## **Methods of Applied Statistics I**

### STA2101H F LEC9101

Week 8

### November 3 2021





- 1. Upcoming events
- 2. Logistic Regression
- 3. Poisson Regression
- 4. In the News
- 5. Homework, Project (Hour 3)



• Thursday Nov 4 3.30

Pragya Sur, Harvard



Precise High-Dimensional Asymptotics for AdaBoost Zoom Link



• Thursday Nov 4 3.30

# Precise High-Dimensional Asymptotics for AdaBoost Zoom Link

Friday Nov 5 Toronto Data Workshop Zoom link

Dear friends,



Toronto Data Workshop this Friday, 5 November, at noon (Toronto time) hosts Yun William Yu on the intersection of math and data science. Yun William Yu is an assistant professor in the math department at UofT whose research focuses on algorithmic methods for computational biology and medical informatics.

Link: https://utoronto.zoom.us/j/84277066292 Meeting ID: 842 7706 6292 Passcode: data\_4\_lyf

- last of linear models: factorial treatment structure, CR and RB designs, interaction plots, estimation of variance, comparison of group means
- regression with binomial response *y*: logistic transform, fitting by ML, interpretation of coefficients, Challenger data, linear predictor, variance-covariance matrix
- estimation of β; estimation of var(β), based on likelihood theory statistics secret sauce

y= Jin, yn iadep. f: (y:; 0) deusity f(y; Q] = T(f(y; Q)) T t ; i= varlar fixed Q, "frue value" lik  $f=L(0,g) \propto f(4,0)$ of fixed = c(y) f(y; 0)lloig) = lop L (O-, y) + a(y) maningles 20) = I lopf: (y.; 0) h(5, g°) L(4; yo)

How to use 
$$L(\theta;y)$$
 or  $L(\theta;y)$ ? provide site  
1.  $\rightarrow$  Bayesia analysis  $\pi(\theta/\theta) = \frac{L(\theta;y)\pi(\theta)}{\int L(\theta;y)\pi(\theta)d\theta} = m(y)$   
post. busily  $\int L(\theta;y)\pi(\theta)d\theta = m(y)$   
2.  $\frac{1}{2}L(\theta;y) = 0$   $\hat{\theta} = max.$  Lit. ast.  
 $\partial \theta = \frac{1}{2} \ln \theta + \frac{1}{2} \ln \theta$ 

- model:  $y_i \sim f(y_i; \theta), i = 1, \dots, n$
- joint density:  $f(\underline{y}; \theta) = \prod_{i=1}^{n} f(y_i; \theta)$
- likelihood function  $L(\theta; \underline{y}) = f(\underline{y}; \theta)$

independent

• model: 
$$y_i \sim f(y_i; \theta), i = 1, \dots, n$$

- joint density:  $f(\underline{y}; \theta) = \prod_{i=1}^{n} f(y_i; \theta)$
- likelihood function  $L(\theta; \underline{y}) = f(\underline{y}; \theta)$
- log-likelihood function  $\ell(\theta; \underline{y}) = \log L(\theta; \underline{y}) = \sum_{i=1}^{n} \log f(y_i; \theta)$
- maximum likelihood estimate  $\hat{\theta} = \arg \sup \ell(\theta; \underline{y});$
- Fisher information  $j(\theta) = -\ell''(\theta)$
- two theorems:

$$(\hat{\theta} - \theta) j^{1/2}(\hat{\theta}) \stackrel{d}{\rightarrow} N(\mathsf{O}, I)$$

asymptotically normal

• likelihood ratio statistic

*p* is dimension of  $\theta$ 

Applied Statistics I

L10.14

November 3 2021

independent

 $\ell'(\hat{\theta}) = \mathbf{0}$ 

• two theorems:

$$(\hat{\theta} - \theta) j^{1/2}(\hat{\theta}) \stackrel{d}{\to} N(0, I)$$
  
 $w(\theta) = 2\{\ell(\hat{\theta}) - \ell(\theta)\} \stackrel{d}{\to} \chi_p^2$ 

• two theorems:  

$$\begin{aligned}
(\hat{\theta} - \theta)j^{1/2}(\hat{\theta}) \stackrel{d}{\rightarrow} N(0, I) \\
W(\theta) &= 2\{\ell(\hat{\theta}) - \ell(\theta)\} \stackrel{d}{\rightarrow} \chi_{p}^{2}
\end{aligned}$$
• two approximations  

