A Generalized Linear Model for Poisson Response Data

An Example Experiment

Consider an experiment designed to evaluate the effectiveness of an anti-fungal chemical on plants.

A total of 60 plant leaves were randomly assigned to treatment with 0, 5, 10, 15, 20, or 25 units of the anti-fungal chemical, with 10 plant leaves for each amount of anti-fungal chemical.

All leaves were infected with a fungus.

Following a two-week period, the leaves were studied under a microscope, and the number of infected cells was counted and recorded for each leaf.

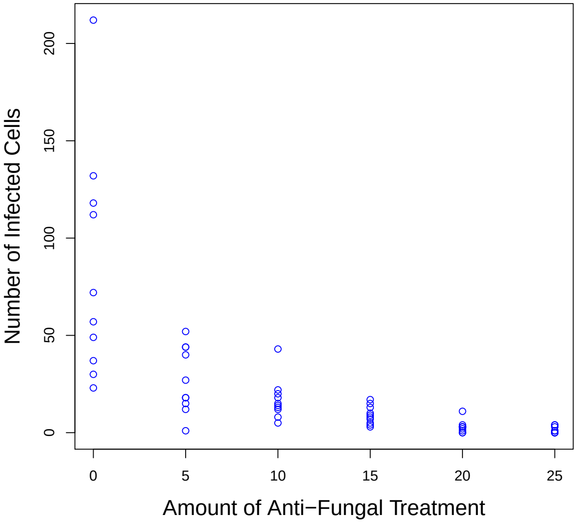
Number of Infected Cells
0
50
100
150
200
0
5
10
15
20
25
Amount of Anti−Fungal Treatment

Questions of Interest

Is there a significant association between the amount of anti-fungal chemical and the number of infected cells per plant leaf?

Can we estimate the mean number of infected cells per leaf as a function of the amount of anti-fungal chemical?

What amount of chemical should be applied to leaves so that the mean number of infected cells is approximately 5 per leaf?

A Generalized Linear Model for Poisson Count Data

For all $i = 1, ..., n$,

$$y_i \sim \text{Poisson}(\lambda_i),$$

$$\log(\lambda_i) = \boldsymbol{x}_i' \boldsymbol{\beta},$$

and $y_1, ..., y_n$ are independent.

The Poisson Distribution

Recall that for $y_i \sim \text{Poisson}(\lambda_i)$, the probability mass function of y_i is

$$P(y_i = y) = \begin{cases} \frac{\lambda_i^y \exp(-\lambda_i)}{y!} & \text{for } y \in \{0, 1, 2, \ldots\} \\ 0 & \text{otherwise} \end{cases},$$

$$E(y_i) = \lambda_i, \quad \text{and} \quad \text{Var}(y_i) = \lambda_i.$$

The Poisson Log Likelihood for the Generalized LM

The Poisson log likelihood function is

$$\begin{aligned}\ell(\boldsymbol{\beta} \mid \boldsymbol{y}) &= \sum_{i=1}^{n} [y_i \log(\lambda_i) - \lambda_i - \log(y_i!)] \\ &= \sum_{i=1}^{n} [y_i \, \boldsymbol{x}_i' \boldsymbol{\beta} - \exp(\boldsymbol{x}_i' \boldsymbol{\beta}) - \log(y_i!)].\end{aligned}$$

Inference for the Poisson Generalized LM

The function $\ell(\boldsymbol{\beta} \mid \boldsymbol{y})$ can be maximized over $\boldsymbol{\beta} \in \mathbb{R}^p$ using Fisher's scoring method to obtain an MLE $\hat{\boldsymbol{\beta}}$.

An estimate of the inverse Fisher information matrix can be used for Wald inference concerning $\boldsymbol{\beta}$.

Alternatively, we can conduct likelihood ratio based inference of reduced vs. full models.

Interpretation of Parameters

Let $\boldsymbol{x} = [x_1, x_2, \ldots, x_{j-1}, x_j \qquad , x_{j+1}, \ldots, x_p]'$.

Let $\tilde{\boldsymbol{x}} = [x_1, x_2, \ldots, x_{j-1}, x_j + 1, x_{j+1}, \ldots, x_p]'$.

In other words, $\tilde{\boldsymbol{x}}$ is the same as $\boldsymbol{x}$ except that the jth explanatory variable has been increased by one unit.

Let $\lambda = \exp(\boldsymbol{x}'\boldsymbol{\beta})$ and $\tilde{\lambda} = \exp(\tilde{\boldsymbol{x}}'\boldsymbol{\beta})$.

