

1.3 Likelihoods for Regression Models

We will start with linear regression and then talk about more general models.

1.3.1 Linear Model

↓
nonlinear
& LM

Consider the familiar linear model

$$Y_i = \mathbf{x}_i^\top \boldsymbol{\beta} + \varepsilon_i, \quad i = 1, \dots, n,$$

where $\mathbf{x}_1, \dots, \mathbf{x}_n$ are known nonrandom vectors.

$$E[\varepsilon_i] = 0 \quad \text{and} \quad \text{Var}[\varepsilon_i] = \sigma^2$$

often estimate $\boldsymbol{\beta}$ by $\hat{\boldsymbol{\beta}}_{OLS}$, which does not require a distribution for ε_i .

For likelihood-based estimation, we need a distribution for ε_i . Start w/ $\varepsilon_i \sim N(0, \sigma^2)$.

$$\begin{aligned} \Rightarrow L(\boldsymbol{\beta}, \sigma | \{Y_i, \mathbf{x}_i\}_{i=1}^n) &= \prod_{i=1}^n \left(\frac{1}{\sqrt{2\pi}\sigma} \right) \exp\left(-\frac{(Y_i - \mathbf{x}_i^\top \boldsymbol{\beta})^2}{2\sigma^2} \right) \\ &= \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^n (Y_i - \mathbf{x}_i^\top \boldsymbol{\beta})^2 \right) \end{aligned}$$

take log,
derivatives, set = 0,
solve

$$\hat{\boldsymbol{\beta}}_{MLE} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \quad \text{same as } \hat{\boldsymbol{\beta}}_{OLS}!$$


$$\hat{\sigma}_{MLE}^2 = \frac{1}{n} \sum_{i=1}^n (Y_i - \mathbf{x}_i^\top \hat{\boldsymbol{\beta}})^2 \quad \text{(only asymptotically unbiased).}$$

What do you do when ϵ_i are not Gaussian?

- transform data so ϵ_i look Gaussian.
- Use a different distribution for ϵ_i !

Example (Venice sea levels): The annual maximum sea levels in Venice for 1931–1981 are :

We know maxima are not Gaussian!



Approach 1: OLS $E[\epsilon_i] = 0, \text{Var}[\epsilon_i] = \sigma^2$ No distributional assumption.

Approach 2: Assume $\epsilon_i \sim \text{Gumbel}$ (extreme value dist), use ML

$$f(y; \sigma) = \frac{1}{\sigma} \exp\left(-\frac{y}{\sigma}\right) \exp\left(-\exp\left(-\frac{y}{\sigma}\right)\right).$$

$$\Rightarrow L(\beta, \sigma | \{y_i, x_i\}_{i=1}^n) = \prod_{i=1}^n f(y_i - x_i^T \beta) = \prod_{i=1}^n \frac{1}{\sigma} \exp\left(-\frac{y_i - x_i^T \beta}{\sigma}\right) \exp\left(-\exp\left(-\frac{y_i - x_i^T \beta}{\sigma}\right)\right)$$

YOUR TURN: Fit both approaches to the Venice data.

$$\text{OLS} \\ \hat{\beta}_0 = 104.8 \quad \hat{\beta}_1 = .56 \quad (\text{SE } .177)$$

$$\text{MLE, GUMBEL} \\ \hat{\beta}_0 = 96.8, \hat{\beta}_1 = 0.56 \quad (\text{SE } .136)$$

OLS vs MLE? IF EV model is correct, more efficient (note: standard errors).

$$\beta_0 \text{ difference: } E[\epsilon_i] = 0.577\sigma = 0.577 \hat{\sigma}_{\text{MLE}} = 0.577(14.5) \neq 0$$

$$96.8 + 0.577(14.5) = 105.1$$

1.3.2 Additive Errors Nonlinear Model

previous example had ① linear trend, ② Non-Gaussian errors.

Non-linear additive model:


$$Y_i = g(x_i, \beta) + \epsilon_i$$

often interested in $\epsilon_i \sim N(0, \sigma^2)$ but $g(x_i, \beta) \neq \sum \beta_j x_j \Rightarrow$ ML required.