$$\begin{aligned}
\hat{\theta}_{k} \approx N(\{\theta_{k}, j^{-1}(\hat{\theta})_{kk}\} &\leftarrow each component \\
W(\theta) \approx \chi_{p}^{2}
\end{aligned}$$

$$\begin{aligned}
\hat{\theta} - \theta_{0} \rangle j^{1/2}(\hat{\theta}) \stackrel{d}{\rightarrow} N(0, T) & \text{when } y \sim f(y, \theta) &= \Pi f_{1}(y; \theta) \\
\vdots &\vdots &\vdots \\
\end{aligned}$$
Applied Statistics November 3 2021
  

$$\begin{aligned}
\hat{\theta} - \theta_{0} \rangle j^{1/2}(\hat{\theta}) \stackrel{d}{\rightarrow} N(0, T) &\text{when } y \sim f(y, \theta) &= \Pi f_{1}(y; \theta) \\
\vdots &\vdots &\vdots \\
\end{aligned}$$

two theorems:

$$(\hat{\theta} - \theta) j^{1/2}(\hat{\theta}) \stackrel{d}{\to} N(0, I)$$
$$w(\theta) = 2\{\ell(\hat{\theta}) - \ell(\theta)\} \stackrel{d}{\to} \chi_p^2$$

two approximations

$$\hat{\theta}_k \sim N(\{\theta_k, j^{-1}(\hat{\theta})_{kk}\} )$$

$$w(\theta) \sim \chi_p^2$$

compare two models using change in likelihood ratio statistic



**Applied Statistics I** 



8



maximum likelihood estimate

 $\partial \ell(\beta; \mathbf{y}) / \partial \beta = \mathbf{0}$ 



### **Nested models**

- Comparing two models:
- likelihood ratio test

$$2\{\ell_A(\hat{\beta}_A) - \ell_B(\hat{\beta}_B)\} \land \chi \qquad \text{where fitter }$$

compares the maximized log-likelihood function under model A and model B

- Comparing two models:
- likelihood ratio test

$$2\{\ell_{\mathsf{A}}(\hat{\beta}_{\mathsf{A}}) - \ell_{\mathsf{B}}(\hat{\beta}_{\mathsf{B}})\}$$

compares the maximized log-likelihood function under model A and model B

• example model A:  $logit(p_i) = \beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i}, \quad \beta_A = (\beta_0, \beta_1, \beta_2)$ model B:  $logit(p_i) = \beta_0 + \beta_1 x_{1i}, \quad \beta_B = (\beta_0, \beta_1)$ 



- Comparing two models:
- likelihood ratio test

$$2\{\ell_{\mathsf{A}}(\hat{\beta}_{\mathsf{A}}) - \ell_{\mathsf{B}}(\hat{\beta}_{\mathsf{B}})\}$$

compares the maximized log-likelihood function under model A and model B

- example model A: logit( $p_i$ ) =  $\beta_0 + \beta_1 x_{1i} + \beta_2 x_{2i}$ ,  $\beta_A = (\beta_0, \beta_1, \beta_2)$ model B: logit( $p_i$ ) =  $\beta_0 + \beta_1 x_{1i}$ ,  $\beta_B = (\beta_0, \beta_1)$
- when model B is nested in model A, LRT is approximately χ<sup>2</sup><sub>ν</sub> distributed, under model B
   ν = dim(A) dim(B)

### ... nested models

>> logitmodcorrect <- glm(cbind(r,m-r) ~ temperature / family = binomial, data = shuttle2)
>> logitmodcorrect2 <- glm(cbind(r,m-r) ~ temperature + pressure, family = binomial, data = shuttle2)
>> summary(logitmodcorrect2)