Ratio of Means

$$
\begin{aligned}
\tilde{\lambda}/\lambda &= \exp(\tilde{\boldsymbol{x}}'\boldsymbol{\beta})/\exp(\boldsymbol{x}'\boldsymbol{\beta}) \\
&= \exp(\tilde{\boldsymbol{x}}'\boldsymbol{\beta} - \boldsymbol{x}'\boldsymbol{\beta}\} \\
&= \exp\{(x_j + 1)\beta_j - x_j\beta_j\} \\
&= \exp(\beta_j).
\end{aligned}
$$

Thus, $\tilde{\lambda} = \exp(\beta_j)\lambda$.

All other explanatory variables held constant, the mean response at $x_j + 1$ is $\exp(\beta_j)$ times the mean response at x_j.

This is true regardless of the initial value x_j.

A one unit increase in the jth explanatory variable (with all other explanatory variables held constant) is associated with a multiplicative change in the mean response by the factor $\exp(\beta_j)$.

If (L_j, U_j) is a $100(1-\alpha)\%$ confidence interval for β_j, then

$$(\exp(L_j), \exp(U_j))$$

is a $100(1-\alpha)\%$ confidence interval for $\exp(\beta_j)$.

Also, note that if (L,U) is a $100(1-\alpha)\%$ confidence interval for $\boldsymbol{x}'\boldsymbol{\beta}$, then a $100(1-\alpha)\%$ confidence interval for $\lambda = \exp(\boldsymbol{x}'\boldsymbol{\beta})$ is

$$(\exp(L), \exp(U))\,.$$

A Generalized LM for the Anti-Fungal Experiment

For $i = 1, ..., 60$, let y_i denote the number of infected cells for leaf i. Suppose $y_1, \ldots, y_{60}$ are independent,

$$y_i \sim \text{Poisson}(\lambda_i), \text{ and}$$

$$\log(\lambda_i) = \beta_0 + \beta_1 x_i,$$

where x_i denotes the amount of anti-fungal chemical applied to leaf i.

$$y_i \overset{ind}{\sim} \text{Poisson}(\lambda_i)$$

$$\log(\lambda_i) = \beta_0 + \beta_1 x_i$$

```
> o=glm(y~x,family=poisson(link = "log"))
> summary(o)

Call:
glm(formula = y ~ x, family = poisson(link = "log"))

Deviance Residuals:
    Min       1Q   Median       3Q      Max
-7.8363  -1.6715  -0.3411   1.2467  12.3127
```

PRETTY EXTREME RESIDUALS

WE WILL IGNORE PROBLEMS WITH FIT FOR NOW.

SQUARE ROOT OF
DIAGONAL OF

$\widehat{I^{-1}(\underset{\sim}{\beta})}$

WALD TESTS STATS
AND p-VALUES

$\hat{\underset{\sim}{\beta}}$

$\hat{\beta}_0$

$\hat{\beta}_1$

```
Coefficients:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)  4.373003   0.032434  134.83   <2e-16 ***
x           -0.162011   0.004933  -32.84   <2e-16 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

(Dispersion parameter for poisson family taken to be 1)
```

$2\hat{\ell}_s - 2\hat{\ell}_0$ $n-1$

```
    Null deviance: 2314.87  on 59  degrees of freedom
Residual deviance:  614.65  on 58  degrees of freedom
AIC: 849.53
```

$2\hat{\ell}_s - 2\hat{\ell}$ $n-2$

```
Number of Fisher Scoring iterations: 5
```

```
> #MLE of beta vector
> b=coef(o)
> b
(Intercept)           x
  4.3730032  -0.1620111
```

$\hat{\beta}_0$ $\hat{\beta}_1$

```
> #Estimated variance of the MLE
> v=vcov(o)
> v
              (Intercept)             x
(Intercept)  1.051970e-03 -9.179444e-05
x           -9.179444e-05  2.433540e-05
```

$\widehat{I^{-1}(\underset{\sim}{\beta})}$

Answering the First Question of Interest

Is there a significant association between the amount of anti-fungal chemical and the number of infected cells per plant leaf?

The null hypothesis of no association between the amount of anti-fungal chemical and the number of infected cells per plant leaf is $H_0 : \beta_1 = 0$ for our Generalized LM.

We can test this null hypothesis using either a Wald test or a LRT as follows.