① non-linear trend, ② Gaussian errors.

Example: exponential growth model

$$g(x, \beta) = \beta_0 \exp(\beta_1 x)$$



1.3.3 Generalized Linear Models

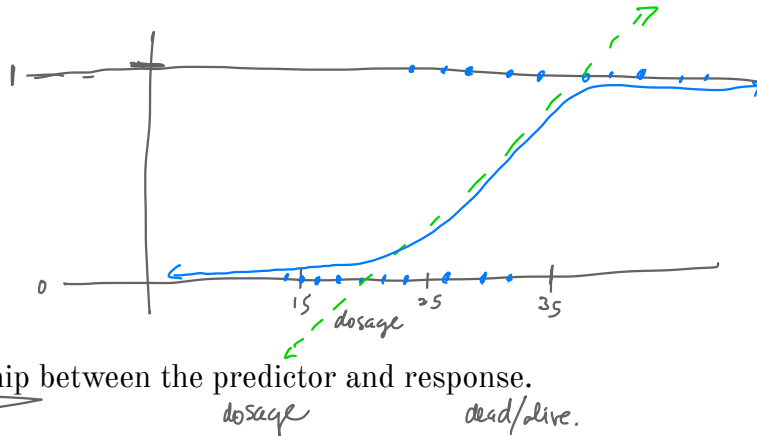
Regression: build a relationship between a parameter (mean) & covariates. ↗ of response.

LM's: stochastic element is additive w/ mean.

GLM's: stochastic element is different.

Imagine an experiment where individual mosquitos are given some dosage of pesticide. The response is whether the mosquito lives or dies. The data might look something like:

x (dosage)	y (0=lives, 1=dies)
15	0
17	0
18	1
20	0
21	1
⋮	⋮
⋮	⋮



Goal: Model the relationship between the predictor and response.

Sounds like regression!

Big difference: y_i 's are not continuous. They only take values of 0 or 1.

Question: What would a curve of best fit look like? Would we want a function that only takes values in $\{0, 1\}$.

It seems sensible to have a curve which takes values near 0 for low doses & near 1 for high doses and intermediate values for middle doses.

what does this curve represent? Probability.

Refined Goal: Model relationship between predictor (dosage) + probability of successful response (mosquito dies).
 Let's build a sensible model. Note: We don't observe the probability.

Step 1: Find a function that behaves the way we want.

like the blue curve.


Consider the logistic function,

$$p = \frac{\exp(\theta)}{1 + \exp(\theta)}$$

$$\text{As } \theta \rightarrow \infty, p \rightarrow 1$$

$$\theta \rightarrow -\infty, p \rightarrow 0$$

$$\theta = 0, p = \frac{1}{2}$$



By changing θ , we can change location, slope, direction of this function.

$$\text{Let } \theta = \beta_0 + \beta_1 x \Rightarrow p = \frac{\exp(\beta_0 + \beta_1 x)}{1 + \exp(\beta_0 + \beta_1 x)}$$

```
# understanding the logistic function
```

```
# first, theta just equals x
```

```
x <- seq(-7, 7, .1)
```

```
theta <- x
```

```
y <- exp(theta)/(1 + exp(theta))
```



```
ggplot() + geom_line(aes(x, y))
```

```
# now, let theta be a linear function of x
```

```
theta <- 1 + 3*x
```

```
y <- exp(theta)/(1 + exp(theta))
```

```
ggplot() + geom_line(aes(x, y))
```



Now we can connect probabilities to covariate x !

We'd be done if we observed probabilities, but our response only takes values of 0 and 1.

Step 2: Build a stochastic mechanism to relate to a binary response.

Recall the Bernoulli distribution

$$Y = \begin{cases} 0 & \text{w.p. } 1-p \\ 1 & \text{w.p. } p. \end{cases}$$

biased
Coin flip example w/ $p = 0.75$. Flip coin, you will observe 0 (tails) or 1 (heads).