### Coefficients:

Estimate Std. Error z value Pr(>|z|) 0.27 = pral(Intercept) 2.520195 3.486784 0.723 0.4698 0.0285 \* temperature -0.098297 0.044890 -2.190 1.105 0.2691 pressure 0.008484 0.007677 Check summary ( U) Null deviance; 24.230 on 22 degrees of freedom Residual deviance: 16.546 on 20 degrees of freedom ATC: 36,106  $\frac{23 \text{ obs}^{-5} \text{ s}}{23 \text{ obs}^{-5} \text{ s}} : \text{ just fit } \beta_0 \quad \text{``null model'' rors dual df } 23 - 1 = 22$ Applied Statistics 1 November 3 2021 june fit } \beta\_0 \beta\_1 \beta\_2 \quad \text{full model} \quad \text{``23 - 3 = 20 ft} Number of Fisher Scoring iterations: 5 « 23-3 = 20 € ... nested models

> logitmodcorrect2 <- glm(cbind(r,m-r) ~ temperature + pressure, family = binomial, data = shuttle2)</pre>

B. B. reduced model ris. dt 23-2=21

```
> anova(logitmodcorrect,logitmodcorrect2)
Analysis of Deviance Table
```

now fit



**Applied Statistics I** 

November 3 2021

### ...nested models

- Model A:  $logit(p_i) = \beta_0 + \beta_1 temp_i + \beta_2 pressure_i$
- Model B:  $logit(p_i) = \beta_0 + \beta_1 temp_i$
- nested: Model B is obtained by setting  $\beta_2 = 0$
- Under Model B, the change in deviance is (approximately) an observation from a  $\chi_1^2$

### ...nested models

- Model A:  $logit(p_i) = \beta_0 + \beta_1 temp_i + \beta_2 pressure_i$
- Model B:  $logit(p_i) = \beta_0 + \beta_1 temp_i$
- nested: Model B is obtained by setting  $\beta_2 = 0$
- Under Model B, the change in deviance is (approximately) an observation from a  $\chi^2_1$

Robel A = Bo E

Rodel B= Bie B, temp

13

•  $\Pr(\chi_1^2 \ge 1.5407) = 0.22$ : this is a p-value for testing  $H_0: \beta_2 = 0$ (other fect  $\beta_2 \pm \cdots + \beta_2$  p she 0.27 ELM-1 p.30 Applied Statistics 1 November 3 2021 resid drive  $A - 2 \log L(\hat{\Theta}) + constant \int_A^{-1} \frac{1}{2} \log L(\hat{\Theta}) + constant \int_A^{-1} \frac{1}{2$ 

constant E B  $\frac{1}{2}\left(\hat{o}_{A}\right) = 2\left(\hat{o}_{B}\right) - 2\left(\hat{o}_{B}\right)^{2}$ BAAI

• confidence intervals for  $\beta_1$ 



- confidence intervals for  $\beta_1$
- based on normal approximation:  $\hat{\beta}_1 \pm \widehat{\text{s.e.}}(\hat{\beta}_1) * 1.96$



- confidence intervals for  $\beta_1$
- based on normal approximation:  $\hat{\beta}_1 \pm \widehat{\text{s.e.}}(\hat{\beta}_1) * 1.96$
- (-0.208, -0.023)



**Applied Statistics I** 

- confidence intervals for  $\beta_1$
- based on normal approximation:  $\hat{\beta}_1 \pm \widehat{\text{s.e.}}(\hat{\beta}_1) * 1.96$
- (-0.208, -0.023)
- based on profile log-likelihood

model A  $w_A^{(\Theta)} = 2 \int R_A(\hat{\Theta}_A) - R_A(\Theta) \int$ " B  $W_R = 2 \int l_B(\hat{o}) - l_B(\hat{o}) \int$ 



November 3 2021 e.g. 1 = 3 ( Bot Bit Bo) - 2 (Bo, Bi)

- confidence intervals for  $\beta_1$
- based on normal approximation:  $\hat{\beta}_1 \pm \widehat{\text{s.e.}}(\hat{\beta}_1) * 1.96$
- (-0.208, -0.023)
- based on profile log-liketihood

(-0.2122262, -0.0244701)

confint(logitmodcorrect):



ß,

 $\mathcal{L}(\mathcal{B}_{\circ},\mathcal{B}_{\circ}) =$ 



 $\mathcal{L}(\widehat{\beta}_{o}(\beta,\gamma,\beta,\gamma) = \mathcal{L}_{p}(\beta,\gamma)$ 