Wald Test of $H_0 : \beta_1 = 0$

```
> #Test Statistic
> z=b[2]/sqrt(v[2,2])
> z
        x
-32.84169

> #p-value
> 2*pnorm(z)
                x
```

~~1.496932e-236~~ ≈ 0

$$z = \frac{\hat{\beta}_1}{SE(\hat{\beta}_1)}$$

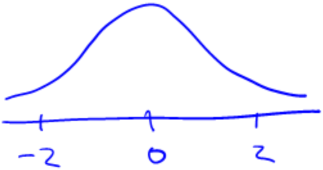

Likelihood Ratio Test of $H_0 : \beta_1 = 0$

```
> anova(o,test="Chisq")
Analysis of Deviance Table

Model: poisson, link: log

Response: y

Terms added sequentially (first to last)

     Df Deviance Resid. Df Resid. Dev  Pr(>Chi)
NULL                    59    2314.87
x     1   1700.2        58     614.65 < 2.2e-16 ***
```

$2-1$

$2\hat{\ell} - 2\hat{\ell}_0$

$2\hat{\ell}_s - 2\hat{\ell}_0$

$n-1$

LRT p-VALUE

$n-2$

$2\hat{\ell}_s - 2\hat{\ell}$

Answering the Second Question of Interest

$\log(\lambda) = \beta_0 + \beta_1 x \Rightarrow \lambda = \exp(\beta_0 + \beta_1 x)$

Can we estimate the mean number of infected cells per leaf as a function of the amount of anti-fungal chemical?

According to our Poisson Generalized LM, the mean number of infected cells for a leaf treated with x units of the anti-fungal chemical is

$$\exp(\beta_0 + \beta_1 x), \text{ which is estimated by } \exp(\hat{\beta}_0 + \hat{\beta}_1 x).$$

We can add the estimated mean function to our scatterplot with the following code:

```
xgrid=seq(0,25,by=.1)
lines(xgrid,exp(b[1]+b[2]*xgrid),col=2,lwd=1.5)
```

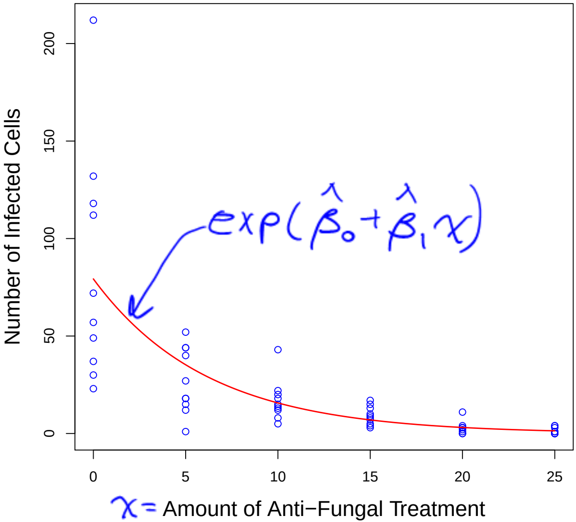
$\exp(\hat{\beta}_0 + \hat{\beta}_1 x)$
Number of Infected Cells
0
50
100
150
200
0
5
10
15
20
25
x = Amount of Anti−Fungal Treatment

Example Confidence Interval for a Mean

An approximate 95% Wald confidence interval for the mean number of infected cells when 15 units of the chemical are applied is $(6.17, 7.89)$.

$$\underset{\sim}{c} = \begin{bmatrix} 1 \\ 15 \end{bmatrix} \qquad \underset{\sim}{c}'\underset{\sim}{\beta} = \beta_0 + \beta_1(15)$$

```
> cc=c(1,15)
> se=sqrt(t(cc)%*%v%*%cc)
> exp(t(cc)%*%b-2*se)
          [,1]
[1,] 6.171721
> exp(t(cc)%*%b+2*se)
         [,1]
[1,] 7.89079
```

$$SE = \sqrt{\underset{\sim}{c}'\,\widehat{I^{-1}(\underset{\sim}{\beta})}\,\underset{\sim}{c}}$$

$$\exp\left(\underset{\sim}{c}'\hat{\underset{\sim}{\beta}} - 2\,SE\right)$$

$$\exp\left(\underset{\sim}{c}'\hat{\underset{\sim}{\beta}} + 2\,SE\right)$$

The Effect of a One Unit Increase in Chemical

```
> exp(b[2])
        x
0.8504318
```

$\exp(\hat{\beta}_1)$

An increase by one unit in the amount of anti-fungal chemical applied to a leaf is associated with a decrease of approximately 15% in the mean number of infected cells.

