

2019 SISG Bayesian Statistics for Genetics

R Notes: Generalized Linear Modeling

Jon Wakefield

Departments of Statistics and Biostatistics, University of
Washington

2019-07-15

Overview

In this set of notes a number of generalized linear models (GLMs) and generalized linear mixed models (GLMMs) will be fitted using Bayesian methods.

Two primary computational techniques will be illustrated:

- integrated nested Laplace approximation (INLA)
- Markov chain Monte Carlo (MCMC) using Stan

Case control example: Data

We analyze a case control example using logistic regression models, first using likelihood methods.

The data concern the numbers of cases (of the disease Leber Hereditary Optic Neuropathy) and controls as a function of genotype at a particular location (rs6767450).

```
x <- c(0, 1, 2)
# Case data for CC CT TT
y <- c(6, 8, 75)
# Control data for CC CT TT
z <- c(10, 66, 163)
```

Case control example: Likelihood analysis

We fit the logistic regression model as a generalized linear model and then examine the estimate and an asymptotic (large sample) 95% confidence interval.

```
logitmod <- glm(cbind(y, z) ~ x, family = "binomial")
thetahat <- logitmod$coeff[2] # Log odds ratio
thetahat
##          x
## 0.4787428
exp(thetahat) # Odds ratio
##          x
## 1.614044
V <- vcov(logitmod)[2, 2] # standard error2
# Asymptotic confidence interval for odds ratio
exp(thetahat - 1.96 * sqrt(V))
##          x
## 0.9879159
exp(thetahat + 1.96 * sqrt(V))
##          x
## 2.637004
```

Case control example: Likelihood analysis

Now let's look at a likelihood ratio test of $H_0 : \theta = 0$ where θ is the log odds ratio associated with the genotype (multiplicative model).

```
logitmod
##
## Call:  glm(formula = cbind(y, z) ~ x, family = "binomial")
##
## Coefficients:
## (Intercept)          x
##   -1.8077         0.4787
##
## Degrees of Freedom: 2 Total (i.e. Null);  1 Residual
## Null Deviance:      15.01
## Residual Deviance: 10.99    AIC: 27.79
dev <- logitmod$null.deviance - logitmod$deviance
dev
## [1] 4.01874
pchisq(dev, df = logitmod$df.residual, lower.tail = F)
## [1] 0.04499731
```

So just significant at the 5% level.

FTO Example: Data

We reproduce the least squares analysis of the FTO data.

```
fto <- structure(list(y = c(2.14709237851534, 6.16728664844416,
  6.5287427751799, 13.7905616042756, 13.6590155436307,
  1.75906323176397, 6.77485810485697, 9.67664941025843,
  11.751562703307, 12.3892232256873, 9.7235623369017,
  12.1796864728229, 14.8575188389164, 16.370600225645,
  27.7498618362862, 6.61013278196954, 11.3676194738021,
  17.9876724213706, 22.4424423901962, 26.687802642435),
  X = structure(c(1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
    1, 1, 1, 1, 1, 1, 1, 1, 1, 0, 0, 0, 0, 0, 0,
    0, 0, 0, 0, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1, 1,
    2, 3, 4, 5, 1, 2, 3, 4, 5, 1, 2, 3, 4, 5, 1,
    2, 3, 4, 5, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 1,
    2, 3, 4, 5, 1, 2, 3, 4, 5), .Dim = c(20L, 4L),
    .Dimnames = list(NULL, c("", "xg", "xa", ""))),
  .Names = c("y", "X"))
```

FTO Example: Data and LS Fit

The `lm` function uses MLE, which is equivalent to ordinary least squares.

```
liny <- fto$y
linxg <- fto$X[, "xg"]
linxa <- fto$X[, "xa"]
linxint <- fto$X[, "xg"] * fto$X[, "xa"]
ftodf <- list(liny = liny, linxg = linxg, linxa = linxa,
             linxint = linxint)
```

FTO Example: LS fit

```
ols.fit <- lm(liny ~ linxg + linxa + linxint, data = ftodf)
summary(ols.fit)
##
## Call:
## lm(formula = liny ~ linxg + linxa + linxint, data = ftodf)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -4.8008 -0.8844  0.2993  1.2270  2.4819
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept) -0.06822    1.42230  -0.048  0.9623
## linxg        2.94485    2.01143   1.464  0.1625
## linxa        2.84421    0.42884   6.632 5.76e-06 ***
## linxint      1.72948    0.60647   2.852  0.0115 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 1.918 on 16 degrees of freedom
## Multiple R-squared:  0.9393, Adjusted R-squared:  0.9279
## F-statistic: 82.55 on 3 and 16 DF,  p-value: 5.972e-10
```


INLA

Integrated nested Laplace approximation (INLA) is a technique for carrying out Bayesian computation.

It is not a standard R package and must be downloaded from the development website.

The `inla` function is the work horse.

```
# install.packages('INLA',  
# repos='http://www.math.ntnu.no/inla/R/stable')  
library(INLA)  
# Data should be input to INLA as either a list or  
# a dataframe  
formula <- liny ~ linxg + linxa + linxint  
lin.mod <- inla(formula, data = ftodf, family = "gaussian")
```

We might wonder, where are the priors?

