

# Stat 5421 Lecture Notes: To Accompany Agresti Ch 4

Charles J. Geyer

October 16, 2020

## Section 4.2.2

```
heart.disease <- rbind(c(24, 1355),
                     c(35, 603),
                     c(21, 192),
                     c(30, 224))
snoring <- c("never", "occasionally", "nearly.every.night", "every.night")
x <- c(0, 2, 4, 5)
```

Then we fit the logistic GLM (usually called *logistic regression*) by

```
gout <- glm(heart.disease ~ x, family = binomial)
summary(gout)
```

```
##
## Call:
## glm(formula = heart.disease ~ x, family = binomial)
##
## Deviance Residuals:
##      1       2       3       4
## -0.8346  1.2521  0.2758 -0.6845
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept) -3.86625    0.16621 -23.261 < 2e-16 ***
## x              0.39734    0.05001   7.945 1.94e-15 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##    Null deviance: 65.9045  on 3  degrees of freedom
## Residual deviance:  2.8089  on 2  degrees of freedom
## AIC: 27.061
##
## Number of Fisher Scoring iterations: 4
```

We show this model first because it is the one we and Agresti recommend and is by far the most widely used.

Now we do the nonsense model, that no one ever uses, except teachers sometimes show it to show why all your intuitions developed from learning about linear models do not transfer to GLM.

```
gout.goofy <- glm(heart.disease ~ x, family = binomial(link="identity"))
summary(gout.goofy)
```

```
##
## Call:
## glm(formula = heart.disease ~ x, family = binomial(link = "identity"))
##
## Deviance Residuals:
##      1      2      3      4
## 0.04478 -0.21322  0.11010  0.09798
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept) 0.017247   0.003451   4.998 5.80e-07 ***
## x           0.019778   0.002805   7.051 1.77e-12 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
## Null deviance: 65.904481  on 3  degrees of freedom
## Residual deviance:  0.069191  on 2  degrees of freedom
## AIC: 24.322
##
## Number of Fisher Scoring iterations: 3
```

We make a plot like Figure~4.1 in Agresti.

```
mu <- predict(gout, type = "response")
mu.goofy <- predict(gout.goofy, type = "response")
mu
```

```
##      1      2      3      4
## 0.02050742 0.04429511 0.09305411 0.13243885
```

```
mu.goofy
```

```
##      1      2      3      4
## 0.01724668 0.05680231 0.09635793 0.11613574
```

```
plot(x, mu, ylab = "probability", pch = 19, ylim = range(mu, mu.goofy))
curve(predict(gout, newdata = data.frame(x = x), type = "response"),
      add = TRUE)
grep("red", colors(), value = TRUE)
```

```
## [1] "darkred"      "indianred"     "indianred1"    "indianred2"
## [5] "indianred3"    "indianred4"    "mediumvioletred" "orangered"
## [9] "orangered1"    "orangered2"    "orangered3"    "orangered4"
## [13] "palevioletred" "palevioletred1" "palevioletred2" "palevioletred3"
## [17] "palevioletred4" "red"           "red1"          "red2"
## [21] "red3"          "red4"          "violetred"     "violetred1"
## [25] "violetred2"    "violetred3"    "violetred4"
```

```
points(x, mu.goofy, col = "red", pch = 19)
curve(predict(gout.goofy, newdata = data.frame(x = x), type = "response"),
      add = TRUE, col = "red")
```




Figure 1: logit link (black) versus identity link (red)

## Section 4.2.2 Comments

### Linearity versus Constraints

Linearity on the mean scale (identity link) fights the constraints that probabilities have to be between zero and one. Let us look at a few examples where we have probabilities near zero and one for our regression function.

Make up some data where the logistic regression model is true and probabilities go from near zero to near one.

```
z <- 1:100
theta <- z - mean(z)
theta <- 3 * theta / max(theta)
p <- 1 / (1 + exp(- theta))
y <- rbinom(length(p), 1, p)
rm(theta, p)
```


Now fit the logistic regression model and plot it.

```
gout.logit <- glm(y ~ z, family = binomial)
summary(gout.logit)
```

```
##
## Call:
## glm(formula = y ~ z, family = binomial)
##
## Deviance Residuals:
##      Min       1Q   Median       3Q      Max
## -2.2313  -0.6284   0.2706   0.6255   2.0897
##
```

```
## Coefficients:
##           Estimate Std. Error z value Pr(>|z|)
## (Intercept) -3.24630    0.66584  -4.876 1.09e-06 ***
## z           0.06569    0.01228   5.347 8.93e-08 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
## Null deviance: 138.589  on 99  degrees of freedom
## Residual deviance:  88.277  on 98  degrees of freedom
## AIC: 92.277
##
## Number of Fisher Scoring iterations: 5
```

```
plot(z, y)
curve(predict(gout.logit, newdata = data.frame(z = x),
            type = "response"), add = TRUE)
```