ELM-1 p. 31

Applied Statistics I November 3 2021

### **Binary data**

15

- each response is Binary:  $y_i = 0, 1$
- explanatory variables  $x_i^T$  as usual
- same model

instead of  $0, 1, \ldots, m_i$ 

 $\operatorname{pr}(y_i = 1 \mid x_i) = p_i(\beta) = \frac{\exp(x_i^T \beta)}{1 + \exp(x_i^T \beta)}$ 2 dr -h(y) STO ncipi yER": s(x) **Applied Statistics I** November 3 2021 (S.,.... Sp, a, ... a, ...) [y.,..., yr)

### **Binary data**

- each response is Binary:  $y_i = 0, 1$
- explanatory variables  $x_i^T$  as usual
- same model

$$\operatorname{pr}(y_i = 1 \mid x_i) = p_i(\beta) = \frac{\exp(x_i^T \beta)}{1 + \exp(x_i^T \beta)}$$

- example: SM 10.18
- example HW6: "The math group, the single dependent variable of this work, was coded as a dichotomous variable (1: math group vs. 0: nonmath group)."
- "To classify the math vs. nonmath groups, we also executed a binary logistic regression."

instead of  $0, 1, \ldots, m_i$ 

15

O, P  $\mathcal{Q}(\Theta_{1},\ldots,\Theta_{p}, y)$ interest Oz ... Op are new rance parameter  $l(\theta_1, \theta_2; \gamma)$ was over Dz, keep D,  $l_p(o_i)$ foxe A Star Constant  $nux l_3(0,)$  $l_{p}(\hat{o})$ 8

### **Binary data**

instead of  $0, 1, \ldots, m_i$ 

- each response is Binary:  $y_i = 0, 1$
- explanatory variables  $x_i^T$  as usual
- same model

$$\operatorname{pr}(y_i = 1 \mid x_i) = p_i(\beta) = \frac{\exp(x_i^T \beta)}{1 + \exp(x_i^T \beta)}$$

- example: SM 10.18
- example HW6: "The math group, the single dependent variable of this work, was coded as a dichotomous variable (1: math group vs. 0: nonmath group)."
- "To classify the math vs. nonmath groups, we also executed a binary logistic regression."
- example wcgs data, ELM-2, Ch.2

### ELM-2 Ch.2

## **Binary data**

> dat	ta(wc	gs, pac	kage="fara	away")					
> hea	ad(wc	gs); he	elp(wcgs) #	#latte:	r not sh	lown	$\sim$		
	age 1	height	weight sdp	dbp	chol ber	nave)	cigs )		
<u> </u>	49	73	150 110	76	225	A2	25		
2002	42	70	160 154	l 84	177	A2	20		
2003	42	69	160 110	) 78	181	BЗ	0		
2004	41	68	152 124	l 78	132	B4	20		
2005	59	,70	150 144	86	255	B3	20		
2006	44	72	204 150	) 90	182	B4	0		1
/	dibe	p chd	typechd ti	mechd	arcus	3		- y = c	nd
2001		B no	none	1664	absent	5			
2002	]	B no	none	3071	present	5		VI VI	Ø
2003	L	A no	none	3071	absent	5		v ye	2S *
2004		A no	none	3064	absent	5		ŀ	
Applied Stati 2005	istics I	A yes i	nideath	1885	present	5			

## **Binary and Count data**

Nancy

2/11/2021

## **Binary data**

data(wcgs, package = "faraway")

head(wcgs) #not run: str(wcgs); plot(wcgs); help(wcgs)

##		age h	eight	weight	sdp	dbp	chol	behave	cigs	dibep	chd	typechd	timechd
##	2001	49	73	150	110	76	225	A2	25	В	no	none	1664
##	2002	42	70	160	154	84	177	A2	20	В	no	none	3071
##	2003	42	69	160	110	78	181	в3	0	A	no	none	3071
##	2004	41	68	152	124	78	132	В4	20	A	no	none	3064
##	2005	59	70	150	144	86	255	в3	20	A	yes	infdeath	1885
##	2006	44	72	204	150	90	182	В4	0	A	no	none	3102
##		arc	us										
##	2001	abse	nt										
##	2002	prese	nt										
##	2003	abse	nt										
##	2004	abse	nt										
##	2005	prese	nt										

• where's the epsilon?