We didn't specify any... but INLA has default choices.

FTO example via INLA: Lots of output available!

```
names(lin.mod)
## [1] "names.fixed"           "summary.fixed"
## [3] "marginals.fixed"      "summary.lincomb"
## [5] "marginals.lincomb"    "size.lincomb"
## [7] "summary.lincomb.derived" "marginals.lincomb.derived"
## [9] "size.lincomb.derived" "mlik"
## [11] "cpo"                  "po"
## [13] "waic"                 "model.random"
## [15] "summary.random"      "marginals.random"
## [17] "size.random"         "summary.linear.predictor"
## [19] "marginals.linear.predictor" "summary.fitted.values"
## [21] "marginals.fitted.values" "size.linear.predictor"
## [23] "summary.hyperpar"    "marginals.hyperpar"
## [25] "internal.summary.hyperpar" "internal.marginals.hyperpar"
## [27] "offset.linear.predictor" "model.spde2.blc"
## [29] "summary.spde2.blc"   "marginals.spde2.blc"
## [31] "size.spde2.blc"     "model.spde3.blc"
## [33] "summary.spde3.blc"  "marginals.spde3.blc"
## [35] "size.spde3.blc"    "logfile"
## [37] "misc"               "dic"
## [39] "mode"              "neffp"
## [41] "joint.hyper"      "nhyper"
## [43] "version"          "Q"
## [45] "graph"            "ok"
## [47] "cpu.used"         "all.hyper"
## [49] ".args"           "call"
## [51] "model.matrix"
```

FTO example: INLA analysis

The posterior means and posterior standard deviations are in very close agreement with the OLS fits presented earlier.

```
coef(ols.fit)
## (Intercept)      linxg      linxa      linxint
## -0.06821632  2.94485495  2.84420729  1.72947648
sqrt(diag(vcov(ols.fit)))
## (Intercept)      linxg      linxa      linxint
##  1.4222970  2.0114316  0.4288387  0.6064695
lin.mod$summary.fixed
##              mean          sd 0.025quant    0.5quant 0.975quant
## (Intercept) -0.06158133  1.4304264 -2.8994511 -0.06200648  2.774212
## linxg        2.93317518  2.0204934 -1.0787238  2.93377097  6.934625
## linxa        2.84236014  0.4313641  1.9859120  2.84245094  3.696808
## linxint      1.73264070  0.6094298  0.5236601  1.73244076  2.940853
##              mode          kld
## (Intercept) -0.06259335  3.215535e-08
## linxg        2.93495278  2.555277e-08
## linxa        2.84264195  3.315511e-08
## linxint      1.73215893  2.751837e-08
```

FTO example: INLA analysis

```
summary(lin.mod)
##
## Call:
## "inla(formula = formula, family = \"gaussian\", data = ftdof)"
##
## Time used:
##   Pre-processing      Running inla Post-processing      Total
##           2.7230           0.3899           0.3089           3.4218
##
## Fixed effects:
##           mean      sd 0.025quant 0.5quant 0.975quant      mode kld
## (Intercept) -0.0616 1.4304   -2.8995  -0.0620    2.7742  -0.0626  0
## linxg         2.9332 2.0205   -1.0787   2.9338    6.9346   2.9350  0
## linxa         2.8424 0.4314    1.9859   2.8425    3.6968   2.8426  0
## linxint       1.7326 0.6094    0.5237   1.7324    2.9409   1.7322  0
##
## The model has no random effects
##
## Model hyperparameters:
##
##           mean      sd 0.025quant 0.5quant
## Precision for the Gaussian observations 0.3059 0.1017    0.1402    0.2947
##           0.975quant      mode
## Precision for the Gaussian observations    0.5359 0.2719
##
## Expected number of effective parameters(std dev): 3.995(0.0018)
## Number of equivalent replicates : 5.006
```

FTO Posterior marginals

We now examine the posterior marginal distributions.

The posterior marginal distribution for the vector of regression coefficients (including the intercept) is given below, and then we examine the posterior marginal on the precision, $1/\sigma_\epsilon$.

Check out the files that are written.

```
plot(lin.mod, plot.hyperparameter = FALSE, plot.fixed.effects = F  
      prefix = "linmodplot2", postscript = T)  
plot(lin.mod, plot.hyperparameter = TRUE, plot.fixed.effects = F  
      prefix = "linmodplot1", postscript = T)
```

FTO example via INLA

In order to carry out model checking we rerun the analysis, but now switch on a flag to obtain fitted values.

```
lin.mod <- inla(liny ~ linxg + linxa + linxint, data = ftodf,  
  family = "gaussian", control.predictor = list(compute = TRUE))  
fitted <- lin.mod$summary.fitted.values[, 1]  
# Now extract the posterior median of the  
# measurement error sd  
sigmamed <- 1/sqrt(lin.mod$summary.hyperpar[, 4])
```

FTO: Residual analysis

With the fitted values we can examine the fit of the model. In particular:

- ▶ Normality of the errors (sample size is relatively small).
- ▶ Errors have constant variance (and are uncorrelated).