Now fit the same data with identity link and add it to the plot.

```
gout.identity <- glm(y ~ z, family = binomial(link = "identity"))
```

```
## Error: no valid set of coefficients has been found: please supply starting values
```

This model is so bad that R function glm has trouble fitting it. Try two.

```
gout.identity <- glm(y ~ z, family = binomial(link = "identity"),
                    start = c(0.5, 0))
```

```
## Warning: step size truncated: out of bounds
```

```
## Warning: step size truncated: out of bounds
```

```
## Warning: step size truncated: out of bounds
```


```
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: step size truncated: out of bounds
## Warning: glm.fit: algorithm did not converge
## Warning: glm.fit: algorithm stopped at boundary value
summary(gout.identity)

##
## Call:
## glm(formula = y ~ z, family = binomial(link = "identity"), start = c(0.5,
##    0))
##
```

```
## Deviance Residuals:
##      Min       1Q   Median       3Q      Max
## -1.97808 -0.76781  0.07124  0.80224  1.87631
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept) -0.0097589  0.0204365  -0.478   0.633
## z            0.0100976  0.0002044  49.410 <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
##      Null deviance: 138.589  on 99  degrees of freedom
## Residual deviance:  87.811  on 98  degrees of freedom
## AIC: 91.811
##
## Number of Fisher Scoring iterations: 25
```

Now plot.

```
plot(z, y)
curve(predict(gout.logit, newdata = data.frame(z = x),
  type = "response"), add = TRUE)
curve(predict(gout.identity, newdata = data.frame(z = x),
  type = "response"), add = TRUE, col = "red")
```



- The identity link model is so bad that R function `glm` has a lot of problems with it and
- we can make the identity link look as bad as we want: it clearly is only paying attention to the endpoints of the data and ignoring everything in between.

## Hypothesis Tests of Non-Nested Models

Although theory exists for comparing non-nested models (Cox, 1961, 1962, 2013), it is not widely known or used. The simplest procedure that makes sense is to compare both to the smallest model containing both. This is the model with four parameters that fits the data perfectly.

```
gout.super <- glm(heart.disease ~ snoring, family = binomial)
summary(gout.super)
```

```
##
## Call:
## glm(formula = heart.disease ~ snoring, family = binomial)
##
## Deviance Residuals:
## [1] 0 0 0 0
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)    -2.0104    0.1944 -10.341 < 2e-16 ***
## snoringnearly.every.night -0.2025    0.3010  -0.673  0.50111
## snoringnever    -2.0231    0.2832  -7.144  9.1e-13 ***
## snoringoccasionally  -0.8361    0.2608  -3.206  0.00135 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
## Null deviance:  6.5904e+01 on 3 degrees of freedom
## Residual deviance: -1.2457e-13 on 0 degrees of freedom
## AIC: 28.253
##
## Number of Fisher Scoring iterations: 3
```

```
gout.goofy.super <- glm(heart.disease ~ snoring,
  family = binomial(link="identity"))
summary(gout.goofy.super)
```

```
##
## Call:
## glm(formula = heart.disease ~ snoring, family = binomial(link = "identity"))
##
## Deviance Residuals:
## [1] 0 0 0 0
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)    0.11811    0.02025    5.832 5.46e-09 ***
## snoringnearly.every.night -0.01952    0.02876  -0.679  0.49739
## snoringnever    -0.10071    0.02055  -4.900 9.61e-07 ***
## snoringoccasionally  -0.06325    0.02217  -2.853  0.00432 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for binomial family taken to be 1)
##
```

```
## Null deviance: 6.5904e+01 on 3 degrees of freedom
## Residual deviance: -1.5765e-14 on 0 degrees of freedom
## AIC: 28.253
##
## Number of Fisher Scoring iterations: 2
mu.super <- predict(gout.super, type = "response")
mu.goofy.super <- predict(gout.goofy.super, type = "response")
all.equal(mu.super, mu.goofy.super)
```

```
## [1] TRUE
```