• where's the epsilon? There isn't one

## ... Binary responses

- where's the epsilon? There isn't one
- what's the model? It has two parts

### ... Binary responses

- where's the epsilon? There isn't one
- what's the model? It has two parts
- Regression.

$$\mathbb{E}(y_i) = p_i = \frac{\exp(x_i^{\mathrm{T}}\beta)}{1 + \exp(x_i^{\mathrm{T}}\beta)}$$

• Probability distribution.

 $y_i \sim Bernoulli(p_i)$ 

### ... Binary responses



- if you add a lot of Bernoulli's together, all with the same  $p_i$ , you get
- how could they have the same *p<sub>i</sub>* in our model?



### **Binomial responses**

- if you add a lot of Bernoulli's together, all with the same  $p_i$ , you get
- how could they have the same *p<sub>i</sub>* in our model?
- $p_i = function(x_i^{T}\beta)$
- different observations with the same *p<sub>i</sub>* are called <u>covariate classes</u>
- Example 10.18 in SM Table 10.8 has 23 rows of binomials sample sizes vary from 1 to 6
- data(nodal) in library(SMPracticals) has 53 rows of binary observations

### **Binomial responses**

- if you add a lot of Bernoulli's together, all with the same  $p_i$ , you get
- how could they have the same *p<sub>i</sub>* in our model?
- $p_i = function(x_i^{T}\beta)$
- different observations with the same *p<sub>i</sub>* are called covariate classes
- Example 10.18 in SM Table 10.8 has 23 rows of binomials sample sizes vary from 1 to 6
- data(nodal) in library(SMPracticals) has 53 rows of binary observations
- R expects cbind(r, m-r) in glm with binomial data
- but if all observations are binary you can get away with r only
- see ?family (check Details)
- you can also specify proportions  $y_i/n_i$ , but then you need to use weights

## **Binomial/Binary**

10.4 · Prope	491							
Table 10.8     Data on       nodal involvement     (Brown, 1980).	m	r	age	stage	grade	xray	acid	
	~>6	>5	0	1	1	1	1 6	- covarate
	$\rightarrow 6$	-1	0	0	0	0	1	
	-94	-50	1	1	1	0	0	alasses
	4	2	1	1	0	0	1	0411.5
	4	0	0	0	0	0	0	
	3	2	0	1	1	0	1	
	3	1	1	1	0	0	0	
	3	0	1	0	0	0	1	
	3	0	1	0	0	0	0	
Can we predict hodal	2	0	1	0	0	1	0	
involvement from other	2	1	0	1	0	0	1	
	2	1	0	0	1	0	0	
measurements?	1	1	1	1	1	1	1	
	1	1	1	1	0	1	1	
	1	1	1	0	1	1	1	
	1	1	1	0	0	1	1	
	1	0	1	0	1	0	0	
	1	1	0	1	1	1	0	
	1	0	0	1	1	0	0	
Applied Statistics I November 3 2021	1	1	U	1	0	I	0	
	1	1	0	0	1	0	1	

20

## **Binary and Count data**

Nancy

2/11/2021

### **Binary data**

data(wcgs, package = "faraway")
head(wcgs) #not run: str(wcgs); plot(wcgs); help(wcgs)

##		age 1	height	weight	sdp	dbp	chol	behave	cigs	dibep	chd	typechd	timechd	
##	2001	49	73	150	110	76	225	A2	25	В	no	none	1664	
##	2002	42	70	160	154	84	177	A2	20	в	no	none	3071	
##	2003	42	69	160	110	78	181	в3	0	А	no	none	3071	
##	2004	41	68	152	124	78	132	в4	20	A	no	none	3064	
##	2005	59	70	150	144	86	255	В3	20	A	ves	infdeath	1885	
##	2006	44	72	204	150	90	182	B4	0	A	no	none	3102	
##	2000	ar	2115	204	100	50	102	24	0		110	none	0102	
##	2001	abe	ant											
##	2001	absi	enc											
##	2002	pres	ent											
##	2003	abs	ent											
##	2004	abs	ent											
##	2005	pres	ent											
##	2006	abs	ent											