FTO Residual analysis

The code below forms residuals and then forms


- ▶ a QQ plot to assess normality,
- ▶ a plot of residuals versus age, to assess linearity,
- ▶ a plot of residuals versus fitted values, to see if an unmodeled mean-variance relationship) and
- ▶ a plot of fitted versus observed for an overall assessment of fit.

FTO: Residual analysis

```
residuals <- (liny - fitted)/sigmamed
par(mfrow = c(2, 2))
qqnorm(residuals, main = "")
abline(0, 1, lty = 2, col = "red")
plot(residuals ~ linya, ylab = "Residuals", xlab = "Age")
abline(h = 0, lty = 2, col = "red")
plot(residuals ~ fitted, ylab = "Residuals", xlab = "Fitted")
abline(h = 0, lty = 2, col = "red")
plot(fitted ~ liny, xlab = "Observed", ylab = "Fitted")
abline(0, 1, lty = 2, col = "red")
```


The model assumptions do not appear to be greatly invalidated here.

Sample Quantiles




Theoretical Quantiles

Residuals




Age

Residuals



Fitted

Fitted



Observed

Case-Control Example: INLA Analysis

We perform two analyses.

The first analysis uses the default priors in INLA (which are relatively flat).

```
x <- c(0, 1, 2)
y <- c(6, 8, 75)
z <- c(10, 66, 163)
cc.dat <- as.data.frame(rbind(y, z, x))
cc.mod <- inla(y ~ x, family = "binomial", data = cc.dat,
  Ntrials = y + z)
summary(cc.mod)
##
## Call:
## c("inla(formula = y ~ x, family = \"binomial\", data = cc.dat, Ntrials = y +
##
## Time used:
##   Pre-processing      Running inla Post-processing      Total
##           2.3800           0.2447           0.2232           2.8479
##
## Fixed effects:
##           mean      sd 0.025quant 0.5quant 0.975quant      mode kld
## (Intercept) -1.8076 0.4549      -2.7496  -1.7905  -0.9629  -1.7566  0
## x             0.4803 0.2503       0.0099   0.4727   0.9936   0.4577  0
##
## The model has no random effects
```

Case-Control Example: Stan Analysis

Analysis with default priors: uses code in file
LogisticExample.stan

```
library(rstan)
stanlogist <- stan("LogisticExample.stan",
  data = list(x = c(0, 1, 2), y = c(6,
    8, 75), n = c(16, 74, 238)), iter = 1000,
  chains = 3, seed = 1234)
```


Case-Control Example: Stan Analysis

Close agreement with INLA analysis

```
summary(stanlogist)
## $summary
##           mean      se_mean      sd      2.5%      25%
## beta0    -1.8490348  0.02303110  0.4531242  -2.81385331  -2.1468026
## beta1     0.4968895  0.01271371  0.2490170   0.03327967   0.3310994
## lp__    -190.7326650  0.04582565  1.0302145  -193.64647207 -191.0619379
##           50%      75%      97.5%      n_eff      Rhat
## beta0    -1.8159325  -1.5536910  -1.009631  387.0839  1.004305
## beta1     0.4791666   0.6598048   1.045553  383.6308  1.004206
## lp__    -190.4109834 -190.0056870 -189.754643 505.4033  1.005734
##
## $c_summary
## , , chains = chain:1
##
##           stats
## parameter      mean      sd      2.5%      25%      50%
## beta0    -1.8358007  0.4856247  -2.91991418  -2.1687854  -1.7921155
## beta1     0.4896008  0.2645468   0.02720093   0.3098106   0.4658704
## lp__    -190.8097762  1.0555923  -193.51522775 -191.2157142 -190.4730521
##           stats
## parameter      75%      97.5%
## beta0    -1.5286527  -0.929760
## beta1     0.6669322   1.065317
## lp__    -190.0423988 -189.769143
##
## , , chains = chain:2
##
##           stats
## parameter      mean      sd      2.5%      25%      50%
## beta0    -1.8711435  0.4388380  -2.74588342  -2.1732235  -1.8347331
## beta1     0.5095493  0.2469336   0.04767094   0.3435566   0.4824003
## lp__    -190.6891500  1.0515562  -193.88374525 -190.9973433 -190.3977556
##           stats
```


Case-Control Example: Stan Analysis

```
traceplot(stanlogist, pars = c("beta1"), inc_warmup = TRUE)
```



Case-Control Example: Stan Analysis

```
plot(stanlogist, color = "green")
```



Prior choice


Suppose that for the odds ratio e^β we believe there is a 50% chance that the odds ratio is less than 1 and a 95% chance that it is less than 5; with $q_1 = 0.5$, $\theta_1 = 1.0$ and $q_2 = 0.95$, $\theta_2 = 5.0$, we obtain lognormal parameters $\mu = 0$ and $\sigma = (\log 5)/1.645 = 0.98$.