We fit the model with two different link functions, even though they are the same model, as proved by the `all.equal` statement. The reason for this is we are not sure how R function `anova` deals with different link functions. We compare the two models of interest to their least common supermodel as follows.

```
anova(gout, gout.super, test = "LRT")
```

```
## Analysis of Deviance Table
##
## Model 1: heart.disease ~ x
## Model 2: heart.disease ~ snoring
##   Resid. Df Resid. Dev Df Deviance Pr(>Chi)
## 1         2     2.8089
## 2         0     0.0000  2   2.8089  0.2455
```

```
anova(gout.goofy, gout.goofy.super, test = "LRT")
```

```
## Analysis of Deviance Table
##
## Model 1: heart.disease ~ x
## Model 2: heart.disease ~ snoring
##   Resid. Df Resid. Dev Df Deviance Pr(>Chi)
## 1         2   0.069191
## 2         0   0.000000  2 0.069191  0.966
```

We see that both models apparently fit these data, but what we are calling the “goofy” model actually fits better.

## 4.2.5

There are other link functions that can be used.

```
gout.probit <- glm(heart.disease ~ x, family = binomial(link="probit"))
gout.cauchit <- glm(heart.disease ~ x, family = binomial(link="cauchit"))
mu.probit <- predict(gout.probit, type = "response")
mu.cauchit <- predict(gout.cauchit, type = "response")

plot(x, mu, ylab = "probability", pch = 19,
     ylim = range(mu, mu.goofy, mu.probit, mu.cauchit))
curve(predict(gout, newdata = data.frame(x = x), type = "response"),
      add = TRUE)
points(x, mu.goofy, col = "red", pch = 19)
curve(predict(gout.goofy, newdata = data.frame(x = x), type = "response"),
      add = TRUE, col = "red")
points(x, mu.probit, col = "yellow2", pch = 19)
```



```

curve(predict(gout.probit, newdata = data.frame(x = x), type = "response"),
      add = TRUE, col = "yellow2")
grep("yellow", colors(), value = TRUE)

```

```

## [1] "greenyellow"          "lightgoldenrodyellow" "lightyellow"
## [4] "lightyellow1"         "lightyellow2"         "lightyellow3"
## [7] "lightyellow4"         "yellow"                "yellow1"
## [10] "yellow2"              "yellow3"               "yellow4"
## [13] "yellowgreen"

```

```

points(x, mu.cauchit, col = "blue", pch = 19)
curve(predict(gout.cauchit, newdata = data.frame(x = x), type = "response"),
      add = TRUE, col = "blue")

```




Figure 2: logit link (black) versus identity link (red) versus probit link (yellow) versus cauchit link (blue)

Your humble instructor recommends logit link for reasons that will become apparent when we discuss exponential families. One really has to like what Agresti calls the “latent tolerance” motivation to use probit or cauchit. And even then, logit link is also a “latent tolerance” model based on the logistic distribution. So it is unclear that there is any reason for avoiding the logit link even if one likes the “latent tolerance” story.

## Confidence Intervals

### Wald Intervals

These are produced by R function `predict`. From now on we will only use the fit in R object `gout`, which uses the default logit link and has 2 parameters.

First, the output of `summary` gives standard errors, so if one wanted confidence intervals for the betas, one would do them as follows.

```

sout <- summary(gout)
class(sout)

## [1] "summary.glm"
class(unclass(sout))

## [1] "list"
names(unclass(sout))

## [1] "call"          "terms"          "family"         "deviance"
## [5] "aic"           "contrasts"     "df.residual"   "null.deviance"
## [9] "df.null"      "iter"          "deviance.resid" "coefficients"
## [13] "aliased"      "dispersion"    "df"            "cov.unscaled"
## [17] "cov.scaled"

cout <- sout$coefficients
class(cout)

## [1] "matrix" "array"
cout

##           Estimate Std. Error   z value    Pr(>|z|)
## (Intercept) -3.8662481 0.16621436 -23.260614 1.110885e-119
## x           0.3973366 0.05001066  7.945039  1.941304e-15

conf.level <- 0.95
crit <- qnorm((1 + conf.level) / 2)
crit

## [1] 1.959964

lower <- cout[ , "Estimate"] - crit * cout[ , "Std. Error"]
upper <- cout[ , "Estimate"] + crit * cout[ , "Std. Error"]
foo <- data.frame(lower, upper)
foo

##           lower      upper
## (Intercept) -4.1920223 -3.5404739
## x           0.2993175  0.4953557

However, the regression coefficients (betas) are meaningless quantities. They are very far from the data, and the model can be reparameterized without changing the model. Thus we look at confidence intervals for probabilities.


pout <- predict(gout, type = "response", se.fit = TRUE)
lower <- pout$fit - crit * pout$se.fit
upper <- pout$fit + crit * pout$se.fit
data.frame(lower, upper)

##           lower      upper
## 1 0.01396364 0.02705120
## 2 0.03561897 0.05297125
## 3 0.07330823 0.11279999
## 4 0.09798190 0.16689580

library(sfsmisc)
errbar(x, mu, upper, lower, ylab = "probability")