Aside: We could instead think about binomial dsn, which counts # of successes for n trials.

$$X = \sum_{i=1}^n Y_i, \quad Y_i \stackrel{\text{iid}}{\sim} \text{Bern}(p). \quad X \text{ takes values in } \{0, 1, \dots, n\}.$$

$$P(X=k) = \binom{n}{k} p^k (1-p)^{n-k}.$$

Step 3: Put Step 1 and Step 2 together.

$Y_i \stackrel{\text{iid}}{\sim} \text{Bernoulli}(p_i)$

↑
outcome of i^{th} observation (observed).

↑
prob of i^{th} observation having success (unobserved).

$$p_i = \frac{\exp(\theta_i)}{1 + \exp(\theta_i)}, \quad \theta_i = \beta_0 + \beta_1 X_i$$

Goal: estimate β_0 and β_1 . Find the "best" estimates.

Fitting our model: Does OLS make sense? No.

What else can we do? Maximum likelihood!

↳ Find the parameters (β s) which make the density agree best w/ data we observed!

$$\text{pmf of Bernoulli: } f(y_i | p_i) = p_i^{y_i} (1-p_i)^{1-y_i}$$

⇒ take y_i 's to estimate p_i 's.

Consider the likelihood contribution.

$$L_i(p_i | Y_i) = p_i^{Y_i} (1-p_i)^{1-Y_i} \quad (Y_i \text{'s are } 0 \text{ or } 1).$$

So the log-likelihood contribution is

$$l_i(p_i) = Y_i \log p_i + (1-Y_i) \log (1-p_i) = \underbrace{\log (1-p_i) + Y_i \log \frac{p_i}{1-p_i}}_{(*)}$$

Recall, we said $p_i = \frac{\exp(\theta_i)}{1+\exp(\theta_i)}$ was sensible.

Manipulating: $p_i + p_i \exp(\theta_i) = \exp(\theta_i)$

$$p_i = (1-p_i) \exp(\theta_i).$$

OR

$$p_i \exp(-\theta_i) = 1-p_i$$

$$\frac{p_i}{1-p_i} = \exp(\theta_i)$$

$$\log\left(\frac{p_i}{1-p_i}\right) = \theta_i \quad (1)$$

$$\frac{\exp(\theta_i)}{1+\exp(\theta_i)} \exp(-\theta_i) = 1-p_i$$

$$\frac{1}{1+\exp(\theta_i)} = 1-p_i$$

$$-\log(1+\exp(\theta_i)) = \log(1-p_i) \quad (2).$$

Plug in in (1) and (2) into (*).

Which gives us,

$$l_i(\theta_i) = \underbrace{-\log(1+\exp(\theta_i)) + Y_i \theta_i}_{\text{now in terms of } \theta_i \text{ not } p_i}$$

notice now the form of the data is "nice" for MLE things.
Why? Because log-likelihood + "sensible" functions $p_i = \frac{\exp(\theta_i)}{1+\exp(\theta_i)}$
work well together.

Not a coincidence.

So the log-likelihood is

$$l(\theta_1, \dots, \theta_n) = \sum_{i=1}^n l_i(\theta_i)$$

$$= \sum_{i=1}^n \{-\log(1+\exp(\theta_i)) + Y_i \theta_i\}$$

$$\Rightarrow l(\beta_0, \beta_1) = \sum_{i=1}^n \{-\log(1+\exp(\beta_0 + \beta_1 x_i)) + Y_i (\beta_0 + \beta_1 x_i)\}$$

To optimize? *Must be done numerically.*

```
## data on credit default
data("Default", package = "ISLR")
head(Default)
```

```
## default student balance income
## 1 No No 729.5265 44361.625
## 2 No Yes 817.1804 12106.135
## 3 No No 1073.5492 31767.139
## 4 No No 529.2506 35704.494
## 5 No No 785.6559 38463.496
## 6 No Yes 919.5885 7491.559
```

```
## fit model with ML
m0 <- glm(default ~ balance, data = Default, family = binomial)
tidy(m0) |> kable()
```