There is a function in the `SpatialEpi` package to find the parameters, as we illustrate.

```
library(SpatialEpi)
lnprior <- LogNormalPriorCh(1, 5, 0.5, 0.95)
lnprior
## $mu
## [1] 0
##
## $sigma
## [1] 0.9784688
```


Prior choice

```
plot(seq(0, 7, 0.1), dlnorm(seq(0, 7, 0.1), meanlog = lnprior$mu,  
  sdlog = lnprior$sigma), type = "l", xlab = "x",  
  ylab = "LogNormal Density")
```



Case-Control Example: Stan Analysis with Informative Prior

```
# Now with informative priors
W <- LogNormalPriorCh(1, 1.5, 0.5, 0.975)$sigma^2
cc.mod2 <- inla(y ~ x, family = "binomial", data = cc.dat,
  Ntrials = y + z, control.fixed = list(mean.intercept = c(0),
    prec.intercept = c(0.1), mean = c(0), prec = c(1/W)))
summary(cc.mod2)
##
## Call:
## c("inla(formula = y ~ x, family = \"binomial\", data = cc.dat, Ntrials = y +
##
## Time used:
##   Pre-processing      Running inla Post-processing      Total
##           2.5185           0.2891           0.1766           2.9842
##
## Fixed effects:
##           mean      sd 0.025quant 0.5quant 0.975quant      mode kld
## (Intercept) -1.3227 0.2896   -1.9013  -1.3193   -0.7639  -1.3124   0
## x            0.1986 0.1536   -0.1000   0.1975    0.5026   0.1955   0
##
## The model has no random effects
##
## The model has no hyperparameters
##
## Expected number of effective parameters(std dev): 1.441(0.00)
```

Case-Control Example: Stan Analysis with Informative Prior

```
library(rstan)
stanlogist2 <- stan("LogisticExamplePriors.stan",
  data = list(x = c(0, 1, 2), y = c(6,
    8, 75), n = c(16, 74, 238)), iter = 1000,
  chains = 3, seed = 2345)
```

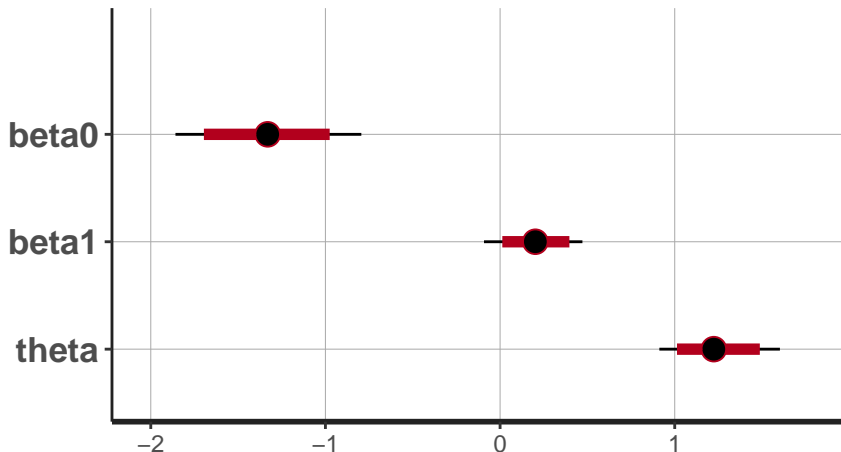
Case-Control Example: Stan Analysis with Informative Prior

Again close agreement with INLA analysis

```
summary(stanlogist2)
## $summary
##          mean      se_mean      sd          2.5%          25%
## beta0    -1.3359973  0.016028376  0.2772953   -1.85844538   -1.545669
## beta1     0.2032825  0.008312775  0.1470454    -0.09203033    0.100697
## theta    1.2386478  0.010231689  0.1813273    0.91207749    1.105942
## lp__    -191.8628757  0.038270460  0.8872137  -194.23853282  -192.223118
##          50%          75%          97.5%    n_eff      Rhat
## beta0    -1.329673   -1.1480507   -0.7944808  299.2995  1.010971
## beta1     0.202189    0.3081545    0.4713988  312.9039  1.009105
## theta    1.224079    1.3609113    1.6022340  314.0739  1.009217
## lp__    -191.601556  -191.2268363  -190.9636514  537.4389  1.001055
##
## $c_summary
## , , chains = chain:1
##
##          stats
## parameter      mean      sd          2.5%          25%          50%
## beta0    -1.3252153  0.2813771   -1.82581135   -1.5305189   -1.3346949
## beta1     0.2009809  0.1520604   -0.09485358    0.1017572    0.1932956
## theta    1.2366177  0.1855242    0.90950938    1.1071148    1.2132414
## lp__    -191.8643627  0.8963543  -194.30850108  -192.2225769  -191.5839966
##          stats
## parameter      75%          97.5%
## beta0    -1.1338300   -0.7655444
## beta1     0.3059642    0.4756200
## theta    1.3579340    1.6090120
## lp__    -191.2246514  -190.9652067
##
## , , chains = chain:2
```

