

# Lecture 1

## Spatio-temporal data & Linear Models

---

Colin Rundel

1/16/2018

## (Bayesian) Linear Models

---

Pretty much everything we are going to see in this course will fall under the umbrella of either linear or generalized linear models.

In previous classes most of your time has likely been spent with the simple iid case,

$$Y_i = \beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip} + \epsilon_i$$
$$\epsilon_i \sim N(0, \sigma^2)$$

these models can also be expressed simply as,

$$Y_i \stackrel{iid}{\sim} N(\beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip}, \sigma^2)$$

We can also express using matrix notation as follows,

$$\begin{aligned} \mathbf{Y} &= \mathbf{X} \boldsymbol{\beta} + \boldsymbol{\epsilon} \\ n \times 1 & \quad n \times p \quad p \times 1 \quad n \times 1 \\ \boldsymbol{\epsilon} &\sim N(\mathbf{0}, \sigma^2 \mathbf{1}_n) \\ n \times 1 & \quad n \times 1 \quad n \times n \end{aligned}$$

or alternative as,

$$\mathbf{Y} \sim N \left( \begin{array}{cc} \mathbf{X} & \boldsymbol{\beta} \\ n \times p & p \times 1 \end{array}, \begin{array}{c} \sigma^2 \mathbf{1}_n \\ n \times n \end{array} \right)$$

## Multivariate Normal Distribution - Review

For an  $n$ -dimension multivariate normal distribution with covariance  $\Sigma$  (positive semidefinite) can be written as

$$\mathbf{Y}_{n \times 1} \sim N\left(\boldsymbol{\mu}_{n \times 1}, \boldsymbol{\Sigma}_{n \times n}\right) \text{ where } \{\boldsymbol{\Sigma}\}_{ij} = \rho_{ij}\sigma_i\sigma_j$$

$$\begin{pmatrix} Y_1 \\ Y_2 \\ \vdots \\ Y_n \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_n \end{pmatrix}, \begin{pmatrix} \rho_{11}\sigma_1\sigma_1 & \rho_{12}\sigma_1\sigma_2 & \cdots & \rho_{1n}\sigma_1\sigma_n \\ \rho_{21}\sigma_2\sigma_1 & \rho_{22}\sigma_2\sigma_2 & \cdots & \rho_{2n}\sigma_2\sigma_n \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{n1}\sigma_n\sigma_1 & \rho_{n2}\sigma_n\sigma_2 & \cdots & \rho_{nn}\sigma_n\sigma_n \end{pmatrix}\right)$$

## Multivariate Normal Distribution - Density

For the  $n$  dimensional multivariate normal given on the last slide, its density is given by

$$(2\pi)^{-n/2} \det(\Sigma)^{-1/2} \exp\left(-\frac{1}{2}(\mathbf{Y} - \boldsymbol{\mu})' \Sigma^{-1} (\mathbf{Y} - \boldsymbol{\mu})\right)$$

$1 \times n$        $n \times n$        $n \times 1$

and its log density is given by

$$-\frac{n}{2} \log 2\pi - \frac{1}{2} \log \det(\Sigma) - \frac{1}{2}(\mathbf{Y} - \boldsymbol{\mu})' \Sigma^{-1} (\mathbf{Y} - \boldsymbol{\mu})$$

$1 \times n$        $n \times n$        $n \times 1$







Likelihood:

$$\mathbf{Y} | \boldsymbol{\beta}, \sigma^2 \sim N(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{1}_n)$$

Likelihood:

$$\mathbf{Y} \mid \boldsymbol{\beta}, \sigma^2 \sim N(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{1}_n)$$

