Inference in Binomial Logistic Regression

Tate Jacobson

Department of Statistics
Oregon State University, Corvallis, Oregon

Spring 2026

Reading: Ch 21 of *The Sleuth*

Outline

1. Inference in Binomial Logistic Regression
   - Inference for Regression Coefficients
   - Deviance Goodness-of-Fit

Wald's Test and CIs

As in binary logistic regression, *the coefficient estimates* $\hat{\beta}_j$ *for binomial logistic regression are the MLEs*.

Therefore

$$Z = \frac{\hat{\beta}_j - \beta_j}{SE(\hat{\beta}_j)}$$

is approximately standard normal for large sample sizes and we can use it to perform Wald's test and construct CIs for the regression coefficients.

Drop in Deviance Test

Likewise, we can use the drop-in-deviance test to compare two nested binomial logistic regression models.

As in binary logistic regression, the test statistic is

$$LRT = \text{deviance}_{reduced} - \text{deviance}_{full}$$

and it has an approximate χ^2_d sampling distribution where $d =$ the difference in the number of parameters between the reduced and full models.

Inference for Krunnit Island extinction model

Let's return to our binomial logistic regression model for extinction based on the Krunnit Islands data with log(Area) as our explanatory variable:

$$\text{logit}(\pi) = \beta_0 + \beta_1 \log(\mathit{Area})$$

where π denotes the probability of extinction.

Inference based on the Krunnit Island R output

```
Coefficients:
            Estimate Std. Error z value Pr(>|z|)
(Intercept) -1.19620    0.11845 -10.099  < 2e-16 ***
larea       -0.29710    0.05485     ???      ???
---

(Dispersion parameter for binomial family taken to be 1)

Null deviance: 45.338  on 17  degrees of freedom
Residual deviance: 12.062  on 16  degrees of freedom
AIC: 75.394

Number of Fisher Scoring iterations: 4
```

Practice (5 minutes): Use this model output to

- Test $H_0 : \beta_1 = 0$ vs $H_a : \beta_1 \neq 0$ using Wald's test
- Construct a 95% CI for β_1.
- Use the drop-in-deviance test to compare this model with the null model (*Hint:* The "Null deviance" in the output is the deviance for the null model.)

Assessing Goodness-of-Fit

Suppose you have we have a Binomial Logistic regression model with p explanatory variables with observations $i = 1, \ldots, n$:

$$\text{logit}(\pi_i) = \beta_0 + \beta_1 X_{i1} + \cdots + \beta_p X_{ip}$$

We would like to be able to measure **how well this model fits the data.**

The Deviance Goodness-of-Fit Test

The idea behind the **Deviance Goodness-of-Fit Test** is to compare the model of interest

$$\operatorname{logit}(\pi_i) = \beta_0 + \beta_1 X_{i1} + \cdots + \beta_p X_{ip} \tag{1}$$

with the *saturated model*, which has one parameter for each observation:

$$\operatorname{logit}(\pi_i) = \alpha_i \tag{2}$$

for $i = 1, \ldots, n$.

Informally, *the null hypothesis is that the model of interest is adequate* and the alternative hypothesis is that more structure is needed to fit the data adequately.

The Deviance Goodness-of-Fit Test

The test statistic is the *drop in deviance test statistic for comparing the model of interest and the saturated model*:

$$D = \text{deviance}_{\text{model of interest}} - \text{deviance}_{saturated}.$$

The saturated model has n parameters and the model of interest has $p + 1$ parameters.
Q: What is the sampling distribution of D? (Hint: Remember the drop in deviance test)

Details of The Deviance Goodness-of-Fit Test

It turns out that the deviance from the saturated model (2) is zero, so our test statistic reduces to

$$D = \text{deviance}_{\text{model of interest}} - 0 \sim \chi^2_{n-p-1}.$$

The Krunnit Archipelago Example

From the R output from the logistic regression:

```
Null deviance: 45.338  on 17  degrees of freedom
Residual deviance: 12.062  on 16  degrees of freedom
```

In the R output:

- the residual deviance = the deviance for the model
- the residual degrees of freedom = $n - p - 1$

Exercise (2 minutes): perform the deviance goodness of fit test for our model for extinction based on $\log(Area)$ using this output.

- The deviance goodness-of-fit statistic is
- its sampling distribution is
- therefore the p-value is

What's our conclusion?

Caveats for the Deviance Goodness of Fit Test

1. As a general rule, the deviance goodness of fit test is not reliable when a substantial proportion of the binomial sizes m_i are less than 5. (In the Krunnit Islands data, the m_i's range from 6 to 75, so we're good.)
2. A small p-value for the deviance goodness of fit test can mean several things:
   1. the model for the mean is incorrect (e.g. more explanatory variables are needed)
   2. the binomial model is inadequate to characterize the distribution of response
   3. there are a few severely outlying observations

   For this reason, it can be difficult to draw strong conclusions based solely on the deviance goodness of fit test. Further exploration is needed.

Let’s head over to R to run through all of these inference procedures in the Krunnit Islands example.

Recap

Material covered: Ch 21 of *The Sleuth*

- Wald's test and CIs
- Drop in deviance test
- Deviance goodness of fit test