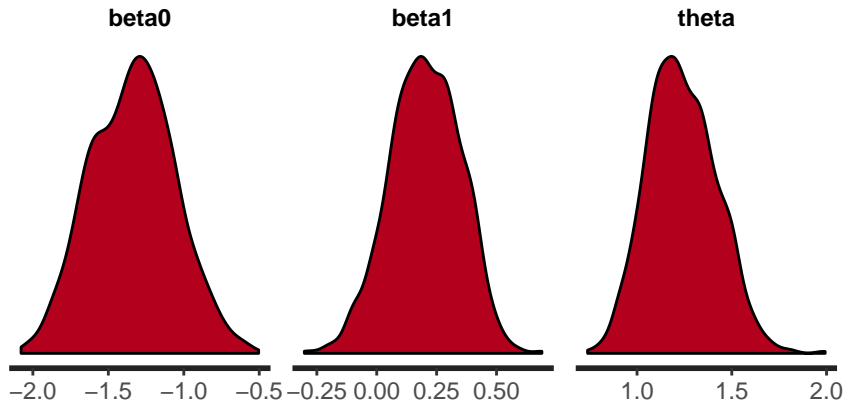
Case-Control Example: Stan Analysis with Informative Prior

```
plot(stanlogist2, color = "green", parameter = "theta")
```



Case-Control Example: Stan Analysis with Informative Prior

```
stan_dens(stanlogist2)
```



A Simple ANOVA Example

We begin with simulated data from the simple one-way ANOVA model example:

$$\begin{aligned} Y_{ij} | \beta_0, b_i &= \beta_0 + b_i + \epsilon_{ij} \\ \epsilon_{ij} | \sigma_\epsilon^2 &\sim_{iid} \text{normal}(0, \sigma_\epsilon^2) \\ b_i | \sigma_b^2 &\sim_{iid} \text{normal}(0, \sigma_b^2) \end{aligned}$$

$i = 1, \dots, 10; j = 1, \dots, 5$, with $\beta_0 = 0.5$, $\sigma_\epsilon^2 = 0.2^2$ and $\sigma_b^2 = 0.3^2$.

b_i are random effects and ϵ_{ij} are **measurement errors** and there are two variances to estimate, σ_ϵ^2 and σ_b^2 .

In a fixed effects Bayesian model, the variance σ_b^2 would be fixed in advance.

A Simple ANOVA Example

Simulation is described and analyzed below. We fit the one-way ANOVA model and see reasonable recovery of the true values that were used to simulate the data.

Not a big surprise, since we have fitted the model that was used to simulate the data!

```
set.seed(133)
sigma.b <- 0.3
sigma.e <- 0.2
m <- 10
ni <- 5
beta0 <- 0.5
b <- rnorm(m, mean = 0, sd = sigma.b)
e <- rnorm(m * ni, mean = 0, sd = sigma.e)
Yvec <- beta0 + rep(b, each = ni) + e
simdata <- data.frame(y = Yvec, ind = rep(1:m, each = ni))
result <- inla(y ~ f(ind, model = "iid"), data = simdata)
sigma.est <- 1/sqrt(result$summary.hyperpar[, 4])
sigma.est
## [1] 0.1927622 0.2455392
```

`sigma.est` corresponds to the posterior medians of σ_ϵ and σ_b , respectively.

A Simple ANOVA Example

```
summary(result)
##
## Call:
## "inla(formula = y ~ f(ind, model = \"iid\"), data = simdata)"
##
## Time used:
##   Pre-processing      Running inla Post-processing      Total
##           3.1749           0.3259           0.2812           3.7820
##
## Fixed effects:
##           mean      sd 0.025quant 0.5quant 0.975quant  mode kld
## (Intercept) 0.525 0.088      0.3491   0.525      0.7006 0.525  0
##
## Random effects:
## Name      Model
## ind      IID model
##
## Model hyperparameters:
##
##           mean      sd 0.025quant 0.5quant
## Precision for the Gaussian observations 27.42 6.026      17.184   26.91
## Precision for ind                       18.45 9.368      6.048   16.59
##           0.975quant  mode
## Precision for the Gaussian observations      40.72 25.98
## Precision for ind                          41.78 13.15
##
## Expected number of effective parameters(std dev): 8.892(0.587)
```

RNA Seq Analysis

We fit a hierarchical logistic regression model starting with first stage:

$$Y_{ij} | N_{ij}, p_{ij} \sim \text{binomial}(N_{ij}, p_{ij})$$

so that p_{ij} is the probability of seeing an A read for gene i and replicate j , $i = 1, \dots, 10$, $j = 1, 2$.

Then the odds of an A read is $\frac{p_{ij}}{1-p_{ij}}$.

At the second stage:

$$\text{logit } p_{ij} = \theta_i + \epsilon_{ij}$$

where $\epsilon_{ij} | \sigma^2 \sim \text{normal}(0, \sigma^2)$ represent random effects that allow for excess-binomial variation; there are a pair for each gene.

The θ_i parameters are taken as fixed effects with a relatively flat prior (the default choice in INLA).

$\exp(\theta_i)$ is the odds of seeing an A read for gene i .

