SISG Bayes: Session 4 INLA example

Ken Rice

2026-05-27

Implementing the LHON example via INLA

```
lhon <- data.frame( y=c(1,1,1,0,0,0), x=c(0,1,2,0,1,2), n=c(6,8,75,10,66,163) )
dmat <- rbind(lhon$n[4:6], lhon$n[1:3])
names(dmat) <- c("CC","CT","TT")

barplot(dmat, names.arg=c("CC","CT","TT"), legend.text=c("case","control")[2:1],
args.legend=list(x="topleft", bty="n"))
```

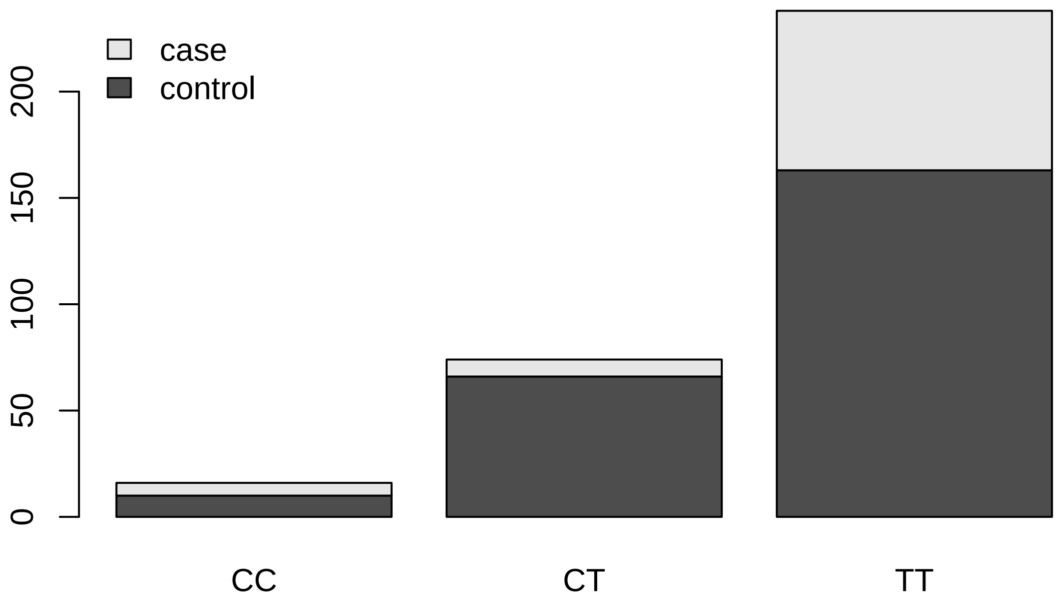

Setting this up for a non-Bayesian analysis, fitting a GLM:

```
# setup data
cc.dat <- data.frame(x=c(0,1,2), success=c(6,8,75), fail=c(10,66,163))
```

```r
# estimates and confidence interval from GLM analysis
logitmod <- glm(cbind(success,fail)~x,family="binomial", data=cc.dat)
coef(summary(logitmod))
```

```
##               Estimate Std. Error   z value     Pr(>|z|)
## (Intercept) -1.8076928  0.4553938 -3.969515 7.201898e-05
## x            0.4787428  0.2504594  1.911459 5.594568e-02
```

```r
confint.default(logitmod)
```

```
##                   2.5 %     97.5 %
## (Intercept) -2.70024830 -0.9151373
## x           -0.01214865  0.9696342
```

Now using INLA – note the similar syntax, and for now using its default priors:

```r
library("INLA")
```

```
## Loading required package: Matrix

## Warning in file(con, "r"): URL 'https://inla.r-inla-download.org/VERSIONS':
## status was 'Could not resolve hostname'

## This is INLA_25.10.19 built 2025-10-19 19:10:20 UTC.
##  - See www.r-inla.org/contact-us for how to get help.
##  - List available models/likelihoods/etc with inla.list.models()
##  - Use inla.doc(<NAME>) to access documentation
```

```r
cc.mod <- inla(success~x,family="binomial",data=cc.dat,Ntrials=success+fail)
summary(cc.mod)
```

```
## Time used:
##     Pre = 0.697, Running = 0.293, Post = 0.0632, Total = 1.05
## Fixed effects:
##               mean    sd 0.025quant 0.5quant 0.975quant   mode kld
## (Intercept) -1.852 0.455     -2.744   -1.852     -0.959 -1.852   0
## x            0.499 0.250      0.009    0.499      0.990  0.499   0
##
## Marginal log-Likelihood:  -17.89
##  is computed
## Posterior summaries for the linear predictor and the fitted values are computed
## (Posterior marginals needs also 'control.compute=list(return.marginals.predictor=TRUE)')
```