```

```
\begin{figure}
```




```
\caption{95% Wald confidence intervals, not simultaneous} \end{figure}
```

These are not, of course, corrected for multiple comparisons.

We can also get Wald intervals at  $x$  values other than those in the observed data. It is not clear that these make any sense in these data, but we will illustrate them anyway.

```
plot(x, mu, ylim = range(mu, upper, lower), ylab = "probability")
curve(predict(gout, newdata = data.frame(x = x), type = "response"),
  add = TRUE)
sally <- function(x) {
  pout <- predict(gout, type = "response", se.fit = TRUE,
    newdata = data.frame(x = x))
  pout$fit - crit * pout$se.fit
}
curve(sally, add = TRUE, lty = "dashed")
sally <- function(x) {
  pout <- predict(gout, type = "response", se.fit = TRUE,
    newdata = data.frame(x = x))
  pout$fit + crit * pout$se.fit
}
curve(sally, add = TRUE, lty = "dashed")
```

```
\begin{figure}
```



\caption{95% Wald confidence bands, not simultaneous} \end{figure}

## Likelihood Intervals

### R Function Confint

R has a built-in function for doing likelihood intervals

```
library(MASS)
confint(gout)
```

```
## Waiting for profiling to be done...
##           2.5 %      97.5 %
## (Intercept) -4.2072190 -3.5544117
## x           0.2999362  0.4963887
```

Unfortunately these intervals are for the meaningless parameters (regression coefficients) rather than the meaningful parameters (success probability as a function of  $x$ ). To calculate those we need to use the methods of the handout on likelihood computation.

### Log Likelihood

But even before we try that, we need the log likelihood function, which we have not had explicitly, because we were using R function `glm`, which knows the log likelihood function but doesn't return it to the user.

For this problem, the log likelihood is

$$\sum_i [x_i \log(p_i) - (n_i - x_i) \log(1 - p_i)]$$

where each row of the data matrix `heart.disease` is  $(x_i, n_i - x_i)$  and the  $p_i$  are components of a vector  $\mathbf{p}$  that satisfies

$$\boldsymbol{\theta} = \mathbf{M}\boldsymbol{\beta}$$

where  $\mathbf{M}$  is the model matrix,  $\boldsymbol{\beta}$  is the vector of regression coefficients, and  $\boldsymbol{\theta}$  is the so-called “linear predictor” which is related to  $\mathbf{p}$  by the inverse link function.

$$p_i = \frac{e^{\theta_i}}{1 + e^{\theta_i}} = \frac{1}{1 + e^{-\theta_i}}$$

so

$$1 - p_i = \frac{1}{1 + e^{\theta_i}} = \frac{e^{-\theta_i}}{1 + e^{-\theta_i}}$$

and

$$\begin{aligned} \log(p_i) &= \theta_i - \log(1 + e^{\theta_i}) = -\log(1 + e^{-\theta_i}) \\ \log(1 - p_i) &= -\log(1 + e^{\theta_i}) = -\theta_i - \log(1 + e^{-\theta_i}) \end{aligned}$$

The reason why we have two expressions for each of these is that  $e^{\theta}$  may overflow (be a number too big for the computer to represent) when  $\theta$  is moderately large

```
exp(800)
```

```
## [1] Inf
```

and, of course,  $e^{-\theta}$  can overflow when  $\theta$  is moderately large negative. Thus we want to use whichever formula cannot overflow, the one containing  $e^{-\theta}$  when  $\theta > 0$  and the one containing  $e^{\theta}$  when  $\theta$  is negative.

Also we never want to calculate  $1 - p_i$  by subtraction because that can cause “catastrophic cancellation”, loss of all significant figures, when  $p_i$  is near zero.

```
1 + 1e-20 - 1
```

```
## [1] 0
```

Hence we need all 4 formulas above.

