STATISTICS 538, Lecture #5

Deviance

November 8, 2010

A closer look at deviance

Write

$$2\{\max l(\theta_1, \ldots, \theta_n) - \max l(\beta)\} = \frac{D}{\phi}$$

D = Deviance
ϕ = Dispersion
D/ϕ = **scaled deviance**

Deviance **residuals**

Since deviance is like RSS...

Deviance residuals ex.

```
n <- 40; x <- rnorm(n); ni <- rep(5, n)

y <- rbinom(n, size=ni, prob=expit(-1 + 0.75*x))

ft <- glm(cbind(y,ni-y)~x, family=binomial)

plot(x, resid(ft, type="deviance"),ylim=4*c(-1,1))
plot(x, resid(ft,type="pearson"),ylim=4*c(-1,1))
```

Residuals versus Explanatory Variable

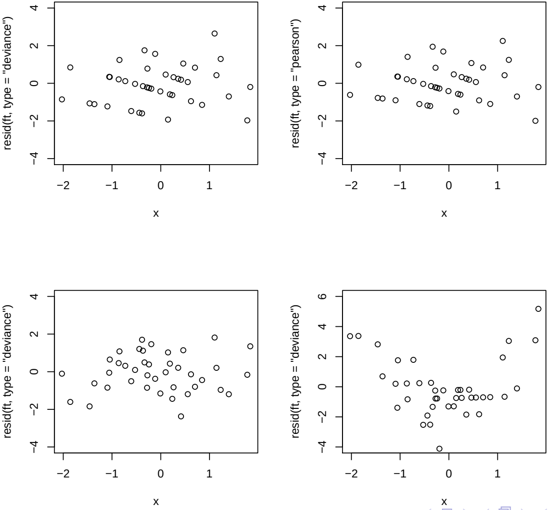

Deviance residuals ex., continued

```
ni <- rep(15, n)

y <- rbinom(n, size=ni, prob=expit(-1 + 0.75*x))

ft <- glm(cbind(y,ni-y)~x, family=binomial)
plot(x, resid(ft, type="deviance"),ylim=4*c(-1,1))

y <- rbinom(n, size=ni, prob=expit(-1 + 0.75*x^2))

ft <- glm(cbind(y,ni-y)~x, family=binomial)
plot(x, resid(ft, type="deviance"),ylim=c(-4,6))
```

Deviance similar to Pearson's Chi-squared statistic?

Dispersion unknown?

Rough intuition: estimate it to make the scaled deviance compatible with χ^2_{n-p}:

$$\hat{\phi} = \frac{D}{n-p}$$

Certainly works for normal models - applicability elsewhere???

Investigate the chi-sq approx. to sampling dist. of deviance . . . meaning of 'large-sample'?

```
### start with 5 trials at each of 5 design points
x <- rep(1:5, each=5)

set.seed(17); dvres <- matrix(0, 100, 2)

for (lp in 1:100) {
  y <- rbinom(length(x), size=1, prob=expit(-2+0.4*x))
  fit.a <- glm(y~x, family="binomial")
  fit.b <- glm(table(x,1-y)~unique(x), family="binomial")
  dvres[lp,] <- c(fit.a$dev, fit.b$dev)
}
```

E.g., for first simulated dataset

```
> summary(fit.a)
Coefficients:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)    5.410      2.172   2.490   0.0128 *
x             -1.280      0.545  -2.349   0.0188 *
---
Residual deviance: 20.198  on 23  degrees of freedom

> summary(fit.b)
Coefficients:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)    5.410      2.172   2.490   0.0128 *
unique(x)     -1.280      0.545  -2.349   0.0188 *
---
Residual deviance:  1.7335  on 3  degrees of freedom
```

Compare actual sampling dist. and χ^2 approx

```
hist(dvres[,1], prob=T,
     main="", xlab="Deviance", ylab="Density")
points(sort(dvres[,1]),
       dchisq(sort(dvres[,1]), length(x)-2), type="l")

hist(dvres[,2], prob=T,
  main="", xlab="Deviance", ylab="Density")
points(sort(dvres[,2]),
  dchisq(sort(dvres[,2]), length(unique(x))-2), type="l")
```

Consider more data in two different senses

```
### now 5 trials at each of twenty design points
x <- rep(seq(from=1,to=5,length=20), each=5)
...

### now 20 trials at each of 5 design points
x <- rep(1:5, each=20)
...
```

Simulation results