$100(1-0.85)=15$

Confidence Interval for the Multiplicative Factor

$$\left(e^{\hat{\beta}_1 - 2SE(\hat{\beta}_1)} , e^{\hat{\beta}_1 + 2SE(\hat{\beta}_1)} \right)$$

```
> c(exp(b[2]-2*sqrt(v[2,2])),exp(b[2]+2*sqrt(v[2,2])))
        x         x
(0.8420825, 0.8588638)
```

An approximated 95% confidence interval for the mean associated with $x + 1$ units of chemical divided by the mean associated with x units of the chemical ($\tilde{\lambda}/\lambda$) is $(0.842, 0.859)$.

Answering the Third Question of Interest

What amount of chemical should be applied to leaves so that the mean number of infected cells is approximately 5 per leaf?

$$\begin{aligned}\lambda = 5 &\iff \log(\lambda) = \log(5)\\ &\iff \boldsymbol{x}'\boldsymbol{\beta} = \log(5)\\ &\iff \beta_0 + \beta_1 x = \log(5)\\ &\iff x = (\log(5) - \beta_0)/\beta_1\end{aligned}$$

Thus, we seek an estimate of the nonlinear function of $\boldsymbol{\beta}$ given by

$$h(\boldsymbol{\beta}) = (\log(5) - \beta_0)/\beta_1.$$

Answering the Third Question of Interest

By the invariance property, the MLE of $h(\boldsymbol{\beta})$ is

$$\widehat{h(\boldsymbol{\beta})} = h(\hat{\boldsymbol{\beta}}) = (\log(5) - \hat{\beta}_0)/\hat{\beta}_1.$$

By the Delta Method, $\widehat{h(\boldsymbol{\beta})}$ is approximately normal with mean $h(\boldsymbol{\beta})$ and variance

$$\widehat{\text{Var}}(\widehat{h(\boldsymbol{\beta})}) = \hat{\boldsymbol{D}}'\widehat{\text{Var}}(\hat{\boldsymbol{\beta}})\hat{\boldsymbol{D}},$$

where

$$\hat{\boldsymbol{D}}' = \left[\left.\frac{\partial h(\boldsymbol{\beta})}{\partial \beta_0}\right|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}}, \left.\frac{\partial h(\boldsymbol{\beta})}{\partial \beta_1}\right|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}} \right].$$

Answering the Third Question of Interest

$$h(\underset{\sim}{\beta}) = \frac{\log(5) - \beta_0}{\beta_1}$$

$$\begin{aligned}\hat{\boldsymbol{D}}' &= \left[\left.\frac{\partial h(\boldsymbol{\beta})}{\partial \beta_0}\right|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}}, \left.\frac{\partial h(\boldsymbol{\beta})}{\partial \beta_1}\right|_{\boldsymbol{\beta}=\hat{\boldsymbol{\beta}}} \right] \\ &= \left[\frac{-1}{\hat{\beta}_1}, \frac{\hat{\beta}_0 - \log(5)}{\hat{\beta}_1^2} \right]\end{aligned}$$

```
> h=(log(5)-b[1])/b[2]
> h
(Intercept)
   17.05788
>
> Dhat=c(-1/b[2],(b[1]-log(5))/b[2]^2)
>
> seh=sqrt(t(Dhat)%*%v%*%Dhat)
>
> ci=c(h-2*seh,h+2*seh)
> ci
[1] 16.18486 17.93090
```

$h(\hat{\underset{\sim}{\beta}})$

$$\hat{D} = \begin{bmatrix} -1/\hat{\beta}_1 \\ \dfrac{\hat{\beta}_0 - \log(5)}{\hat{\beta}_1^2} \end{bmatrix}$$

$$\sqrt{\hat{D}' \, \widehat{I^{-1}(\beta)} \, \hat{D}} = SE$$

$$h(\hat{\underset{\sim}{\beta}}) - 2SE, \; h(\hat{\underset{\sim}{\beta}}) + 2SE$$

Answering the Third Question of Interest

The estimated amount of chemical that should be applied to leaves so that the mean number of infected cells is approximately 5 per leaf is 17.1 units.

