}$.)

Having detected that there is a problem with the logistic model, we can ask where it lies. We *could* just use the GAM, but it's more interesting to try to diagnose what's going on.

In this respect Figure 11.9 is actually a little misleading, because it leads the eye to emphasize the disagreement between the models at large X, when actually there are very few data points there, and so even large differences in predicted probabilities there contribute little to the over-all likelihood difference. What is actually more important is what happens at $X = 0$, which contains a very large number of observations (about 47% of all observations), and which we have reason to think is a special value anyway.

Let's try introducing a dummy variable for $X = 0$ into the logistic regression, and see what happens. It will be convenient to augment the data frame with an extra column, recording 1 whenever $X = 0$ and 0 otherwise.

```
snoq2 <- data.frame(snoq, dry = ifelse(snoq$today == 0, 1, 0))
snoq2.logistic <- glm((tomorrow > 0) ~ today + dry, data = snoq2, family = binomial)
snoq2.gam <- gam((tomorrow > 0) ~ s(today) + dry, data = snoq2, family = binomial)
```

Notice that I allow the GAM to treat zero as a special value as well, by giving it access to that dummy variable. In principle, with enough data it can decide whether or not that is useful on its own, but since we have guessed that it is, we might as well include it. The new GLM has a deviance of 1.5×10^4, lower than even the GAM before, and the new GAM has a deviance of 1.48×10^4. I will leave repeating the specification test as an exercise. Figure 11.10 shows the data and the two new models. These are *extremely* close to each other at low percipitation, and diverge thereafter. The new GAM is the smoothest model we've seen yet, which suggests that before the it was being under-smoothed to help capture the special value at zero.

Let's turn now to looking at calibration. The actual fraction of no-precipitation days which are followed by precipitation is

```
signif(mean(snoq$tomorrow[snoq$today == 0] > 0), 3)
## [1] 0.287
```

What does the new logistic model predict?

```
signif(predict(snoq2.logistic, newdata = data.frame(today = 0, dry = 1), type = "response"),
    3)
##     1
## 0.287
```

This should not be surprising — we've given the model a special parameter

dedicated to getting this one probability exactly right! The hope however is that this will change the predictions made on days *with* precipitation so that they are better.

Looking at a histogram of fitted values (`hist(fitted(snoq2.logistic))`) shows a gap in the distribution of predicted probabilities below 0.63, so we'll look first at days where the predicted probability is between 0.63 and 0.64.

```
signif(mean(snoq$tomorrow[(fitted(snoq2.logistic) >= 0.63) & (fitted(snoq2.logistic) <
    0.64)] > 0), 3)
## [1] 0.526
```

Not bad — but a bit painful to write out. Let's write a function:

```
frequency.vs.probability <- function(p.lower, p.upper = p.lower + 0.01, model = snoq2.logistic,
    events = (snoq$tomorrow > 0)) {
  fitted.probs <- fitted(model)
  indices <- (fitted.probs >= p.lower) & (fitted.probs < p.upper)
  ave.prob <- mean(fitted.probs[indices])
  frequency <- mean(events[indices])
  se <- sqrt(ave.prob * (1 - ave.prob)/sum(indices))
  return(c(frequency = frequency, ave.prob = ave.prob, se = se))
}
```

I have added a calculation of the average predicted probability, and a crude estimate of the standard error we should expect if the observations really are binomial with the predicted probabilities[11]. Try the function out before doing anything rash:

```
frequency.vs.probability(0.63)
## frequency    ave.prob           se
## 0.52603037 0.63414568 0.01586292
```

This agrees with our previous calculation.
Now we can do this for a lot of probability brackets:

```
f.vs.p <- sapply(c(0.28, (63:100)/100), frequency.vs.probability)
```

This comes with some unfortunate R cruft, removable thus

```
f.vs.p <- data.frame(frequency = f.vs.p["frequency", ], ave.prob = f.vs.p["ave.prob",
    ], se = f.vs.p["se", ])
```

and we're ready to plot (Figure 11.11). The observed frequencies are generally reasonably near the predicted probabilities. While I wouldn't want to say this was the last word in weather forecasting[12], it's surprisingly good for such a simple model. I will leave calibration checking for the GAM as another exercise.

---

[11] This could be improved by averaging predicted variances for each point, but using probability ranges of 0.01 makes it hardly worth the effort.

[12] There is an extensive discussion of this data in Guttorp (1995, ch. 2), including many significant refinements, such as dependence across multiple days.

```
library(mgcv)
plot((tomorrow > 0) ~ today, data = snoq, xlab = "Precipitation today (1/100 inch)",
    ylab = "Positive precipitation tomorrow?")
rug(snoq$today, side = 1, col = "grey")
pred.bands(snoq.logistic, data.plot)
lines(snoq.spline, col = "red")
snoq.gam <- gam((tomorrow > 0) ~ s(today), data = snoq, family = binomial)
pred.bands(snoq.gam, data.plot, "blue")
```

Figure 11.9 As Figure 11.8, but with the addition of a generalized additive model (blue line) and its confidence limits (dashed blue lines).

```
plot((tomorrow > 0) ~ today, data = snoq, xlab = "Precipitation today (1/100 inch)",
    ylab = "Positive precipitation tomorrow?")
rug(snoq$today, side = 1, col = "grey")
data.plot = data.frame(data.plot, dry = ifelse(data.plot$today == 0, 1, 0))
lines(snoq.spline, col = "red")
pred.bands(snoq2.logistic, data.plot)
pred.bands(snoq2.gam, data.plot, "blue")
```

