Poisson Generalized Linear Model

Filippo Gambarota

University of Padova

2022/2023

Updated on 2023-05-09

Outline

Poisson GLM

Everything that we discussed for the binomial GLM is also relevant for the Poisson GLM. We are gonna focus on specificity of the Poisson model in particular:

- **Poisson distribution** and **link function**
- **Parameters interpretation**
- **Overdispersion** causes, consequences and remedies

Poisson distribution

Poisson distribution

The Poisson distribution is defined as:

$$p(y; \lambda) = \frac{\lambda^y e^{-\lambda}}{y!}$$

Where the mean is λ and the variance is λ

Poisson distribution

y ~ Poisson($\lambda = 5$)

Poisson distribution

As the mean increases also the variance increase and the distributions is approximately normal:

Poisson distribution

$y = 0.69 + 2.71*x$

Link function

The common (and default in R) link function ($g(\lambda)$) for the Poisson distribution is the **log** link function and the inverse of link function is the exponential.

$$log(\lambda) = \beta_0 + \beta_1 X_1 + ...\beta_p X_p$$
$$\lambda = e^{\beta_0 + \beta_1 X_1 + ...\beta_p X_p}$$

Parameters intepretation

An example...

Let's start by a simple example trying to explain the number of errors of math exercises by students ($N = 100$) as a function of the number of hours of study.

id	studyh	solved
1	54	36
2	48	36
3	8	12
4	9	15
...	...	...
97	100	78
98	73	46
99	85	55
100	45	26

An example...

- There is a clear non-linear and positive relationship
- The both the mean and variance increase as a function of the predictor

Model fitting

We can fit the model using the `glm` function in R setting the appropriate **random component** (`family`) and the **link function** (`link`):

```
fit <- glm(solved ~ studyh, family = poisson(link = "log"), data = dat)

summary(fit)
```
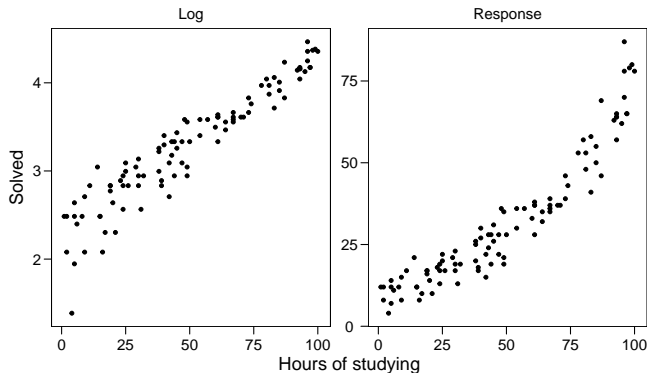
```
##
## Call:
## glm(formula = solved ~ studyh, family = poisson(link = "log"),
##     data = dat)
##
## Deviance Residuals:
##      Min        1Q    Median        3Q       Max
## -2.42031  -0.71240   0.01224   0.65541   2.02636
##
## Coefficients:
##              Estimate Std. Error z value Pr(>|z|)
## (Intercept) 2.3140977  0.0456575   50.68   <2e-16 ***
## studyh      0.0200667  0.0006341   31.64   <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: 1166.213  on 99  degrees of freedom
## Residual deviance:   82.262  on 98  degrees of freedom
## AIC: 599.4
##
## Number of Fisher Scoring iterations: 4
```

Parameters intepretation

- The (Intercept) 2.314 is the log of the expected number of solved exercises for a student with 0 hours of studying. Taking the exponential we obtain the estimation on the response scale 10.116
- the studyh 0.020 is the increase in the expected increase of (log) solved exercises for a unit increase in hours of studying. Taking the exponential we obtain the ratio of increase of the number of solved exercises 1.020

Parameters intepretation

Again, as in the binomial model, the effects are linear on the log scale but non-linear on the response scale.