An approximate 95% confidence interval for the required amount of chemical is 16.2 to 17.9 units.

```
plot(x,y,xlab="Level of Anti-Fungal Treatment",
     ylab="Number of Infected Cells",col=4,cex.lab=1.5,
     xlim=c(14,21),ylim=c(0,18))
xgrid=seq(0,25,by=.1)
lines(xgrid,exp(b[1]+b[2]*xgrid),col=2,lwd=1.5)
abline(h=5,lty=2)
lines(c(h,h),c(-1,5),lwd=1.5)
lines(c(ci[1],ci[1]),c(-1,5),lwd=1.5,col="purple")
lines(c(ci[2],ci[2]),c(-1,5),lwd=1.5,col="purple")
```

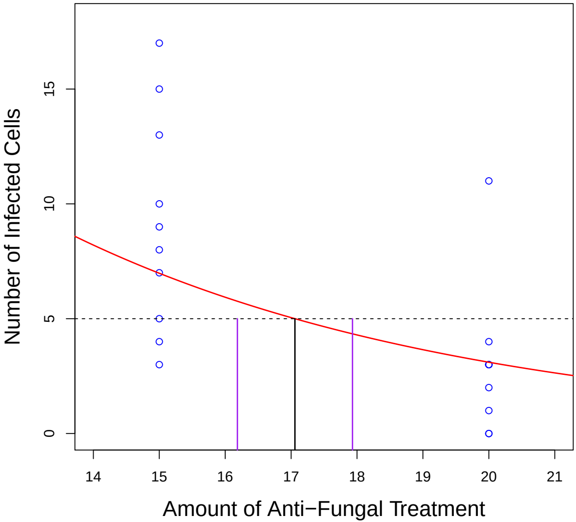
Number of Infected Cells
0
5
10
15
14
15
16
17
18
19
20
21
Amount of Anti−Fungal Treatment

Checking for Lack of Fit

We can compare the fit of a Poisson Generalized LM to the fit of a *saturated model*.

The saturated model uses one parameter for each observation.

In this case, the saturated model has one free parameter λ_i for each y_i.

Poisson Generalized LM	Saturated Model
$y_i \sim \text{Poisson}(\lambda_i)$	$y_i \sim \text{Poisson}(\lambda_i)$
$y_1, \ldots, y_n$ independent	$y_1, \ldots, y_n$ independent
$\lambda_i = \exp(\boldsymbol{x}_i'\boldsymbol{\beta})$ for some $\boldsymbol{\beta} \in \mathbb{R}^p$	$\lambda_i \in [0, \infty)$ for $i = 1, \ldots, n$ with no other restrictions
p parameters	n parameters

For all $i = 1, \ldots, n$,

the MLE of λ_i under the Poisson Generalized LM is

$$\hat{\lambda}_i = \exp(\boldsymbol{x}_i' \hat{\boldsymbol{\beta}}),$$

and the MLE of λ_i under the saturated model is simply y_i.

Then the likelihood ratio statistic for testing the Poisson Generalized LM as the reduced model vs. the saturated model as the full model is

$$2 \sum_{i=1}^{n} \left[y_i \ \log\left(\frac{y_i}{\hat{\lambda}_i} \right) - (y_i - \hat{\lambda}_i) \right].$$

This is the *Deviance Statistic* for the Poisson case.

A Lack-of-Fit Test

When the Poisson means $\lambda_1, \ldots, \lambda_n$ are suitably large, the Deviance Statistic is approximately χ^2_{n-p} if the Poisson Generalized LM is correct.

Thus, the Deviance Statistic can be compared to the χ^2_{n-p} distribution to test for lack of fit of the logistic regression model.

As indicated in Chapter 22 of *The Statistical Sleuth*, the χ^2_{n-p} approximation may not be very reliable if a substantial proportion of the $\hat{\lambda}_i$ values are less than 5.

Deviance Residuals

In the Poisson case, the deviance residuals are given by

$$d_i \equiv \text{sign}(y_i - \hat{\lambda}_i)\sqrt{2\left[y_i\,\log\left(\frac{y_i}{\hat{\lambda}_i}\right) - (y_i - \hat{\lambda}_i)\right]}.$$

The Deviance Statistic is the sum of the squared deviance residuals ($\sum_{i=1}^{n} d_i^2$).