Finally, to avoid more catastrophic cancellation, we use R function `log1p` which calculates the function  $x \mapsto \log(1 + x)$  more accurately when  $x$  is small than just using the `log` function and addition. From calculus  $\log(1 + x) \approx x$  when  $x$  is near zero.

```
xx <- c(1e-20, 1e-10, 1, 10)
log1p(xx)
```

```
## [1] 1.000000e-20 1.000000e-10 6.931472e-01 2.397895e+00
```

```
log(1 + xx)
```

```
## [1] 0.0000000000 0.0000000001 0.6931471806 2.3978952728
```

Such finicky calculation is also something we do not expect for homework but which is necessary to write good code useable by others. This subject of computer arithmetic is really the subject of another course, so we won’t say much about it in this course. Those that are interested can look at my Stat 3701 notes on this subject.

We also need the model matrix  $M$ . We could figure it out for ourselves, but we can also get it from the output of R function `glm` if we ask for it.

```
gout <- glm(heart.disease ~ x, family = binomial, x = TRUE)
gout$x
```

```
## (Intercept) x
## 1          1 0
## 2          1 2
## 3          1 4
## 4          1 5
## attr("assign")
## [1] 0 1
```

Now we are ready to code the log likelihood function. We use a factory function, but, of course, the global variables approach can also be used.

```
mlogl.factory <- function(modmat, response) {
  stopifnot(is.numeric(modmat))
  stopifnot(is.finite(modmat))
  stopifnot(is.matrix(modmat))
  stopifnot(is.numeric(response))
  stopifnot(is.finite(response))
  stopifnot(is.matrix(response))
  stopifnot(round(response) == response) # check for integer-valued
  stopifnot(response >= 0)
  stopifnot(ncol(response) == 2)
  stopifnot(nrow(response) == nrow(modmat))
  success <- response[, 1]
  failure <- response[, 2]
  function(beta) {
    stopifnot(is.numeric(beta))
    stopifnot(is.finite(beta))
    stopifnot(length(beta) == ncol(modmat))
    theta <- as.vector(modmat %*% beta)
    logp <- ifelse(theta > 0, - log1p(exp(- theta)),
                  theta - log1p(exp(theta)))
    logq <- ifelse(theta > 0, - theta - log1p(exp(- theta)),
                  - log1p(exp(theta) ))
    - sum(success * logp + failure * logq)
  }
}
```

We should have said when we introduced the function factory approach the reasons why it is better than the global variables approach are

- global variables are evil — they are easily clobbered accidentally before they are used, or they can clobber other variables of the same name being used by the user for other purposes and
- the factory function can check the valid of what would be the global variables in the other approach — and this check is done once at the time the `mlogl` function is created rather than every time the `mlogl` function is invoked (as would have to be done in the global variables approach where there is no factory function to do this checking)

```
mlogl <- mlogl.factory(gout$x, heart.disease)
```

Let us check that our function works and gives the same MLE as R function `glm`

```
nout <- nlm(mlogl, c(0, 0))
nout$code <= 2
```

```
## [1] TRUE
```

```
all.equal(nout$estimate, gout$coefficients, check.attributes = FALSE)
```

```
## [1] TRUE
```

Good! The reason why we did not bother to find a “good” starting point for the optimization is that we know this log likelihood function has a unique local maximum, which is also the global maximum, that is found from any starting point. But the explanation for this will have to wait until we study exponential families of distributions.

Now we also need the function that maps  $\beta$  to  $\mathbf{p}$ .

```
pooh.factory <- function(x) {  
  stopifnot(is.numeric(x))  
  stopifnot(is.finite(x))  
  stopifnot(length(x) == 1)  
  function(beta) {  
    stopifnot(is.numeric(beta))  
    stopifnot(is.finite(beta))  
    stopifnot(length(beta) == 2)  
    theta <- sum(c(1, x) * beta)  
    1 / (1 + exp(- theta))  
  }  
}
```

Try it out.

```
pooh <- pooh.factory(0)  
pooh(nout$estimate)
```

```
## [1] 0.02050742
```

```
mu[1]
```

```
##          1  
## 0.02050742
```