Parameters intepretation

The non-linearity can be easily seen using the `predict()` function:

```r
linear <- predict(fit, newdata = data.frame(studyh = c(10, 11)))
diff(linear) # same as the beta0
```

```
##          2
## 0.02006669
```

```r
nonlinear <- exp(linear) # or predict(..., type = "response")
diff(nonlinear)
```

```
##         2
## 0.2506045
```

```r
# ratio of increase when using the response scale
nonlinear[2]/nonlinear[1]
```

```
##        2
## 1.020269
```

```r
# equivalent to exp(beta1)
exp(coef(fit)[2])
```
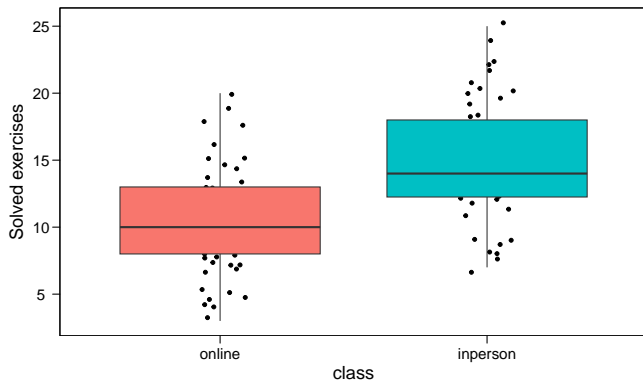
```
##    studyh
## 1.020269
```

Parameters intepretation - Categorical variable

Let's make a similar example with the number of solved exercises comparing students who attended online classes and students attending in person.

id	class	class.c	solved
1	online	0	3
2	inperson	1	12
3	online	0	19
4	inperson	1	15
...	...	...	...
97	online	0	8
98	inperson	1	15
99	online	0	11
100	inperson	1	14

Parameters intepretation - Categorical variable

Parameters intepretation - Categorical variable

R by default set the categorical variables using **dummy-coding**. In this case we set the reference category to `online`.

```
fit <- glm(solved ~ class, family = poisson(link = "log"), data = dat)
summary(fit)
```

```
##
## Call:
## glm(formula = solved ~ class, family = poisson(link = "log"),
##     data = dat)
##
## Deviance Residuals:
##     Min      1Q   Median      3Q      Max
## -2.6831  -0.7591  -0.2560   0.7691   2.6727
##
## Coefficients:
##                Estimate Std. Error z value Pr(>|z|)
## (Intercept)     2.33214    0.04407  52.925  < 2e-16 ***
## classinperson   0.37457    0.05724   6.543 6.01e-11 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for poisson family taken to be 1)
##
##     Null deviance: 183.98  on 99  degrees of freedom
## Residual deviance: 140.41  on 98  degrees of freedom
## AIC: 575.43
##
## Number of Fisher Scoring iterations: 4
```

Parameters intepretation - Categorical variable

- Similarly to the previous example, the intercept is the expected number of solved exercises when the `class` is 0. Thus the expected number of solved exercises for online students.
- the `classinperson` is the difference in log solved exercises between online and in person classes. In the response scale is the expected increase in the ratio of solved exercises. People doing in person classes solve 145.4368932 of the exercises of people doing online classes

```
c(coef(fit)["(Intercept)"], exp(coef(fit)["(Intercept)"]))
```

```
## (Intercept) (Intercept)
##    2.332144   10.300000
```

```
c(coef(fit)["classinperson"], exp(coef(fit)["classinperson"]))
```

```
## classinperson classinperson
##     0.3745721     1.4543689
```

Overdispersion

Overdispersion

Overdispersion concerns observing a greater variance compared to what would have been expected by the model.

The **overdispersion** ϕ can be estimated using Pearson Residuals:

$$\hat{\phi} = \frac{\sum_{i=1}^{n} \frac{(y_i - \hat{y}_i)^2}{\hat{y}_i}}{n - p - 1}$$

Where the numerator is the sum of squared Pearson residuals, n is the number of observations and k the number of predictors. For standard Binomial and Poisson models $\phi = 1$.

Applied Regression Analysis and Generalized Linear Models, Fox (2016), Ch. 15, pp. 432

Overdispersion

If the model is correctly specified for binomial and poisson models the
ratio is equal to 1, of the ratio is > 1 there is evidence for overdispersion.
In practical terms, if the residual deviance is higher than the residuals
degrees of freedom, there is evidence for overdispersion.

```
P <- sum(residuals(fit, type = "pearson")^2)
P / df.residual(fit) # nrow(fit$dat) - length(fit_p$coefficients)
```

```
## [1] 1.423248
```