Pearson's Chi-Square Statistic

Another lack of fit statistic that is approximately χ^2_{n-p} under the null is Pearson's Chi-Square Statistic:

$$\begin{aligned} X^2 &= \sum_{i=1}^{n} \left(\frac{y_i - \widehat{E}(y_i)}{\sqrt{\widehat{\text{Var}}(y_i)}} \right)^2 \\ &= \sum_{i=1}^{n} \left(\frac{y_i - \hat{\lambda}_i}{\sqrt{\hat{\lambda}_i}} \right)^2 . \end{aligned}$$

Pearson Residuals

The term

$$r_i = \frac{y_i - \hat{\lambda}_i}{\sqrt{\hat{\lambda}_i}}$$

is a *Pearson residual* for a Poisson Generalized LM.

The Pearson statistic is the sum of the squared Pearson residuals ($\sum_{i=1}^{n} r_i^2$).

Residual Diagnostics

For large λ_i, both d_i and r_i should be approximately distributed as standard normal random variables if the Poisson Generalized LM is correct.

Thus, either set of residuals can be used to diagnose problems with model fit by, e.g., identifying outlying observations.

R Code to Generate a Deviance Residual Plot

```
d=resid(o,type="deviance")
plot(fitted(o),d,
     xlab="Estimated Mean",
     ylab="Deviance Residual",
     cex.lab=1.4,
     pch=16,col=4)
abline(h=0,lty=2)
```

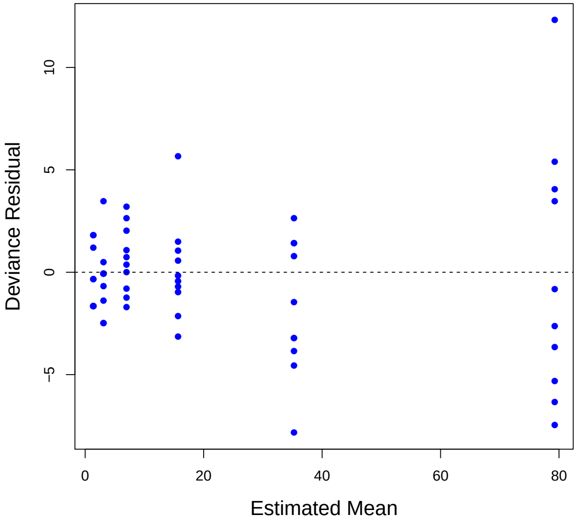
Deviance Residual
−5
0
5
10
0
20
40
60
80
Estimated Mean

R Code to Generate a Pearson Residual Plot

```
r=resid(o,type="pearson")
plot(fitted(o),r,
     xlab="Estimated Mean",
     ylab="Pearson Residual",
     cex.lab=1.4,
     pch=16,col=4)
abline(h=0,lty=2)
```

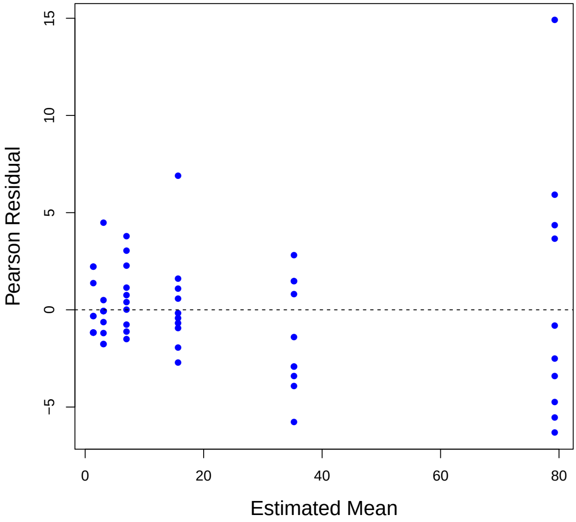
Pearson Residual
Estimated Mean
15
10
5
0
−5
0
20
40
60
80

Overdispersion

In the Generalized Linear Models framework, its often the case that $\mathrm{Var}(y)$ is a function of $E(y)$.

For the Poisson case,

$$\mathrm{Var}(y) = E(y) = \lambda.$$

Thus, when we fit a Poisson Generalized LM and obtain estimates of the mean of the response, we get estimates of the variance of the response as well.

If the variability of our response is greater than we should expect based on our estimates of the mean, we say that there is *overdispersion*.