Priors:

$$\beta_i \sim N(0, \sigma_\beta^2) \text{ or } \boldsymbol{\beta} \sim N(\mathbf{0}, \sigma_\beta^2 \mathbf{1}_p)$$

$$\sigma^2 \sim \text{Inv-Gamma}(a, b)$$

$$\begin{aligned}[\beta, \sigma^2 | \mathbf{Y}, \mathbf{X}] &= \frac{[\mathbf{Y} | \mathbf{X}, \beta, \sigma^2]}{[\mathbf{Y} | \mathbf{X}]} [\beta, \sigma^2] \\ &\propto [\mathbf{Y} | \mathbf{X}, \beta, \sigma^2] [\beta, \sigma^2] \\ &\propto [\mathbf{Y} | \mathbf{X}, \beta, \sigma^2] [\beta | \sigma^2] [\sigma^2]\end{aligned}$$

## Deriving the posterior

$$\begin{aligned}[\boldsymbol{\beta}, \sigma^2 \mid \mathbf{Y}, \mathbf{X}] &= \frac{[\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2]}{[\mathbf{Y} \mid \mathbf{X}]} [\boldsymbol{\beta}, \sigma^2] \\ &\propto [\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2] [\boldsymbol{\beta}, \sigma^2] \\ &\propto [\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2] [\boldsymbol{\beta} \mid \sigma^2] [\sigma^2]\end{aligned}$$

where,

$$f(\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2) = (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2}(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})'(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})\right)$$

## Deriving the posterior

$$\begin{aligned}[\boldsymbol{\beta}, \sigma^2 \mid \mathbf{Y}, \mathbf{X}] &= \frac{[\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2]}{[\mathbf{Y} \mid \mathbf{X}]} [\boldsymbol{\beta}, \sigma^2] \\ &\propto [\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2] [\boldsymbol{\beta}, \sigma^2] \\ &\propto [\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2] [\boldsymbol{\beta} \mid \sigma^2] [\sigma^2]\end{aligned}$$

where,

$$f(\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2) = (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2}(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})'(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})\right)$$

$$f(\boldsymbol{\beta} \mid \sigma_\beta^2) = (2\pi\sigma_\beta^2)^{-p/2} \exp\left(-\frac{1}{2\sigma_\beta^2}\boldsymbol{\beta}'\boldsymbol{\beta}\right)$$

$$\begin{aligned}[\boldsymbol{\beta}, \sigma^2 \mid \mathbf{Y}, \mathbf{X}] &= \frac{[\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2]}{[\mathbf{Y} \mid \mathbf{X}]} [\boldsymbol{\beta}, \sigma^2] \\ &\propto [\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2] [\boldsymbol{\beta}, \sigma^2] \\ &\propto [\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2] [\boldsymbol{\beta} \mid \sigma^2] [\sigma^2]\end{aligned}$$

where,

$$f(\mathbf{Y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2) = (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2}(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})'(\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})\right)$$

$$f(\boldsymbol{\beta} \mid \sigma^2) = (2\pi\sigma^2)^{-p/2} \exp\left(-\frac{1}{2\sigma^2}\boldsymbol{\beta}'\boldsymbol{\beta}\right)$$

$$f(\sigma^2 \mid a, b) = \frac{b^a}{\Gamma(a)} (\sigma^2)^{-a-1} \exp\left(-\frac{b}{\sigma^2}\right)$$

## Deriving the Gibbs sampler ( $\sigma^2$ step)

$$\begin{aligned} [\beta, \sigma^2 | \mathbf{Y}, \mathbf{X}] &\propto (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2}(\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta)\right) \\ &\quad (2\pi\sigma_\beta^2)^{-p/2} \exp\left(-\frac{1}{2\sigma_\beta^2}\beta'\beta\right) \\ &\quad \frac{b^a}{\Gamma(a)}(\sigma^2)^{-a-1} \exp\left(-\frac{b}{\sigma^2}\right) \end{aligned}$$