*↑
broom
package.*

term	$\hat{\beta}_0, \hat{\beta}_1$ estimate	$se(\hat{\beta}_0), se(\hat{\beta}_1)$ std.error	$\frac{\hat{\beta}_i}{se(\hat{\beta}_i)}$ statistic	p.value
(Intercept)	-10.6513306	0.3611574	-29.49221	0
balance	0.0054989	0.0002204	24.95309	0

data's are 0,1

*optimize
likelihood
numerically.*

*H₀: β_i = 0
H_a: β_i ≠ 0
i=1,2*

```
glance(m0) |> kable()
```


null.deviance	df.null	logLik	AIC	BIC	deviance	df.residual	nobs
2920.65	9999	-798.2258	1600.452	1614.872	1596.452	9998	10000

```
## plot the curve
```

```
x_new <- seq(0, 2800, length.out = 200)
theta <- m0$coefficients[1] + m0$coefficients[2]*x_new
p_hat <- exp(theta)/(1 + exp(theta))
```

```
ggplot() +
  geom_point(aes(balance, as.numeric(default) - 1), alpha = 0.5, data
    = Default) +
  geom_line(aes(x_new, p_hat), colour = "blue") +
  ylab("Probability of Defaulting")
```

$l(\hat{\beta}_0, \hat{\beta}_1)$.



never outside of $[0, 1]$ → valid probabilities!

In general, a GLM is three pieces:

1. The random component

probability dsr from
exponential family.

Ex: Logistic Regression

$$Y_i \sim \text{Binom}(p_i)$$

Explanation:

describe the generating
mechanism of observed data

2. The systemic component

A function relating the parameter
of interest (mean!) to θ

$$E[Y] = \bar{g}'(\eta)$$

$$p_i = \frac{\exp(\theta_i)}{1 + \exp(\theta_i)} = \bar{g}'(\theta_i)$$

↑
Note $Y_i \sim \text{Bern}(p_i)$
 $E[Y_i] = p_i$

transforming linear relationship
to be on a scale that makes
sense for the parameter of interest
"linking" linear relationship to mean.

3. A linear predictor

$$\theta = X\beta$$

$$\theta_i = \alpha_i \beta$$

describing how θ is a linear
function of predictor variables.

Remarks:

(1) standard formulation denotes function by \bar{g}' : $p = \bar{g}'(\theta) = \frac{\exp(\theta)}{1 + \exp(\theta)}$

$$\Rightarrow \theta = g(p) = \log\left(\frac{p}{1-p}\right)$$

(2) Parameter of interest is still the mean, just like linear regression.

(3) Theoretical reasons for exponential family ... relationship btw/ param of interest & Variance.

→ for count data.

Example (Poisson regression):

(1) Poisson(λ).

$$(3) \theta = X\beta$$

$$(2) \lambda = \bar{g}'(\theta) = \exp(\theta) \quad (\lambda > 0)$$

$$\downarrow$$

$$\theta = \log(\lambda)$$

Some background
on choice of g .

Consider a general family of distributions:

subfamily of exponential
family dets that includes
Binomial, Poisson, etc.

$$\log f(y_i; \theta_i, \phi) = \frac{y_i \theta_i - b(\theta_i)}{a_i(\phi)} + c(y_i, \phi).$$

$$\begin{aligned} f(y_i; \theta_i, \phi) &= \exp \left\{ \frac{y_i \theta_i - b(\theta_i)}{a_i(\phi)} + c(y_i, \phi) \right\} \\ &= \exp \left\{ \underbrace{\frac{y_i \theta_i}{a_i(\phi)} + c(y_i, \phi)}_* - \underbrace{\frac{b(\theta_i)}{a_i(\phi)}}_{**} \right\}. \end{aligned}$$

recall exponential family w/ parameter $\underline{\theta} = (\theta_1, \dots, \theta_s)^T$ is of the form:

$$f(y; \underline{\theta}) = h(y) \exp \left\{ \underbrace{\sum_{j=1}^s g_j(\underline{\theta}) T_j(y)}_* - \underbrace{B(\underline{\theta})}_{**} \right\}$$

assumes $T_1(y_i) = y_i$
 $g_1(\underline{\theta}) = \frac{\theta_i}{a_i(\phi)}$

subfamily of exponential family.

similar to single param exp family except dispersion term $a_i(\phi)$.