Now implementing prior $N(0, 10)$ for the intercept (i.e. log odds for $X = 0$) and – informatively – $N(0, \sigma^2)$ for the log odds ratio, with σ such that the upper 95% point of the distribution is at $OR = 5$

```r
Upper95 <- log(5) # note log(5) = 1.61
sigma <- Upper95/qnorm(0.95)
cc.inla <- inla(success~x,family="binomial",data=cc.dat,Ntrials=success+fail,
   control.fixed=list(mean.intercept=c(0),prec.intercept=c(1/10),
                      mean=c(0),prec=c(1/sigma^2)))
summary(cc.inla)
```

```
## Time used:
##     Pre = 0.394, Running = 0.292, Post = 0.0403, Total = 0.725
## Fixed effects:
##               mean    sd 0.025quant 0.5quant 0.975quant   mode kld
## (Intercept) -1.760 0.431     -2.605   -1.760     -0.916 -1.760   0
## x            0.449 0.237     -0.016    0.449      0.914  0.449   0
##
## Marginal log-Likelihood:  -16.81
##  is computed
## Posterior summaries for the linear predictor and the fitted values are computed
## (Posterior marginals needs also 'control.compute=list(return.marginals.predictor=TRUE)')
```

For a more informative prior, repeat but set the upper 97.5% point of the prior to $OR = 1.5$:

```
Upper975 <- 1.5 # note log(1.5)=0.41
sigma <- log(Upper975)/qnorm(0.975)
cc.inf.inla <- inla(success~x,family="binomial",data=cc.dat,Ntrials=success+fail,
   control.fixed=list(mean.intercept=c(0),prec.intercept=c(1/10),
                      mean=c(0),prec=c(1/sigma^2)))
summary(cc.inf.inla)
```

```
## Time used:
##     Pre = 0.575, Running = 0.254, Post = 0.0334, Total = 0.862
## Fixed effects:
##               mean    sd 0.025quant 0.5quant 0.975quant   mode kld
## (Intercept) -1.332 0.290     -1.899   -1.332     -0.764 -1.332   0
## x            0.201 0.154     -0.100    0.201      0.502  0.201   0
##
## Marginal log-Likelihood:  -16.64
##  is computed
## Posterior summaries for the linear predictor and the fitted values are computed
## (Posterior marginals needs also 'control.compute=list(return.marginals.predictor=TRUE)')
```

Some graphs to compare them, on the log odds ratio scale:

```
#pdf("inlalhonpost.pdf", w=8, h=4.5)
par(mar=c(4,4,0,0)+0.2)
curve(inla.dmarginal(x, cc.inf.inla$marginals.fixed$x), -1,1.5,
    col="forestgreen", ylab="density", lty=2, xlab=expression(beta[1]))
curve(inla.dmarginal(x, cc.inla$marginals.fixed$x), -1,1.5,
    col="forestgreen", ylab="density", lty=1, add=TRUE)
curve(dnorm(x,0,log(Upper975)/qnorm(0.975)), -1, 1.5, col=4, lty=2, add=TRUE)
curve(dnorm(x,0,Upper95/qnorm(0.95)), -1, 1.5, col=4, lty=1, add=TRUE)
lines(y=c(0,0), x=confint.default(logitmod,2), lwd=2, col="goldenrod")
points(y=0, x=coef(logitmod)[2], lwd=2, col="goldenrod")
legend("topleft", lty=1:2, c("diffuse","informative"), bty="n")
legend("topright", lty=1, bty="n", col=c(4,"forestgreen","goldenrod"), c("prior","posterior","non-Bayes
```