Testing overdispersion

To formally test for overdispersion i.e. testing if the ratio is significantly different from 1 we can use the
performance::check_overdispersion() function.

```
performance::check_overdispersion(fit)
```

```
## # Overdispersion test
##
##        dispersion ratio =   1.423
##    Pearson's Chi-Squared = 139.478
##                 p-value =   0.004
```
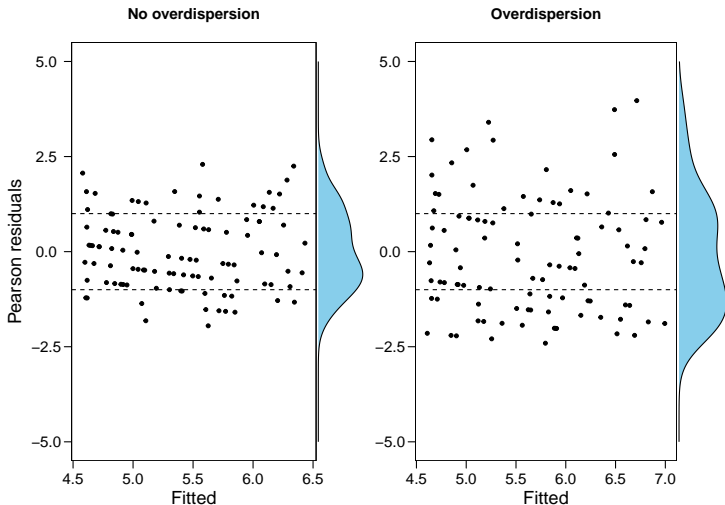
Regression and Multilevel Models, (Gelman and Hill, 2007), Ch. 6, pp. 114-115

Overdispersion plot

Pearson residuals are defined as:

$$r_p = \frac{y_i - \hat{y}_i}{\sqrt{V(\hat{y}_i)}}$$
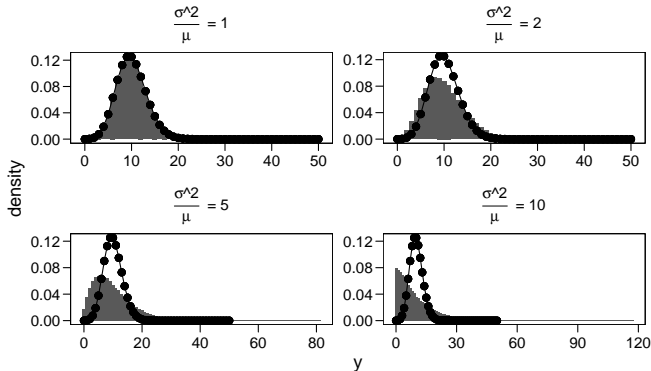
$$V(\hat{y}_i) = \sqrt{\hat{y}_i}$$

Remember that the mean and the variance are the same in Poisson models. If the model is correct, the standardized residuals should be normally distributed with mean 0 and variance 1.

Overdispersion plot

Regression and Multilevel Models, (Gelman & Hill, 2007), Ch. 6, pp. 114

Variance-mean relationship

The overdispersion can be expressed also in terms of variance-mean ratio. In fact, when the ratio is 1, there is no evidence of overdispersion.

Causes of overdispersion

There could be multiple causes for overdispersion:

- the **phenomenon itself** cannot be modelled with a Poisson distribution
- **outliers or anomalous obervations** that increases the observed variance
- **missing important variables** in the model

Outliers or anomalous data

This (simulated) dataset contains $n = 30$ observations coming from a poisson model in the form $y = 1 + 2x$ and $n = 7$ observations coming from a model $y = 1 + 10x$.

● Normal ● Outlier

Outliers or anomalous data

Clearly the sum of squared pearson residuals is inflated by these values producing more variance compared to what should be expected.

```
c(mean = mean(datout$y), var = var(datout$y)) # mean and variance should be similar
```

```
##     mean      var
## 2.756757 6.689189
```

```
performance::check_overdispersion(fit)
```

```
## # Overdispersion test
##
##        dispersion ratio =  1.515
##    Pearson's Chi-Squared = 53.019
##                 p-value =  0.026
```

Missing important variables in the model

Let's imagine to analyze again the dataset with the number of solved exercises. We have the effect of the studyh variable. In addition we have the effect of the class variable, without interaction.

id	class	studyh	class.c	lp	solved
1	online	82	0	22.6127119305986	24
2	inperson	65	1	47.7341622043588	50
3	online	12	0	11.2682503013197	15
4	inperson	54	1	42.7852617196088	50
...	...	...	...	...	...
37	online	85	0	23.2978997148077	35
38	inperson	55	1	43.2131143368049	43
39	online	80	0	22.1671521719426	26
40	inperson	77	1	53.7880487642509	54