Figure 11.10 As Figure 11.9, but allowing the two models to use a dummy variable indicating when today is completely dry ($X = 0$).

```
plot(frequency ~ ave.prob, data = f.vs.p, xlim = c(0, 1), ylim = c(0, 1), xlab = "Predicted probabil
    ylab = "Observed frequencies")
rug(fitted(snoq2.logistic), col = "grey")
abline(0, 1, col = "grey")
segments(x0 = f.vs.p$ave.prob, y0 = f.vs.p$ave.prob - 1.96 * f.vs.p$se, y1 = f.vs.p$ave.prob +
    1.96 * f.vs.p$se)
```

Figure 11.11 Calibration plot for the modified logistic regression model
`snoq2.logistic`. Points show the actual frequency of precipitation for each
level of predicted probability. Vertical lines are (approximate) 95% sampling
intervals for the frequency, given the predicted probability and the number
of observations.

11.8 Logistic Regression with More Than Two Classes

If Y can take on more than two values, say k of them, we can still use logistic regression. Instead of having one set of parameters β_0, β, each class c in $0 : (k-1)$ will have its own offset $\beta_0^{(c)}$ and vector $\beta^{(c)}$, and the predicted conditional probabilities will be

$$\Pr\left(Y = c | \vec{X} = x\right) = \frac{e^{\beta_0^{(c)} + x \cdot \beta^{(c)}}}{\sum_c e^{\beta_0^{(c)} + x \cdot \beta^{(c)}}} \tag{11.23}$$

You can check that when there are only two classes (say, 0 and 1), equation 11.23 reduces to equation 11.5, with $\beta_0 = \beta_0^{(1)} - \beta_0^{(0)}$ and $\beta = \beta^{(1)} - \beta^{(0)}$. In fact, no matter how many classes there are, we can always pick one of them, say $c = 0$, and fix its parameters at exactly zero, without any loss of generality (Exercise 11.2)[13].

Calculation of the likelihood now proceeds as before (only with more book-keeping), and so does maximum likelihood estimation.

As for R implementations, for am long time the easiest way to do this was actually to use the `nnet` package for neural networks (Venables and Ripley, 2002). More recently, the `multiclass` function from the `mgcv` package does the same sort of job, with an interface closer to what you will be familiar with from `glm` and `gam`.

---

[13] Since we can arbitrarily chose which class's parameters to "zero out" without affecting the predicted probabilities, strictly speaking the model in Eq. 11.23 is **unidentified**. That is, different parameter settings lead to *exactly* the same outcome, so we can't use the data to tell which one is right. The usual response here is to deal with this by a convention: we *decide* to zero out the parameters of the first class, and then estimate the contrasting parameters for the others.

Exercises

11.1 *"We minimize the mis-classification rate by predicting the most likely class"*: Let $\widehat{\mu}(x)$ be our predicted class, either 0 or 1. Our error rate is then $\Pr(Y \neq \widehat{\mu})$. Show that $\Pr(Y \neq \widehat{\mu}) = \mathbb{E}\left[(Y - \widehat{\mu})^2\right]$. Further show that $\mathbb{E}\left[(Y - \widehat{\mu})^2 \mid X = x\right] = \Pr(Y = 1 | X = x)(1 - 2\widehat{\mu}(x)) + \widehat{\mu}^2(x)$. Conclude by showing that if $\Pr(Y = 1 | X = x) > 0.5$, the risk of mis-classification is minimized by taking $\widehat{\mu} = 1$, that if $\Pr(Y = 1 | X = x) < 0.5$ the risk is minimized by taking $\widehat{\mu} = 0$, and that when $\Pr(Y = 1 | X = x) = 0.5$ both predictions are equally risky.

11.2 A multiclass logistic regression, as in Eq. 11.23 has parameters $\beta_0^{(c)}$ and $\beta^{(c)}$ for each class c. Show that we can always get the same predicted probabilities by setting $\beta_0^{(c)} = 0$, $\beta^{(c)} = 0$ for any one class c, and adjusting the parameters for the other classes appropriately.

11.3 Find the first and second derivatives of the log-likelihood for logistic regression with one predictor variable. Explicitly write out the formula for doing one step of Newton's method. Explain how this relates to re-weighted least squares.

11.4 *Intuition-building with logistic regression*

In a classification problem, we get data like that shown above, with the two classes indicated by whether a point is plotted with $+$ or $-$. We decide to use a logistic regression model. We estimate $\hat{\beta}_0 = 0$, $\hat{\beta}_1 = 1$ and $\hat{\beta}_2 = 1$.

1. Draw the line where the log-odds $= 0$, and so where $p = 1/2$. Does this separate most of the $+$ points from the $-$ points? Does it separate *all* of them?
2. Explain why the points marked A and B will have exactly the same probability of $Y = 1$, and why that probability will be $> 1/2$.
3. Explain why the point marked C will have nearly the same probability for $Y = 1$ as A and B do, while point D will have a much lower probability, even though D is much closer than C to both A and B.