RNA Seq Analysis

Rows 1 and 2 represent the two replicates for gene 1, rows 3 and 4 represent the two replicates for gene 2, etc. . .

rep1 is the variable that defines the random effects.

xvar is the gene number, there are 10 genes in this dataset.

```
rnay <- c(1963, 3676, 249, 110, 78, 92, 1585, 798,
          2525, 2620, 598, 473, 120, 72, 1496, 1291, 397,
          480, 242, 174)
rnan <- c(7617, 10413, 308, 114, 161, 153, 2321, 1527,
          4142, 3861, 910, 778, 160, 85, 2795, 2697, 810,
          928, 466, 313)
rnarep1 <- seq(1, 20)
rnaxvar <- rep(1:10, each = 2)
RNAdat <- as.data.frame(cbind(rnay, rnan, rnarep1,
                              rnaxvar))
names(RNAdat) <- c("y", "n", "rep1", "xvar")
head(RNAdat, n = 5)
##      y      n rep1 xvar
## 1 1963  7617     1     1
## 2 3676 10413     2     1
## 3  249   308     3     2
## 4  110   114     4     2
## 5   78   161     5     3
```

RNA Seq Analysis

Below is the code for fitting the random effects model.

The -1 in the model specification removes the intercept, so that the factor levels are defined with one level for each gene.

```
RNAfit <- inla(y ~ as.factor(xvar) - 1 + f(rep1, model = "iid"),  
              family = "binomial", data = RNAdat, Ntrials = n)
```

RNA Seq Analysis

```
RNAfit$summary.fixed[1:10, c(1:5)]
```

```
##                mean          sd 0.025quant    0.5quant 0.975quant
## as.factor(xvar)1 -0.83151915 0.2409771 -1.31848001 -0.83148290 -0.3454449
## as.factor(xvar)2  2.02846178 0.3062303  1.46843071  2.01201739  2.6830520
## as.factor(xvar)3  0.17275894 0.2660097 -0.35821961  0.17248642  0.7047302
## as.factor(xvar)4  0.43012608 0.2427891 -0.06034110  0.43029959  0.9190930
## as.factor(xvar)5  0.59638990 0.2415450  0.10858844  0.59635990  1.0836064
## as.factor(xvar)6  0.54547388 0.2456901  0.05026705  0.54549687  1.0399062
## as.factor(xvar)7  1.36397585 0.2921559  0.79618494  1.35970461  1.9562787
## as.factor(xvar)8  0.02795368 0.2419176 -0.46055161  0.02794724  0.5157758
## as.factor(xvar)9  0.01491012 0.2451661 -0.47934489  0.01492469  0.5084299
## as.factor(xvar)10 0.14954910 0.2513086 -0.35500469  0.14933012  0.6546026
```

```
RNAfit$summary.hyperpar[c(1:5)]
```

```
##                mean          sd 0.025quant 0.5quant 0.975quant
## Precision for rep1 11.71229 6.704229  3.109531 10.25667  28.63479
```


RNA Seq Analysis

```
RNAfit$summary.random$rep1[, c(2:6)]
```

| ## | mean | sd | 0.025quant | 0.5quant | 0.975quant |
|-------|-------------|-----------|-------------|-------------|------------|
| ## 1 | -0.22463992 | 0.2409749 | -0.71205178 | -0.22441497 | 0.2610189 |
| ## 2 | 0.22454380 | 0.2409747 | -0.26177977 | 0.22430998 | 0.7111898 |
| ## 3 | -0.47628401 | 0.3030389 | -1.14374319 | -0.45205994 | 0.0552154 |
| ## 4 | 0.47652319 | 0.3030982 | -0.05618667 | 0.45215783 | 1.1445510 |
| ## 5 | -0.18449556 | 0.2614869 | -0.72230085 | -0.17924091 | 0.3259034 |
| ## 6 | 0.18451558 | 0.2614880 | -0.32502569 | 0.17926856 | 0.7224917 |
| ## 7 | 0.32953374 | 0.2427763 | -0.15720438 | 0.32820528 | 0.8220183 |
| ## 8 | -0.32948406 | 0.2427759 | -0.82260151 | -0.32818268 | 0.1569429 |
| ## 9 | -0.14883708 | 0.2415324 | -0.63735107 | -0.14855760 | 0.3376903 |
| ## 10 | 0.14890601 | 0.2415325 | -0.33819573 | 0.14859410 | 0.6367141 |
| ## 11 | 0.10000372 | 0.2454109 | -0.39263430 | 0.09916682 | 0.5959377 |
| ## 12 | -0.09994075 | 0.2454106 | -0.59666366 | -0.09913388 | 0.3923015 |
| ## 13 | -0.19098338 | 0.2782504 | -0.77186307 | -0.18176407 | 0.3402557 |
| ## 14 | 0.19114195 | 0.2782662 | -0.34033130 | 0.18194755 | 0.7727922 |
| ## 15 | 0.11135542 | 0.2418948 | -0.37640785 | 0.11105393 | 0.5997830 |
| ## 16 | -0.11135219 | 0.2418948 | -0.60050656 | -0.11107368 | 0.3758144 |
| ## 17 | -0.05148316 | 0.2449316 | -0.54629711 | -0.05110711 | 0.4408796 |
| ## 18 | 0.05148488 | 0.2449316 | -0.44141390 | 0.05108651 | 0.5455838 |
| ## 19 | -0.06566901 | 0.2502098 | -0.57168455 | -0.06467751 | 0.4348675 |
| ## 20 | 0.06568634 | 0.2502104 | -0.43517831 | 0.06467192 | 0.5708872 |