## $\longrightarrow$ .Rmd

- likelihood ratio test for logistic model  $p_i = p_i(\beta) = \exp(x_i^T \beta)$ ,  $\hat{p}_i = p_i(\hat{\beta})$
- this model is nested in the saturated model  $\tilde{p}_i = y_i/n_i$

and Poisson

- likelihood ratio test for logistic model  $p_i = p_i(\beta) = \exp(x_i^T\beta)$ ,  $\hat{p}_i = p_i(\hat{\beta})$
- this model is nested in the saturated model  $\tilde{p}_i = y_i/n_i$
- · residual deviance compares fitted model to saturated model

and Poisson

- likelihood ratio test for logistic model  $p_i = p_i(\beta) = \exp(x_i^T\beta)$ ,  $\hat{p}_i = p_i(\hat{\beta})$
- this model is nested in the saturated model  $\tilde{p}_i = y_i/n_i$
- residual deviance compares fitted model to saturated model
- under the fitted model, approximately distributed as  $\chi^2_{n-q}$

and Poisson

• likelihood ratio test for logistic model  $p_i = p_i(\beta) = \exp(x_i^T\beta)$ ,  $\hat{p}_i = p_i(\hat{\beta})$ 

F.

xiB

- this model is nested in the saturated model  $\tilde{p}_i = y_i/n_i$
- residual deviance compares fitted model to saturated model
- under the fitted model, approximately distributed as  $\chi^2_{n-q}$  if each  $n_i$  "large"
- > summary(Ex1018.glm)

(Dispersion parameter for binomial family taken to be 1)

Null deviance: 40.710 on 22 degrees of freedom Residual deviance: 18.069 on 17 degrees of freedom AIC: 41.69  $h_i = 6, 5, 4, 2,$ 

" laze "

Applied Statistics I November 3 2021

ELM-1 p.29

y:~Bin(ui,pi)

Mi Xi Vi

> step(ex1018binom)

Coefficients:

(Intercept) stage xray acid -3.052 1.645 1.912 1.638 y, ∓ bi (n, p,) :

Degrees of Freedom: 22 Total (i.e. Null); 19 Residual Null Deviance: ^^I 40.71 Residual Deviance: 19.64 ^^IAIC: 39.26

$$P_i = y_i / n_i$$

- we can drop age and grade without affecting quality of the fit  $\sum_{i=1}^{\infty} l(\tilde{p}_i) =$ 

- in other words the model can be simplified by setting two regression coefficients to zero
- several mistakes in text on pp. 491,2;

- deviances in Table 10.9 are incorrect as well http://statwww.epfl.ch/davison/SM/ has corrected version

### ... example 10.18: variable selection

- step implements stepwise regression
- evaluates each fit using AIC =  $-2\ell(\hat{\beta}; y) + 2p$
- penalizes models with larger number of parameters
- we can also compare fits by comparing deviances



 $Ey_i = \mu_i$  i = 1..., n

### ... example 10.18: variable selection

- step implements stepwise regression
- evaluates each fit using AIC =  $-2\ell(\hat{\beta}; y) + 2p$
- penalizes models with larger number of parameters
- we can also compare fits by comparing deviances
- > update(ex1018binom, . ~ . aged stage)

Coefficients:

(Intercept) grade xray acid -2.734 1.420 1.750 1.797

Degrees of Freedom: 22 Total (i.e. Null); 19 Residual Null Deviance: 40.71 Residual Deviance: 21.28 AIC: 40.9

> deviance(ex1018binom)
[1] 18.06869
> pchisq(21.28-18.07,df=2,lower=F)
[1] 0.2008896

ER

Pr, ..., Pn Satid Br, ..., Be consatid

Sullos

- as terms are added to the model, deviance always decreases
- because log-likelihood function always increases
- similar to residual sum of squares

.

- as terms are added to the model, deviance always decreases
- because log-likelihood function always increases
- similar to residual sum of squares
- Akaike Information Criterion penalizes models with more parameters

$$AIC = 2\{-\ell(\hat{\beta}; y) + p\}$$

• comparison of two model fits by difference in AIC

Applied Statistics I November 3 2021

SM (4.57)

> summary(ex1018binom)

Call: glm(formula = cbind(r, m - r) ~ ., family = binomial, data = nodal2)

Deviance Residuals:

Min	1Q	Median	ЗQ	Max
-1.4989	-0.7726	-0.1265	0.7997	1.4351

> summary(ex1018binom)