## Deriving the Gibbs sampler ( $\beta$ step)

$$\begin{aligned} [\beta, \sigma^2 | \mathbf{Y}, \mathbf{X}] &\propto (2\pi\sigma^2)^{-n/2} \exp\left(-\frac{1}{2\sigma^2}(\mathbf{Y} - \mathbf{X}\beta)'(\mathbf{Y} - \mathbf{X}\beta)\right) \\ &\quad (2\pi\sigma_\beta^2)^{-p/2} \exp\left(-\frac{1}{2\sigma_\beta^2}\beta'\beta\right) \\ &\quad \frac{b^a}{\Gamma(a)}(\sigma^2)^{-a-1} \exp\left(-\frac{b}{\sigma^2}\right) \end{aligned}$$



## A Quick Example

---

## Some Fake Data

Lets generate some simulated data where the underlying model is known and see how various regression procedures function.

$$\beta_0 = 0.7, \quad \beta_1 = 1.5, \quad \beta_2 = -2.2, \quad \beta_3 = 0.1$$

$$n = 100, \quad \epsilon_i \sim N(0, 1)$$

## Generating the data

```
set.seed(01162018)
n = 100
beta = c(0.7, 1.5, -2.2, 0.1)
eps = rnorm(n)

d = data_frame(
  X1 = rt(n, df=5),
  X2 = rt(n, df=5),
  X3 = rt(n, df=5)
) %>%
  mutate(Y = beta[1] + beta[2]*X1 + beta[3]*X2 + beta[4]*X3 + eps)
```

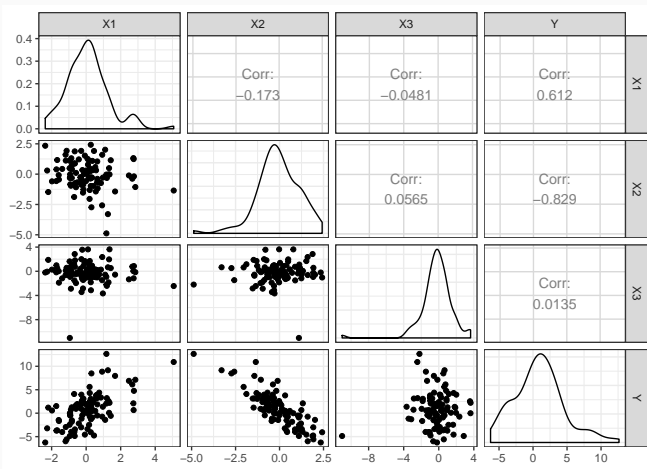
## Model Matrix

```
X = model.matrix(~X1+X2+X3, d)
tbl_df(X)
```

```
## # A tibble: 100 x 4
##   `(Intercept)`      X1      X2      X3
##   <dbl>      <dbl> <dbl> <dbl>
## 1           1  1.26 -3.30  0.664
## 2           1  1.17 -4.88 -2.20
## 3           1 -0.590 -1.08 -1.95
## 4           1  0.358  0.319 -0.582
## 5           1  1.20 -0.314 -2.41
## 6           1  0.856 -0.733  0.274
## 7           1  0.649 -0.385  0.112
## 8           1 -1.58 -0.449 -1.89
## 9           1 -0.424  1.22 -0.193
## 10          1  0.0808  1.27 -0.488
## # ... with 90 more rows
```

# Pairs plot

```
GGally::ggpairs(d, progress = FALSE)
```



Let  $\hat{\mathbf{Y}}$  be our estimate for  $\mathbf{Y}$  based on our estimate of  $\boldsymbol{\beta}$ ,

$$\hat{\mathbf{Y}} = \hat{\beta}_0 + \hat{\beta}_1 \mathbf{X}_1 + \hat{\beta}_2 \mathbf{X}_2 + \hat{\beta}_3 \mathbf{X}_3 = \mathbf{X} \hat{\boldsymbol{\beta}}$$



Let  $\hat{\mathbf{Y}}$  be our estimate for  $\mathbf{Y}$  based on our estimate of  $\boldsymbol{\beta}$ ,

$$\hat{\mathbf{Y}} = \hat{\beta}_0 + \hat{\beta}_1 \mathbf{X}_1 + \hat{\beta}_2 \mathbf{X}_2 + \hat{\beta}_3 \mathbf{X}_3 = \mathbf{X} \hat{\boldsymbol{\beta}}$$