Now we are finally ready to calculate the likelihood intervals.

```
library(alabama)
```

```
## Loading required package: numDeriv
```

```
lower <- double(length(x))  
for (i in seq(along = x)) {  
  aout <- suppressWarnings(auglag(nout$estimate,  
    fn = pooh.factory(x[i]),  
    hin = function(theta) nout$minimum + crit - mlogl(theta),  
    control.outer = list(trace = FALSE)))  
  stopifnot(aout$convergence == 0)  
  lower[i] <- aout$value  
}  
upper <- double(length(x))  
for (i in seq(along = x)) {  
  aout <- suppressWarnings(auglag(nout$estimate,  
    fn = pooh.factory(x[i]),  
    hin = function(theta) nout$minimum + crit - mlogl(theta),  
    control.outer = list(trace = FALSE),  
    control.optim = list(fnscale = -1)))
```

```

stopifnot(aout$convergence == 0)
upper[i] <- aout$value
}
data.frame(lower, upper)

```

```

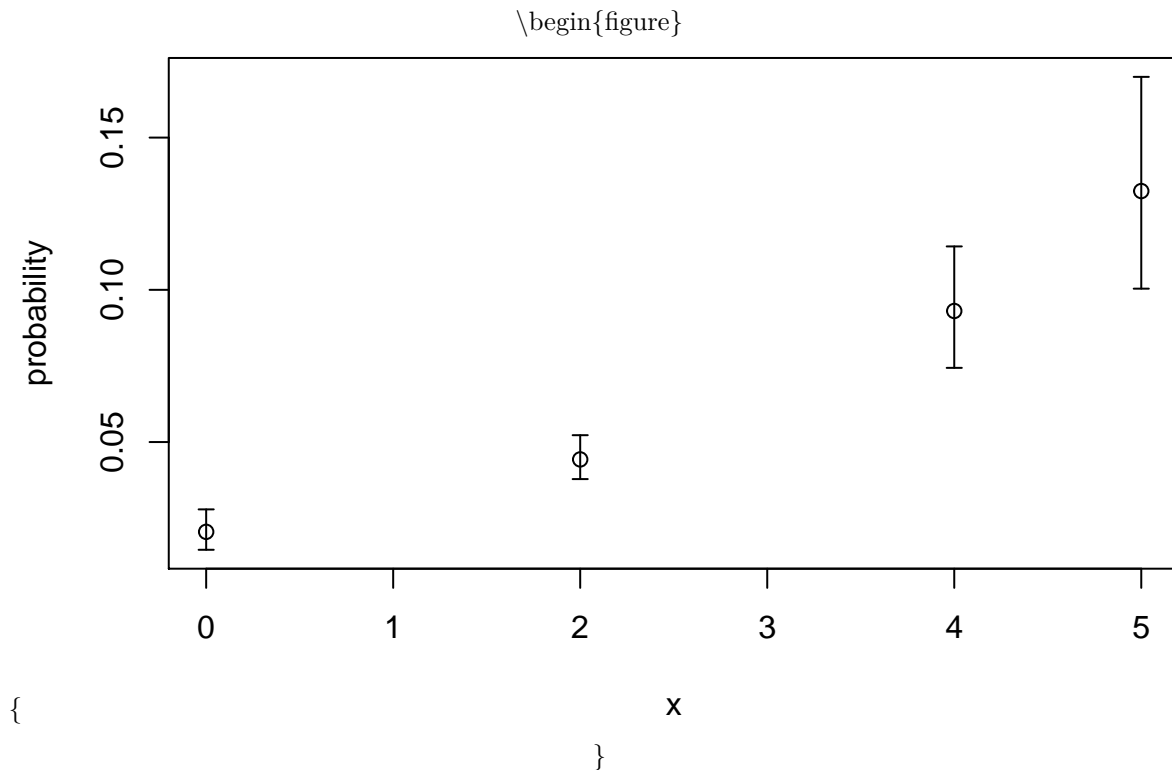
##      lower      upper
## 1 0.01461808 0.02788185
## 2 0.03783432 0.05225022
## 3 0.07436824 0.11426074
## 4 0.10038659 0.16994416

```

```

errbar(x, mu, upper, lower, ylab = "probability")

```



```

\caption{95% likelihood confidence intervals, not simultaneous} \end{figure}

```

Of course, if we want simultaneous coverage, we can just use the appropriate critical value for that, degrees of freedom equal to the length of **beta**, here 2, as discussed in the section on simultaneous coverage of the lecture notes on likelihood computation

## Section 4.3

### A Poisson GLM Example

```

# clean out R global environment and start fresh
rm(list = ls())

```




## Data

There are also GLM for Poisson response. For an example, we will use some data analyzed in my 5102 lecture notes

```
foo <- read.table("http://www.stat.umn.edu/geyer/5102/data/ex6-4.txt",
  header = TRUE)
names(foo)
```

```
## [1] "hour" "count"
```

```
hour <- foo$hour
count <- foo$count
plot(hour, count)
```




These (simulated) data are hourly counts from a not necessarily homogeneous Poisson process. The variables are

- `hour` which counts hours sequentially throughout a 14-day period (running from 1 to  $14 \times 24 = 336$ ) and
- `count` giving the second giving the count for each of these hours.