Call: glm(formula = cbind(r, m - r) ~ ., family = binomial, data = nodal2)

Deviance Residuals:

Min 1Q Median 3Q Max -1.4989 -0.7726 -0.1265 0.7997 1.4351



### Generalized linear models

```
glm has several options for family
binomial(link = "logit")
gaussian(link = "identity")
Gamma(link = "inverse")
inverse.gaussian(link = "1/mu^2")
poisson(link = "log")
quasi(link = "identity", variance = "constant")
quasibinomial(link = "logit")
quasipoisson(link = "log")
```

fouly = bonowal



Each of these is a member of the class of generalized linear models Generalized: distribution of response is not assumed to be normal

Linear: some transformation of  $E(y_i)$  is of the form  $x_i^T \beta$ 

Applied Statistics I November 3 2021

.01

link function

### **Poisson regression**

•

.

• the Poisson distribution is a useful starting point for data that counts events

$$f(y_i \mid x_i) = \frac{1}{y!} \mu_i^{y_i} e^{-\mu_i}, y_i = 0, 1, \dots,$$

$$f(\mathbf{y}_i \mid \mathbf{x}_i) = \exp\{\mathbf{y}_i \log \mu_i - \mu_i - \log(\mathbf{y}_i!)\}$$

• canonical parameter

$$\theta_i = \log(\mu_i)$$

• linear model:

$$\log(\mu_i) = \mathbf{x}_i^{\mathsf{T}} \beta$$

equivalently

$$E(\mathbf{y}_i) = \mu_i = \exp(\mathbf{x}_i^{\mathsf{T}}\beta)$$



- coding 1 for "lack math", 0 otherwise; p.6 + data
- *t*-test with 84 (and 83) df; Fig 2
- how many predictors in logistic regression? p.2,3
- conclusions p. 4



Welch's *t*-test  $S_{b}^{2} = (n_{1}-1)S_{1}^{2} + (n_{2}-1)S_{1}^{2}$ N,+n\_ 30

### HW Question Week 4

STA2101F 2021

#### Due October 14 2021 11.59 pm

#### Homework to be submitted through Quercus

Part 1: Data set for project Okay to submit October 21

Please submit details about the data you will use for your project. Ideally the data will have a single response or outcome variable of interest, and several potential explanatory variables. You should submit with this homework:

- (1) the data source: both bibliographic and a web link
- (2) the number of observations and the number of potential explanatory variables
- (3) a description of the response variable
- (4) a description of the potential explanatory variables
- (5) the scientific question(s) of interest

When you submit the final project, it will consist of the parts listed in Slide 3 on October 6.

#### Part 2: Question(s) for marking

There has been a lot of talk this week about rapid testing in the schools. On one hand there seems no harm in using rapid antigen tests on a regular basis, but on the other hand if a lot of the tests give incorrect results, especially flagging as covid-related too often, then children will unnecessarily miss school. This seems to be the main concern from the public health officials who are cautioning a slower approach.

Tests for Covid19 (or any screening for that matter), are assessed by their false positive and false negative rates, or equivalently by their sensitivity and specificity. Sensitivity of the test is the true positive rate, i.e.  $Pr(T+ \mid C+)$ , and 1 minus sensitivity is the false negative rate Pr(T - | C +). Specificity of the test is the true negative rate, i.e. pr(T - | C -), and 1 minus specificity is the *false positive rate*. (My source is Wikipedia.)

(a) If a given student tests positive, compute the probability that s/he has Covid19 using Bayes theorem. A gift:



Projects

$$(c) - (f(\underline{s}, \underline{q})) = e^{\sum s \top \theta - nc(\theta) \cdot - k(s) \cdot s}$$

$$T f(\underline{y}, \underline{\theta}) = T e^{s(\underline{y}, \overline{\theta}) - k(\theta) - h(\underline{y}, \overline{\theta})}$$

$$= f(\underline{y}, \theta)$$

 $\int f(y; \theta) dy =$   $\begin{cases} \text{sats} \\ f(y; \theta) dy \\ f(y, y_2) \rightarrow f(y, y_2) \\ \text{yerc}^{\uparrow}: \\ \text{yerc}^{\uparrow}: \\ \text{fig}_{i} + y_{\overline{z}}^{\downarrow} \\ \text{fig}_{i} + y_{$