Missing important variables in the model

We can also have a look at the data:

# Missing important variables in the model

Now let's fit the model considering only `studyh` and ignoring the group:

```r
fit <- glm(solved ~ studyh, data = dat, family = poisson(link = "log"))
summary(fit)
```

```
##
## Call:
## glm(formula = solved ~ studyh, family = poisson(link = "log"),
##     data = dat)
##
## Deviance Residuals:
##    Min      1Q  Median      3Q     Max
## -5.047  -2.247  -0.107   2.214   3.246
##
## Coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept) 2.692791   0.068831   39.12   <2e-16 ***
## studyh      0.013624   0.001063   12.82   <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: 417.99  on 39  degrees of freedom
## Residual deviance: 243.99  on 38  degrees of freedom
## AIC: 451.39
##
## Number of Fisher Scoring iterations: 4
```

Missing important variables in the model

Essentially, we are fitting an average relationship across groups but the model ignore that the two groups differs, thus the observed variance is definitely higher because we need two separate means to explain the `class` effect.

Missing important variables in the model

By fitting the appropriate model, the overdispersion is no longer a problem:

```
fit2 <- glm(solved ~ studyh + class, data = dat, family = poisson(link = "log"))
summary(fit2)
```

```
##
## Call:
## glm(formula = solved ~ studyh + class, family = poisson(link = "log"),
##     data = dat)
##
## Deviance Residuals:
##      Min        1Q    Median        3Q       Max
## -1.93006  -0.48422   0.00417   0.43834   1.97814
##
## Coefficients:
##                Estimate Std. Error z value Pr(>|z|)
## (Intercept)    2.275159   0.078239   29.08   <2e-16 ***
## studyh         0.010895   0.001068   10.21   <2e-16 ***
## classinperson  0.907365   0.065373   13.88   <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: 417.993  on 39  degrees of freedom
## Residual deviance:  29.022  on 37  degrees of freedom
## AIC: 238.41
##
## Number of Fisher Scoring iterations: 4
```

Missing important variables in the model

```
performance::check_overdispersion(fit)
```

```
## # Overdispersion test
##
##        dispersion ratio =   6.115
##   Pearson's Chi-Squared = 232.369
##                 p-value = < 0.001
```

```
performance::check_overdispersion(fit2)
```

```
## # Overdispersion test
##
##        dispersion ratio =   0.777
##   Pearson's Chi-Squared =  28.738
##                 p-value =   0.832
```

Missing important variables in the model

Also the residuals plot clearly improved after including all relevant predictors:

Why worring about overdispersion?

Before analyzing the two main strategies to deal with overdispersion, it is important to understand why it is very problematic.

Despite the estimated parameters are not affected by overdispersion, the standard error are very underestimated as a function of the degree of overdispersion.

Underestimated standard errors produce unrealistic precise parameters estimations inflating the type-1 error (i.e., parameters tends to be more significant than when overdispersion is considered)

Why worring about overdispersion?

The estimated overdispersion of `fit` is ~ 6.4208936. The `summary()` function in R has a `dispersion` argument to check how model parameters are affected.

```
phi <- fit$deviance / fit$df.residual
summary(fit, dispersion = 1)$coefficients # the default
```

```
##              Estimate  Std. Error z value    Pr(>|z|)
## (Intercept) 2.69279058 0.068830558 39.12202 0.000000e+00
## studyh      0.01362422 0.001062724 12.82009 1.265575e-37
```

```
summary(fit, dispersion = 2)$coefficients # the default
```

```
##              Estimate  Std. Error z value    Pr(>|z|)
## (Intercept) 2.69279058 0.097341109 27.663447 1.923115e-168
## studyh      0.01362422 0.001502919  9.065171 1.244091e-19
```

```
summary(fit, dispersion = phi)$coefficients # the appropriate
```

```
##              Estimate  Std. Error z value    Pr(>|z|)
## (Intercept) 2.69279058 0.174413070 15.439156 8.926080e-54
## studyh      0.01362422 0.002692888  5.059333 4.207258e-07
```

Why worring about overdispersion?