The least squares estimate,  $\hat{\boldsymbol{\beta}}_{ls}$ , is given by

$$\arg \min_{\boldsymbol{\beta}} \sum_{i=1}^n (Y_i - \mathbf{X}_i \cdot \boldsymbol{\beta})^2$$

Let  $\hat{\mathbf{Y}}$  be our estimate for  $\mathbf{Y}$  based on our estimate of  $\boldsymbol{\beta}$ ,

$$\hat{\mathbf{Y}} = \hat{\beta}_0 + \hat{\beta}_1 \mathbf{X}_1 + \hat{\beta}_2 \mathbf{X}_2 + \hat{\beta}_3 \mathbf{X}_3 = \mathbf{X} \hat{\boldsymbol{\beta}}$$

The least squares estimate,  $\hat{\boldsymbol{\beta}}_{ls}$ , is given by

$$\arg \min_{\boldsymbol{\beta}} \sum_{i=1}^n (Y_i - \mathbf{X}_i \cdot \boldsymbol{\beta})^2$$

Previously we derived,

$$\hat{\boldsymbol{\beta}}_{ls} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{Y}$$

## Frequentist Fit

```
l = lm(Y ~ X1 + X2 + X3, data=d)
l$coefficients
```

```
## (Intercept)          X1          X2          X3
##  0.6566561    1.4657537  -2.2807109    0.1629704
```

```
(beta_hat = solve(t(X) %*% X, t(X)) %*% d$Y)
```

```
##                [,1]
## (Intercept)  0.6566561
## X1           1.4657537
## X2           -2.2807109
## X3           0.1629704
```

## Bayesian model specification (JAGS)

```
model =
"model{
  # Likelihood
  for(i in 1:length(Y)){
    Y[i] ~ dnorm(mu[i], tau)
    mu[i] = beta[1] + beta[2]*X1[i] + beta[3]*X2[i] + beta[4]*X3[i]
  }

  # Prior for beta
  for(j in 1:4){
    beta[j] ~ dnorm(0,1/100)
  }

  # Prior for sigma / tau2
  tau ~ dgamma(1, 1)
  sigma2 = 1/tau
}"
```

# Compiling

```
m = rjags::jags.model(  
  textConnection(model),  
  data = d, n.chains = 2  
)  
  
## Compiling model graph  
##   Resolving undeclared variables  
##   Allocating nodes  
## Graph information:  
##   Observed stochastic nodes: 100  
##   Unobserved stochastic nodes: 5  
##   Total graph size: 810  
##  
## Initializing model
```

```
# Burnin
update(m, n.iter=1000, progress.bar="none")

# Draw samples
samp = rjags::coda.samples(
  m, variable.names=c("beta", "sigma2"),
  n.iter=5000, progress.bar="none"
)
```

```
str(samp)
```







```
## List of 2
## $ : 'mcmc' num [1:5000, 1:5] 0.477 0.533 0.513 0.608 0.47 ...
##   ..- attr(*, "dimnames")=List of 2
##     .. ..$ : NULL
##     .. ..$ : chr [1:5] "beta[1]" "beta[2]" "beta[3]" "beta[4]" ...
##   ..- attr(*, "mcpair")= num [1:3] 1001 6000 1
## $ : 'mcmc' num [1:5000, 1:5] 0.729 0.351 0.576 0.687 0.622 ...
##   ..- attr(*, "dimnames")=List of 2
##     .. ..$ : NULL
##     .. ..$ : chr [1:5] "beta[1]" "beta[2]" "beta[3]" "beta[4]" ...
##   ..- attr(*, "mcpair")= num [1:3] 1001 6000 1
## - attr(*, "class")= chr "mcmc.list"
```

```
head(samp[[1]])
```

```
## Markov Chain Monte Carlo (MCMC) output:  
## Start = 1001  
## End = 1007  
## Thinning interval = 1  
##      beta[1] beta[2] beta[3] beta[4] sigma2  
## [1,] 0.4772517 1.419778 -2.199144 0.28649330 1.0809728  
## [2,] 0.5328665 1.521434 -2.391340 0.10170084 1.2145136  
## [3,] 0.5127307 1.569942 -2.288688 0.09279966 1.2330530  
## [4,] 0.6082786 1.457989 -2.301124 0.11507748 0.7760024  
## [5,] 0.4696911 1.497949 -2.427537 0.15902657 1.1581875  
## [6,] 0.7505117 1.397314 -2.231748 0.04573775 0.9684167  
## [7,] 0.6361385 1.625917 -2.208012 0.21898668 1.3867957
```