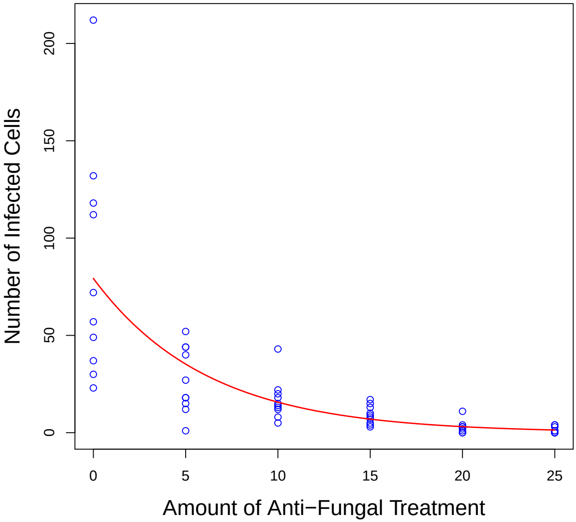
Number of Infected Cells
0
50
100
150
200
0
5
10
15
20
25
Amount of Anti−Fungal Treatment

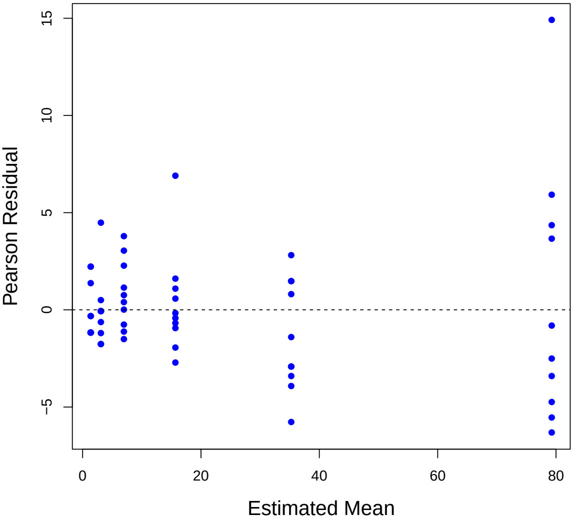
Pearson Residual
15
10
5
0
−5
0
20
40
60
80
Estimated Mean

If either the Deviance Statistic or the Pearson Chi-Square Statistic suggests a lack of fit that cannot be explained by other reasons (e.g., poor model for the mean or a few extreme outliers), overdispersion may be the problem.

Evidence for Overdispersion

$$2\hat{\ell}_s - 2\hat{\ell}$$

```
> deviance(o)
[1] 614.6479
> sum(d^2)
[1] 614.6479
> sum(r^2)
[1] 639.797
> 1-pchisq(deviance(o),60-2)
[1] 0
> 1-pchisq(sum(r^2),60-2)
[1] 0
```

$\sum_{i=1}^{n} d_i^2$

$\sum_{i=1}^{n} r_i^2$

χ^2_{58}

Quasi-Likelihood (QL) Inference

If there is overdispersion, a *quasi-likelihood* approach may be used.

In the Poisson case, we make all the same assumptions as before except that we assume

$$\mathrm{Var}(y_i) = \phi\lambda_i$$

for some unknown dispersion parameter $\phi > 1$.

The dispersion parameter ϕ can be estimated by

$$\hat{\phi} = \frac{\sum_{i=1}^{n} d_i^2}{n-p} = \frac{2\hat{\ell}_s - 2\hat{\ell}}{n-p}$$

or

$$\hat{\phi} = \frac{\sum\limits_{i=1}^{n} r_i^2}{n-p}.$$

All analyses are as before except that

1. The estimated variance of $\hat{\beta}$ is multiplied by $\hat{\phi}$.
2. For Wald type inferences, the standard normal null distribution is replaced by t with $n - p$ degrees of freedom.
3. Any test statistic T that was assumed χ^2_q under H_0 is replaced with $T/(q\hat{\phi})$ and compared to an F distribution with q and $n - p$ degrees of freedom.

These changes to the inference strategy in the presence of overdispersion are analogous to the changes that would take place in normal theory Gauss-Markov linear model analysis if we switched from assuming σ^2 were known to be 1 to assuming σ^2 were unknown and estimating it with MSE.
(Here ϕ is like σ^2 and $\hat{\phi}$ is like MSE.)

Whether there is overdispersion or not, all the usual ways of conducting generalized linear models inference are approximate except for the special case of normal theory linear models.