By using multiple values for ϕ we can see the impact on the standard error:

Dealing with overdispersion

Dealing with overdispersion

If all the variables are included and there are no outliers, the phenomenon itself contains more variability compared to what predicted by the Poisson. There are two main approaches to deal with the situation:

- **quasi-poisson** model
- poisson-gamma model AKA **negative-binomial model**

Quasi-poisson model

The **quasi-poisson** model is essentially a poisson model that estimate the ϕ parameter and adjust the standard errors accordingly. Again, assuming to fit the `studyh` only model (with overdispersion):

```
fit <- glm(solved ~ studyh, data = dat, family = quasipoisson(link = "log"))
summary(fit)
```

```
##
## Call:
## glm(formula = solved ~ studyh, family = quasipoisson(link = "log"),
##     data = dat)
##
## Deviance Residuals:
##    Min      1Q  Median      3Q     Max
## -5.047  -2.247  -0.107   2.214   3.246
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) 2.692791   0.170209  15.821  < 2e-16 ***
## studyh      0.013624   0.002628   5.184 7.46e-06 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for quasipoisson family taken to be 6.115055)
##
##     Null deviance: 417.99  on 39  degrees of freedom
## Residual deviance: 243.99  on 38  degrees of freedom
## AIC: NA
##
## Number of Fisher Scoring iterations: 4
```
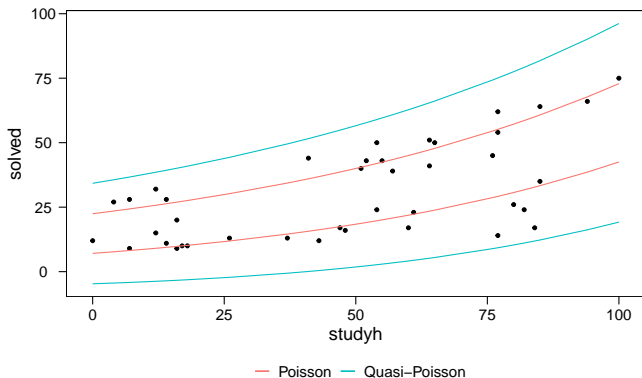
Quasi-poisson model

The quasi-poisson model estimates the same parameter and adjust standard errors as we did on slide 39. All other parameters are the same.

The quasi-poisson model is useful because it is a very simple way to deal with overdispersion.

The variance $(V(\mu))$ of the Poisson model is no longer μ but $V(\mu) = \mu\phi$. When ϕ is close to 1, the quasi-poisson model is the same as a standard poisson model.

Quasi-poisson model

Here we compare the expected variance $(\pm 2\sqrt{V(y_i)})$ of the Quasi-Poisson and Poisson models:

Problems of Quasi-* model

The main problem of quasi-* models is that they are not a specific distribution family and there is not a likelihood function. For this reason, we cannot perform model comparison the standard AIC/BIC. See [1] for an overview.

Negative-binomial model

A negative binomial model is a separated random component with two parameters: the mean as in standard poisson model and the dispersion parameter. Similarly to the quasi-poisson model it estimates a dispersion parameter.

Practically the negative-binomial model is constructed from a hierarchical model:

$$y_i \sim Poisson(\lambda_i)$$
$$\lambda_i \sim Gamma(\mu, \frac{1}{\theta})$$

In this way, the Gamma distribution regulate the dispersion around the expected value under the poisson model.

Negative-binomial model[1]

The Poisson-Gamma mixture can be expressed as:

$$p(y; \mu, \theta) = \frac{\Gamma(y + \theta)}{\Gamma(\theta)\Gamma(y+1)} \left(\frac{\mu}{\mu + \theta}\right)^y \left(\frac{\theta}{\mu + \theta}\right)^\theta$$

Where θ is the overdispersion parameter, $\Gamma()$ is the gamma function. The mean is μ and the variance is $\mu + \frac{\mu^2}{\phi}$. The θ is the inverse of the overdispersion in terms that as $1/\theta$ increase the data are more overdispersed. As $1/\theta$ approaches 0, the negative-binomial reduced to a Poisson model.
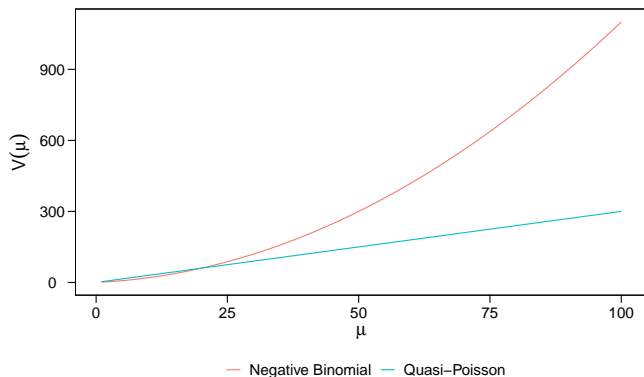
Generalized linear models, Dunn (2018), Ch. 10, pp. 400