The idea of the regression is to estimate the mean as a function of time. Many time series have a daily cycle. If we pool the counts for the same hour of the day over the 14 days of the series, we see a clear pattern in the histogram.

```
hourofday <- (hour - 1) %% 24 + 1
fred <- structure(list(breaks = 0:24,
  counts = sapply(split(count, hourofday), sum),
  mids = 1:24 - 1 / 2, xname = "hour of day",
  equidist = TRUE), class = "histogram")
plot(fred)
```

## Histogram of hour of day



The somewhat mysterious code above makes an R object of class "histogram" and hands it off to R generic function `plot` which has a method for objects of this class. The requirements for such objects are documented in the help page for R function `histogram`.

## First GLM


Since there seems to be a daily cycle with two peaks we make the "linear predictor" for our GLM a Fourier series with frequencies one per day and two per day.

```
w <- hour / 24 * 2 * pi
out <- glm(count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) +
  I(cos(2 * w)), family = poisson)
summary(out)
```

```
##
## Call:
## glm(formula = count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) +
##   I(cos(2 * w)), family = poisson)
##
## Deviance Residuals:
##   Min       1Q   Median       3Q      Max
## -3.2043  -0.7431  -0.0905   0.6129   3.2662
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept)   1.65917    0.02494  66.516 < 2e-16 ***
## I(sin(w))     -0.13916    0.03128  -4.448 8.66e-06 ***
## I(cos(w))     -0.28510    0.03661  -7.787 6.86e-15 ***
## I(sin(2 * w)) -0.42974    0.03385 -12.696 < 2e-16 ***
```

```
## I(cos(2 * w)) -0.30846    0.03346   -9.219 < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for poisson family taken to be 1)
##
## Null deviance: 704.27  on 335  degrees of freedom
## Residual deviance: 399.58  on 331  degrees of freedom
## AIC: 1535.7
##
## Number of Fisher Scoring iterations: 5
```

```
plot(hourofday, count, xlab = "hour of the day")
curve(predict(out, data.frame(w = x / 24 * 2 * pi),
  type="response"), add=TRUE)
```



Not bad. We seem to have a reasonable regression function for these data.