`plot(samp)`



## Tidy Bayes


```
df = samp %>%  
  tidybayes::gather_draws(beta[i], sigma2) %>%  
  mutate(param = ifelse(is.na(i), .variable, paste0(.variable,"[,i,]")))
```

df

```
## # A tibble: 50,000 x 7  
## # Groups:   i, .variable [5]  
##   .chain .iteration .draw     i .variable .value param  
##   <int>      <int> <int> <int> <chr>      <dbl> <chr>  
## 1         1         1     1     1 beta         0.477 beta[1]  
## 2         1         1     1     2 beta         1.42  beta[2]  
## 3         1         1     1     3 beta        -2.20  beta[3]  
## 4         1         1     1     4 beta         0.286 beta[4]  
## 5         1         2     2     1 beta         0.533 beta[1]  
## 6         1         2     2     2 beta         1.52  beta[2]  
## 7         1         2     2     3 beta        -2.39  beta[3]  
## 8         1         2     2     4 beta         0.102 beta[4]  
## 9         1         3     3     1 beta         0.513 beta[1]  
## 10        1         3     3     2 beta         1.57  beta[2]  
## # ... with 49,990 more rows
```


## Tidy Bayes + ggplot - Traceplot

```
ggplot(df, aes(x=.iteration, y=.value, color=as.character(.chain))) +  
  geom_line(alpha=0.5) +  
  facet_wrap(~param, scale="free_y") +  
  labs(color="Chain")
```



## Tidy Bayes + ggplot - Density plot

```
ggplot(df, aes(x=.value, fill=as.character(.chain))) +  
  geom_density(alpha=0.5) +  
  facet_wrap(~param, scale="free_x") +  
  labs(fill="Chain")
```



## Comparing Approaches

```
pt_est = df %>%  
  filter(.chain == 1) %>%  
  group_by(param) %>%  
  summarize(post_mean = mean(.value)) %>%  
  ungroup() %>%  
  mutate(  
    truth = c(0.7, 1.5, -2.2, 0.1, 1),  
    ols   = c(l$coefficients, var(l$residuals))  
  ) %>%  
  select(param, truth, ols, post_mean)
```

pt\_est

```
## # A tibble: 5 x 4  
##   param      truth      ols post_mean  
##   <chr>    <dbl> <dbl>    <dbl>  
## 1 beta[1]    0.7  0.657    0.656  
## 2 beta[2]    1.5  1.47     1.47  
## 3 beta[3]   -2.2 -2.28    -2.28  
## 4 beta[4]    0.1  0.163    0.164  
## 5 sigma2    1    1.08     1.14
```

## Comparing Approaches - code

```
ggplot(df, aes(x=.value, fill=as.character(.chain))) +  
  geom_density(alpha=0.5) +  
  facet_wrap(~param, scale="free_x") +  
  geom_vline(  
    data = tidyr::gather(pt_est, pt_est, value, -param),  
    aes(xintercept = value, color=pt_est)  
  ) +  
  labs(fill="Chain")
```

# Comparing Approaches - plot