Example (Normal model):

We can learn something about this distribution by considering its mean and variance. Because we don't have an explicit form of the density, we rely on two facts:

$$1. \mathbf{E} \left[\frac{\partial \log f(Y_i; \theta_i, \phi)}{\partial \theta_i} \right] = 0.$$

$$2. \mathbf{E} \left[\frac{\partial^2 \log f(Y_i; \theta_i, \phi)}{\partial \theta_i^2} \right] + \mathbf{E} \left[\left(\frac{\partial \log f(Y_i; \theta_i, \phi)}{\partial \theta_i} \right)^2 \right] = 0.$$

$$\text{For } \log f(y_i; \theta_i, \phi) = \frac{y_i \theta_i - b(\theta_i)}{a_i(\phi)} + c(y_i, \phi),$$

Example (Bernoulli model):

$$f(y_i; p_i) = p_i^{y_i} (1 - p_i)^{1-y_i}$$

Finally, back to modelling. Our **goal** is to build a relationship between the mean of Y_i and covariates \boldsymbol{x}_i .

Example (Bernoulli model, cont'd):

1.4 Marginal and Conditional Likelihoods

Consider a model which has $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$, where $\boldsymbol{\theta}_1$ are the parameters of interest and $\boldsymbol{\theta}_2$ are nuisance parameters.

One way to improve estimation for $\boldsymbol{\theta}_1$ is to find a one-to-one transformation of the data \mathbf{Y} to (\mathbf{V}, \mathbf{W}) such that either

The key feature is that one component of each contains only the parameter of interest.

Example (Neyman-Scott problem): Let $Y_{ij}, i = 1, \dots, n, j = 1, 2$ be independent normal random variables with possible different means μ_i but the same variance σ^2 .

Our goal is to estimate σ^2 . Should we be able to?

Following the usual arguments,

$$\hat{\mu}_{i,\text{MLE}} = \frac{Y_{i1} + Y_{i2}}{2}$$
$$\hat{\sigma}_{\text{MLE}}^2 = \frac{1}{2n} \sum_{i=1}^n \sum_{j=1}^2 (Y_{ij} - \hat{\mu}_{i,\text{MLE}})^2$$

$$\mathbb{E}[\hat{\sigma}_{\text{MLE}}^2] =$$

A reworking of the data seems more promising. Let,

$$V_i = \frac{Y_{i1} - Y_{i2}}{\sqrt{2}} \quad \text{and} \quad W_i = \frac{Y_{i1} + Y_{i2}}{\sqrt{2}}$$

For conditional likelihoods, we can often exploit the existence of sufficient statistics for the nuisance parameters under the assumption that the parameter of interest is known.

Example (Exponential Families): The structure of exponential families is such that it is often possible to exploit their properties to eliminate nuisance parameters. Let Y have a density of the form

$$f(\mathbf{y}; \boldsymbol{\eta}) = h(\mathbf{y}) \exp \left\{ \sum_{i=1}^s \eta_i T_i(\mathbf{y}) - A(\boldsymbol{\eta}) \right\},$$

then

Thus, exponential families often provide an automatic procedure for finding \mathbf{W} and \mathbf{U} .

Example (Logistic Regression): For binary Y_i , the standard logistics regression model is

$$P(Y_i = 1) = p_i(\mathbf{x}_i, \boldsymbol{\beta}) = \frac{\exp(\mathbf{x}_i^\top \boldsymbol{\beta})}{1 + \exp(\mathbf{x}_i^\top \boldsymbol{\beta})}$$

and the likelihood is

$$L(\boldsymbol{\beta} | \mathbf{Y}, \mathbf{X}) =$$