RNA Seq Analysis

PostDens [as.factor(xvar)1]




Mean = -0.832 SD = 0.245

PostDens [as.factor(xvar)2]




Mean = 2.031 SD = 0.31

PostDens [as.factor(xvar)3]




Mean = 0.173 SD = 0.27

PostDens [as.factor(xvar)4]




Mean = 0.43 SD = 0.247

PostDens [as.factor(xvar)5]




Mean = 0.596 SD = 0.245

PostDens [as.factor(xvar)6]




Mean = 0.546 SD = 0.25

PostDens [as.factor(xvar)7]




Mean = 1.365 SD = 0.295

PostDens [as.factor(xvar)8]



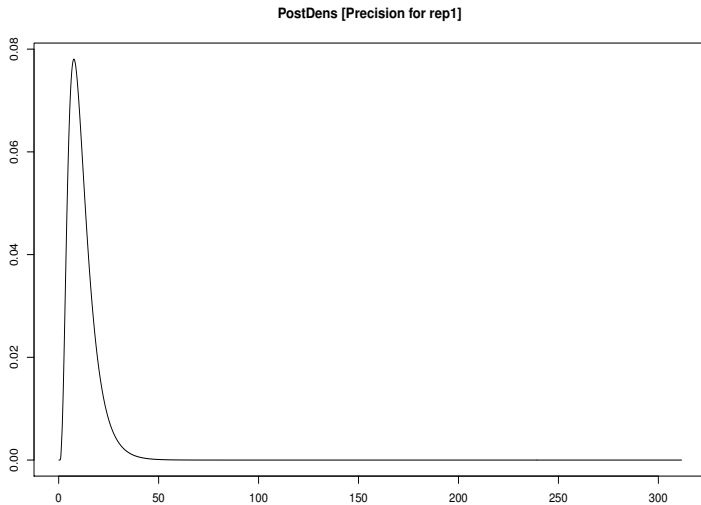
Mean = 0.028 SD = 0.246

PostDens [as.factor(xvar)9]




Mean = 0.015 SD = 0.249

RNA Seq Analysis



RNA Seq Analysis



RNA Seq Analysis

We extract the 95% intervals and posterior medians for the log odds of being an A allele.


Comparison with 0 gives an indication of cis effects.

Genes 1, 2, 5, 6, 7 show evidence of cis effects.

RNA Seq Analysis

```
thetasum <- RNAfit$summary.fixed[, 3:5]
par(mfrow = c(1, 1))
plot(thetasum[, 2], seq(1, 10), xlim = c(min(thetasum),
    max(thetasum)), ylab = "Gene", xlab = "Log Odds")
for (i in 1:10) {
  lines(x = c(thetasum[i, 1], thetasum[i, 3]), y = c(i,
    i))
}
abline(v = 0, col = "red")
# Intervals to the left/right of this line?
```

RNA Seq Analysis



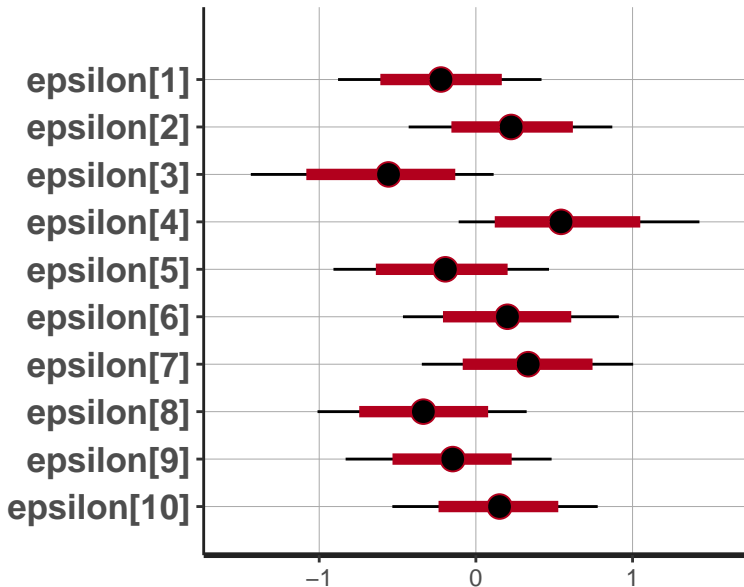
Logistic regression with random effects: Stan analysis

```
N <- 20
rny <- c(1963, 3676, 249, 110, 78, 92, 1585,
        798, 2525, 2620, 598, 473, 120, 72, 1496,
        1291, 397, 480, 242, 174)
rnan <- c(7617, 10413, 308, 114, 161, 153,
        2321, 1527, 4142, 3861, 910, 778, 160,
        85, 2795, 2697, 810, 928, 466, 313)
rnax <- c(1, 1, 2, 2, 3, 3, 4, 4, 5, 5, 6,
        6, 7, 7, 8, 8, 9, 9, 10, 10)
rnadat3 <- list(N = N, rny = rny, rnan = rnan,
               rnax = rnax)
stanlogist3 <- stan("LogisticExampleRandomEffects.stan",
                  data = rnadat3, iter = 3000, chains = 3,
                  seed = 3456)
```