QL Analysis of the Fungal Infection Data

```
> #Estimates of the dispersion parameter
>
> deviance(o)/df.residual(o)
[1] 10.59738
>
> sum(r^2)/df.residual(o)
[1] 11.03098
```

$\hat{\phi}$

```
> oq=glm(y~x,family=quasipoisson(link = "log"))
> summary(oq)

Call:
glm(formula = y ~ x, family = quasipoisson(link = "log"))

Deviance Residuals:
    Min       1Q   Median       3Q      Max
-7.8363  -1.6715  -0.3411   1.2467  12.3127
```

Square root of diagonal of $\hat{\phi}\,\widehat{I^{-1}(\hat{\boldsymbol{\beta}})}$

$\hat{\boldsymbol{\beta}}$ SAME

```
Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)  4.37300    0.10772  40.595  < 2e-16 ***
x           -0.16201    0.01638  -9.888  4.7e-14 ***
---
Signif. codes:  0 *** 0.001 ** 0.01 * 0.05 . 0.1   1

(Dispersion parameter for quasipoisson family
 taken to be 11.03098)
```

$$\hat{\phi} = \frac{\sum_{i=1}^{n} r_i^2}{n-p} = \frac{\sum_{i=1}^{60} r_i^2}{58}$$

SAME AS BEFORE

```
    Null deviance: 2314.87  on 59  degrees of freedom
Residual deviance:  614.65  on 58  degrees of freedom
AIC: NA

Number of Fisher Scoring iterations: 5
```

```
> #MLE of beta vector
> b=coef(oq)
> b
(Intercept)           x
  4.3730032  -0.1620111
>
> #Estimated variance of the MLE
> v=vcov(oq)
> v
             (Intercept)             x
(Intercept)  0.011604262 -0.0010125829
x           -0.001012583  0.0002684434
```

$$\hat{\phi}\,\widehat{I^{-1}(\underset{\sim}{\hat{\beta}})} = \widehat{VAR}(\underset{\sim}{\hat{\beta}})$$

```
> #Test Statistic
> tstat=b[2]/sqrt(v[2,2])
> tstat
        x
-9.888225
>
> #p-value
> 2*pt(tstat,60-2)
           x
4.700095e-14
```

≈ 0

$$t = \frac{\hat{\beta}_1}{\sqrt{\hat{\phi}\left[\widehat{I^{-1}(\beta)}\right]_{11}}}$$

SE($\hat{\beta}_1$) BY QUASILIKELIHOOD APPROACH

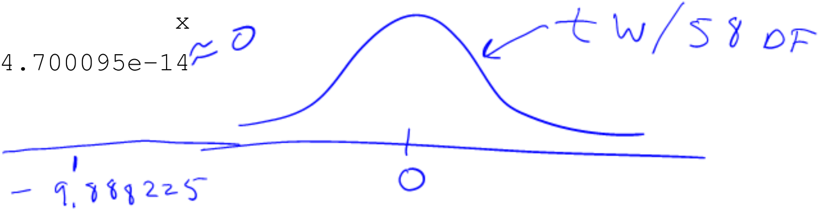

```
> #Likelihood Ratio Test
> anova(oq,test="F")
Analysis of Deviance Table

Model: quasipoisson, link: log

Response: y

Terms added sequentially (first to last)


     Df Deviance Resid. Df Resid. Dev      F    Pr(>F)
NULL                    59    2314.87
x     1   1700.2        58     614.65 154.13 < 2.2e-16 ***
```

$$F = \frac{(2\hat{\ell} - 2\hat{\ell}_0)/(2-1)}{\sum_{i=1}^{60} r_i^2/(60-2)}$$

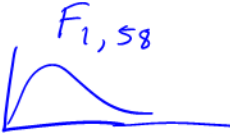

$2\hat{\ell}_s - 2\hat{\ell}_0$

$2\hat{\ell} - 2\hat{\ell}_0$

$p\text{-value} \approx 0$

$2\hat{\ell}_s - 2\hat{\ell}$

Exponential Dispersion Families

The normal, Bernoulli, binomial, and Poisson families of distributions can each be characterized as an *exponential dispersion family* of distributions with a probability density function or probability mass function of the form

$$\exp\left\{\frac{y\theta - b(\theta)}{a(\phi)} + c(y, \phi)\right\},$$

where $a(\cdot)$, $b(\cdot)$, and $c(\cdot)$ are specific functions, θ is an unknown parameter that is a function of the mean, and ϕ is a dispersion parameter that may or may not be known.

Exponential Dispersion Families

The gamma and inverse Gaussian families of distributions are other example distributions sometimes useful for modeling data that can be characterized as exponential dispersion families.

Generalized Linear Models for gamma and inverse Gaussian distributions will be discussed in STAT520 – along with more details about normal, Bernoulli, binomial, and Poisson Generalized Linear Models.