### Hypothesis Tests

But maybe more terms in our Fourier series would be better?

```
form <- "count ~ I(sin(w)) + I(cos(w))"
for (i in 2:10) {
  form <- c(form, paste0(form[length(form)],
    " + I(sin(", i, " * w)) + I(cos(", i, " * w))"))
}
outs <- list()
for (i in seq(along = form))
  outs[[i]] <- glm(as.formula(form[i]), family = poisson)
class(outs) <- "glmlist"
anova(outs, test = "LRT")
```

```

## Analysis of Deviance Table
##
## Model 1: count ~ I(sin(w)) + I(cos(w))
## Model 2: count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) + I(cos(2 * w))
## Model 3: count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) + I(cos(2 * w)) +
##   I(sin(3 * w)) + I(cos(3 * w))
## Model 4: count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) + I(cos(2 * w)) +
##   I(sin(3 * w)) + I(cos(3 * w)) + I(sin(4 * w)) + I(cos(4 *
##   w))
## Model 5: count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) + I(cos(2 * w)) +
##   I(sin(3 * w)) + I(cos(3 * w)) + I(sin(4 * w)) + I(cos(4 *
##   w)) + I(sin(5 * w)) + I(cos(5 * w))
## Model 6: count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) + I(cos(2 * w)) +
##   I(sin(3 * w)) + I(cos(3 * w)) + I(sin(4 * w)) + I(cos(4 *
##   w)) + I(sin(5 * w)) + I(cos(5 * w)) + I(sin(6 * w)) + I(cos(6 *
##   w))
## Model 7: count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) + I(cos(2 * w)) +
##   I(sin(3 * w)) + I(cos(3 * w)) + I(sin(4 * w)) + I(cos(4 *
##   w)) + I(sin(5 * w)) + I(cos(5 * w)) + I(sin(6 * w)) + I(cos(6 *
##   w)) + I(sin(7 * w)) + I(cos(7 * w))
## Model 8: count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) + I(cos(2 * w)) +
##   I(sin(3 * w)) + I(cos(3 * w)) + I(sin(4 * w)) + I(cos(4 *
##   w)) + I(sin(5 * w)) + I(cos(5 * w)) + I(sin(6 * w)) + I(cos(6 *
##   w)) + I(sin(7 * w)) + I(cos(7 * w)) + I(sin(8 * w)) + I(cos(8 *
##   w))
## Model 9: count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) + I(cos(2 * w)) +
##   I(sin(3 * w)) + I(cos(3 * w)) + I(sin(4 * w)) + I(cos(4 *
##   w)) + I(sin(5 * w)) + I(cos(5 * w)) + I(sin(6 * w)) + I(cos(6 *
##   w)) + I(sin(7 * w)) + I(cos(7 * w)) + I(sin(8 * w)) + I(cos(8 *
##   w)) + I(sin(9 * w)) + I(cos(9 * w))
## Model 10: count ~ I(sin(w)) + I(cos(w)) + I(sin(2 * w)) + I(cos(2 * w)) +
##   I(sin(3 * w)) + I(cos(3 * w)) + I(sin(4 * w)) + I(cos(4 *
##   w)) + I(sin(5 * w)) + I(cos(5 * w)) + I(sin(6 * w)) + I(cos(6 *
##   w)) + I(sin(7 * w)) + I(cos(7 * w)) + I(sin(8 * w)) + I(cos(8 *
##   w)) + I(sin(9 * w)) + I(cos(9 * w)) + I(sin(10 * w)) + I(cos(10 *
##   w))
##   Resid. Df Resid. Dev Df Deviance Pr(>Chi)
## 1      333      651.10
## 2      331      399.58  2  251.523 < 2.2e-16 ***
## 3      329      396.03  2    3.546  0.169802
## 4      327      388.16  2    7.870  0.019547 *
## 5      325      377.96  2   10.208  0.006072 **
## 6      323      369.95  2    8.005  0.018269 *
## 7      321      368.57  2    1.383  0.500901
## 8      319      366.65  2    1.917  0.383492
## 9      317      363.37  2    3.279  0.194049
## 10     315      361.98  2    1.393  0.498338
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```

The big improvement is going from frequency 1 per day to 2 per day, but smaller improvements occur up to frequency 6 per day.

Alternative methods of model selection that avoid formal hypothesis tests use the information criteria AIC and BIC.

```

out.aic <- sapply(outs, AIC)
out.aic

## [1] 1783.204 1535.680 1536.134 1532.264 1526.056 1522.051 1524.668 1526.751
## [9] 1527.472 1530.079

which(out.aic == min(out.aic))

## [1] 6

out.bic <- sapply(outs, BIC)
out.bic

## [1] 1794.655 1554.766 1562.854 1566.618 1568.044 1571.673 1581.925 1591.642
## [9] 1599.997 1610.238

which(out.bic == min(out.bic))

## [1] 2

```

- BIC (Bayesian Information Criterion) is for when one believes that the true unknown probability model is in one of the statistical models considered — in this example that the regression function mapped to the linear predictor scale (by taking logs) has the form of a trigonometric series.
- AIC (Akaike Information Criterion) is for when one does not believe the true unknown probability model is in one of the statistical models considered.

## Confidence Intervals

Confidence intervals are very like those done above for binomial response. The only difference is that one uses the log likelihood for Poisson response rather than binomial response.


We will just show some Wald confidence bands.

```

conf.level <- 0.95
crit <- qnorm((1 + conf.level) / 2)
plot(hourofday, count, xlab = "hour of the day")
curve(predict(out, newdata = data.frame(w = x / 24 * 2 * pi),
  type="response"), add=TRUE)
sally <- function(x) {
  pout <- predict(out, type = "response", se.fit = TRUE,
    newdata = data.frame(w = x / 24 * 2 * pi))
  pout$fit - crit * pout$se.fit
}
curve(sally, add = TRUE, lty = "dashed")
sally <- function(x) {
  pout <- predict(out, type = "response", se.fit = TRUE,
    newdata = data.frame(w = x / 24 * 2 * pi))
  pout$fit + crit * pout$se.fit
}
curve(sally, add = TRUE, lty = "dashed")

```

\begin{figure}



{

}

\caption{95% Wald confidence bands, not simultaneous} \end{figure}