Logistic regression with random effects: Stan analysis

```
## $summary
##                mean      se_mean      sd      2.5%
## epsilon[1] -2.246395e-01 0.004829035 0.3251320 -8.793885e-01
## epsilon[2]  2.252049e-01 0.004787228 0.3243642 -4.289727e-01
## epsilon[3] -5.849420e-01 0.006657290 0.3920394 -1.436278e+00
## epsilon[4]  5.726339e-01 0.005256640 0.3832490 -1.098250e-01
## epsilon[5] -2.022849e-01 0.005164586 0.3487287 -9.088247e-01
## epsilon[6]  2.022792e-01 0.004973792 0.3449540 -4.653313e-01
## epsilon[7]  3.306600e-01 0.005096235 0.3414554 -3.449760e-01
## epsilon[8] -3.354515e-01 0.005179654 0.3417241 -1.010461e+00
## epsilon[9] -1.538362e-01 0.004727868 0.3236196 -8.307891e-01
## epsilon[10] 1.456187e-01 0.004695782 0.3235445 -5.331445e-01
## epsilon[11] 1.053189e-01 0.005075729 0.3273046 -5.339340e-01
## epsilon[12] -9.717759e-02 0.004983204 0.3268193 -7.423470e-01
## epsilon[13] -2.244291e-01 0.004469825 0.3525138 -9.597766e-01
## epsilon[14]  2.286210e-01 0.004619784 0.3571658 -4.561853e-01
## epsilon[15] 1.064358e-01 0.005185528 0.3338578 -5.737402e-01
## epsilon[16] -1.176573e-01 0.005095026 0.3351175 -7.795833e-01
## epsilon[17] -6.007708e-02 0.004771030 0.3174358 -7.123277e-01
## epsilon[18]  4.580314e-02 0.004747041 0.3166734 -5.926365e-01
## epsilon[19] -5.977598e-02 0.004560714 0.3337353 -7.299581e-01
## epsilon[20]  7.260294e-02 0.004563275 0.3334286 -6.041017e-01
## sigma
## 4.341919e-01 0.004633383 0.1520200  2.307855e-01
## beta[1] -8.320365e-01 0.004812352 0.3248013 -1.474976e+00
## beta[2]  2.116041e+00 0.005969147 0.3885023  1.396327e+00
## beta[3]  1.750840e-01 0.005046756 0.3510150 -5.177243e-01
## beta[4]  4.315573e-01 0.005127250 0.3428660 -2.412228e-01
## beta[5]  6.008675e-01 0.004715101 0.3236424 -3.618104e-02
## beta[6]  5.421669e-01 0.005032589 0.3274446 -1.093177e-01
## beta[7]  1.381062e+00 0.004454887 0.3587102  6.976079e-01
## beta[8]  3.372643e-02 0.005167797 0.3339123 -6.177861e-01
## beta[9]  2.274081e-02 0.004741969 0.3171578 -5.962635e-01
## beta[10]  1.472542e-01 0.004555396 0.3346755 -5.306240e-01
## lp__ -2.596686e+04 0.177189172 5.3173554 -2.597838e+04
##                25%                50%                75%                97.5%
```

Logistic regression with random effects: Stan analysis



Approximate Bayes

We return to the case control example seen earlier.

Below we construct the posterior by hand

```
x <- c(0, 1, 2)
y <- c(6, 8, 75)
z <- c(10, 66, 163)
logitmod <- glm(cbind(y, z) ~ x, family = "binomial")
thetahat <- logitmod$coef[2]
V <- vcov(logitmod)[2, 2]
# 97.5 point of prior is log(1.5) so that we with
# prob 0.95 we think theta lies in (2/3, 1.5)
W <- LogNormalPriorCh(1, 1.5, 0.5, 0.975)$sigma^2
```


Approximate Bayes: estimation

```
r <- W/(V + W)
r
## [1] 0.4055539
# Not so much data here, so weight on prior is
# high. Bayesian posterior median
exp(r * thetahat)
##          x
## 1.214286
# Shrunk towards prior median of 1 Note: INLA
# estimate (with same prior) is 1.22 and
# approximate posterior SD here is sqrt(rV)=0.159,
# INLA version is 0.154. Bayesian approximate 95%
# credible interval
exp(r * thetahat - 1.96 * sqrt(r * V))
##          x
## 0.8882832
exp(r * thetahat + 1.96 * sqrt(r * V))
##          x
## 1.659932
```

Approximate Bayes: hypothesis testing

Now we turn to testing using Bayes factors.

We examine the sensitivity to the prior on the alternative, π_1 .

```
pi1 <- c(1/2, 1/100, 1/1000, 1/10000, 1/1e+05) # 5 prior probs on the null
source("http://faculty.washington.edu/jonno/BFDP.R")
BFcall <- BFDPfunV(thetahat, V, W, pi1)
BFcall
## $BF
##          x
## 0.6182773
##
## $pH0
##          x
## 0.256323
##
## $pH1
##          x
## 0.4145761
##
## $BFDP
## [1] 0.3820589 0.9839253 0.9983836 0.9998383 0.9999838
```

So data are twice as likely under the alternative (0.502) as compared to the null (0.256).