[1]There are multiple parametrizations of the negative-binomial distribution. The one used by MASS is in terms of poisson-gamma mixture see
https://mc-stan.org/docs/2_20/functions-reference/nbalt.html

Negative-binomial model

Compared to the Poisson model, the negative binomial allows for overdispersion estimating the parameter θ and compared to the quasi-poisson model the variance is not a linear increase of the mean ($V(\mu) = \theta\mu$) but have a quadratic relationship $V(\mu) = \mu + \mu^2/\theta$


— Negative Binomial — Quasi–Poisson

Negative-binomial model

We can use the MASS package to implement the negative binomial distribution using the MASS::rnegbin():

Negative-binomial model

The θ parameter is the estimated dispersion. To note, is not the same as ϕ in the quasi-poisson model. As θ increase, the overdispersion is reduced and the model is similar to a standard Poisson model.

Negative-binomial model

For fitting a negative-binomial model we cannot use the `glm` function but we need to use the `MASS::glm.nb()` function. The syntax is almost the same but we do not need to specify the family because this function only fit negative-binomial models. Let's simulate some data coming from a negative binomial distribution with $\theta = 10$

```
theta <- 10
n <- 100
b0 <- 10
b1 <- 5
dat <- sim_design(n, nx = list(x = runif(n)))
dat$lp <- with(dat, exp(log(b0) + log(b1)*x))
dat$y <- with(dat, MASS::rnegbin(n, lp, theta))
```

id	x	y
1	0.219907104037702	8
2	0.426618238678202	21
...	...	...
99	0.619165255688131	37
100	0.965580604737625	41

Negative-binomial model
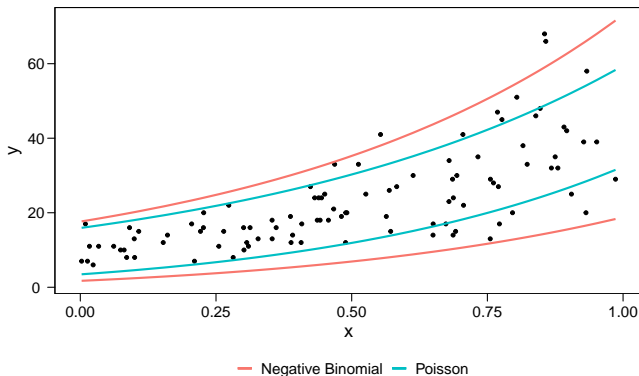
Negative-binomial model

Then we can fit the model:

```
fit_nb <- MASS::glm.nb(y ~ x, data = dat)
summary(fit_nb)
```

```
##
## Call:
## MASS::glm.nb(formula = y ~ x, data = dat, init.theta = 15.26306425,
##     link = log)
##
## Deviance Residuals:
##     Min       1Q   Median       3Q      Max
## -2.3245  -0.6955  -0.1392   0.4985   2.3194
##
## Coefficients:
##             Estimate Std. Error z value Pr(>|z|)
## (Intercept)  2.26925    0.07651   29.66   <2e-16 ***
## x            1.55794    0.12481   12.48   <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## (Dispersion parameter for Negative Binomial(15.2631) family taken to be 1)
##
##     Null deviance: 255.646  on 99  degrees of freedom
## Residual deviance:  94.579  on 98  degrees of freedom
## AIC: 672.73
##
## Number of Fisher Scoring iterations: 1
##
##
##               Theta:  15.26
##           Std. Err.:   3.49
```

Negative-binomial model

Here we compare the expected variance $(\pm 2\sqrt{V(y_i)})$ of the Negative binomial and Poisson models:

Negative Binomial — Poisson

Negative-binomial model

Here we compare the expected variance $(\pm 2\sqrt{V(y_i)})$ of the three models:

Negative Binomial — Poisson — Quasi–Poisson

Negative-binomial model

We can compare the coefficients fitting the Poisson, the Quasi-Poisson
and the Negative-binomial model:

```
fit_p <- glm(y ~ x, data = dat, family = poisson(link = "log"))
fit_qp <- glm(y ~ x, data = dat, family = quasipoisson(link = "log"))
car::compareCoefs(fit_nb, fit_p, fit_qp)
```

```
## Calls:
## 1: MASS::glm.nb(formula = y ~ x, data = dat, init.theta = 15.26306425, link
##    = log)
## 2: glm(formula = y ~ x, family = poisson(link = "log"), data = dat)
## 3: glm(formula = y ~ x, family = quasipoisson(link = "log"), data = dat)
##
##             Model 1 Model 2 Model 3
## (Intercept)  2.2692  2.2703  2.2703
## SE           0.0765  0.0537  0.0874
##
## x            1.5579  1.5561  1.5561
## SE           0.1248  0.0807  0.1312
##
```

Negative-binomial model (NB) vs Quasi-poisson (QP)

- The NB has the likelihood function thus AIC, LRT and other likelihood-based metrics works compared to the QP
- The NB assume a different mean-variance relationship thus estimated coefficients could be different to the P where QP produce the same estimates.
- Both NB and QP estimate higher standard errors in the presence of overdispersion
- If there is evidence of overdispersion, the important is to fit a model that take into account it

If you want to try to simulate NB data you can use the
`MASS::rnegbin(n, mu, theta)` function or the `rnb(n, mu, vmr)`
custom function that could be more intuitive because requires μ (i.e., the
mean) and `vmr` that is the desired variance-mean ratio. Using `message =`
`TRUE` it will tells the θ value:

```
y <- rnb(1e5, mu = 10, vmr = 7, message = TRUE)
```

```
## y ~ NegBin(mu = 10, theta = 1.67), var = 70.00, vmr = 7.00
```

```
nprint(mu = mean(y), v = var(y), vmr = var(y) / mean(y))
```

```
##     mu      v    vmr
## 9.958 69.569  6.986
```

---

[2]Using `vmr = 1` it will use the `rpois()` function

Deviance based pseudo-R^2

The Deviance based pseudo-R^2 is computed from the ratio between the residual deviance and the null deviance[3]:

$$R^2 = 1 - \frac{D_{current}}{D_{null}}$$

```
1 - fit_p$deviance/fit_p$null.deviance
```

```
## [1] 0.6078928
```

---

[3]https://en.wikipedia.org/wiki/Pseudo-R-squared

Applied Regression Analysis and Generalized Linear Models, Fox (2016), Ch. 14.1, pp. 383

McFadden's pseudo-R^2

The McFadden's pseudo-R^2 compute the ratio between the log-likelihood of the intercept-only (i.e., null) model and the current model [2]:

$$R^2 = 1 - \frac{\log(\mathcal{L}_{current})}{\log(\mathcal{L}_{null})}$$

There is also the adjusted version that take into account the number of parameters of the model. In R can be computed manually or using the `performance::r2_mcfadden()`:

```
performance::r2_mcfadden(fit_p)
```

```
## # R2 for Generalized Linear Regression
##        R2: 0.349
##    adj. R2: 0.348
```

Nagelkerke's pseudo-R^2

The Cox and Snell's [3] R^2 is defined as:

$$R^2 = 1 - \left(\frac{\mathcal{L}_{null}}{\mathcal{L}_{current}}\right)^{\frac{2}{n}}$$

Where Nagelkerke [4] provide a correction to set the range of values between 0 and 1.

```
performance::r2_nagelkerke(fit_p)
```

```
## Nagelkerke's R2
##       0.9827151
```

References

[1]   J. M. Ver Hoef and P. L. Boveng, "Quasi-Poisson vs. Negative binomial regression: How should we model overdispersed count data?" *Ecology*, vol. 88, no. 11, pp. 2766–2772, Nov. 2007, doi: 10.1890/07-0043.1.

[2]   D. McFadden, "Regression-based specification tests for the multinomial logit model," *J. Econom.*, vol. 34, no. 1, pp. 63–82, Jan. 1987, doi: 10.1016/0304-4076(87)90067-4.

[3]   D. R. Cox and E. J. Snell, *Analysis of binary data, second edition*. CRC Press, 1989.

[4]   N. J. D. Nagelkerke, "A note on a general definition of the coefficient of determination," *Biometrika*, vol. 78, no. 3, pp. 691–692, 1991, doi: 10.1093/biomet/78.3.691.

✉ filippo.gambarota@unipd.it

⭘ github.com/filippogambarota