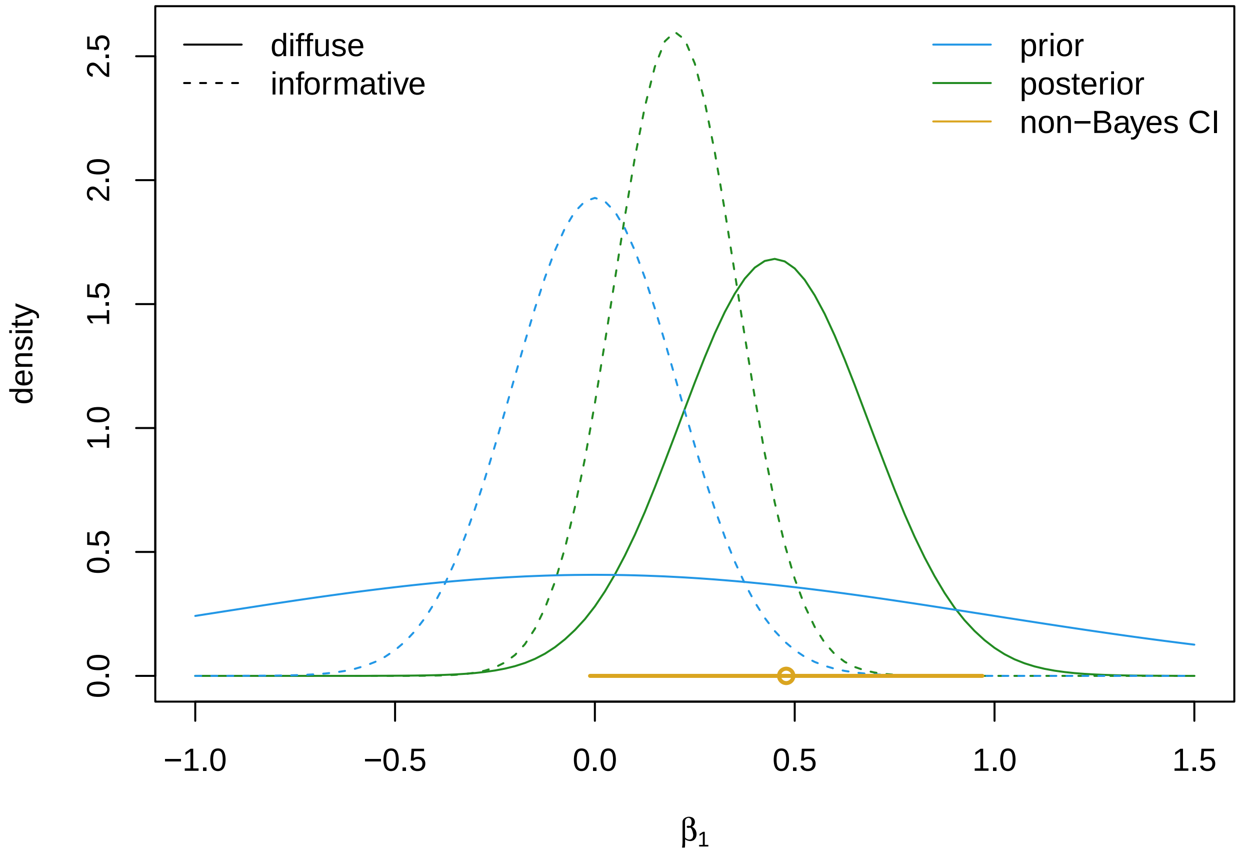

```
#dev.off()
```

Add on the odds ratio scale:

```
#pdf("inlalhonpostexp.pdf", w=8, h=4.5)
par(mar=c(4,4,0,0)+0.2)
plot( inla.tmarginal(exp, cc.inla$marginals.fixed$x) , ylab="density",
type="l", xlab=expression("Odds ratio, "*e^beta[1]), col="forestgreen",
ylim=c(0,2.2), xlim=c(0.5, 3.5))
lines( inla.tmarginal(exp, cc.inf.inla$marginals.fixed$x) ,
col="forestgreen", type="l", lty=2)
curve(dlnorm(x,0,log(Upper975)/qnorm(0.975)), exp(-1), exp(1.5), col=4, lty=2, add=TRUE)
curve(dlnorm(x,0,Upper95/qnorm(0.95)), exp(-1), exp(1.5), col=4, lty=1, add=TRUE)
lines(y=c(0,0), x=exp(confint.default(logitmod,2)), lwd=2, col="goldenrod")
points(y=0, x=exp(coef(logitmod)[2]), lwd=2, col="goldenrod")
legend("topright", lty=1:2, c("diffuse","informative"), bty="n")
legend("right", lty=1, bty="n", col=c(4,"forestgreen","goldenrod"), c("prior","posterior","non-Bayes CI
```

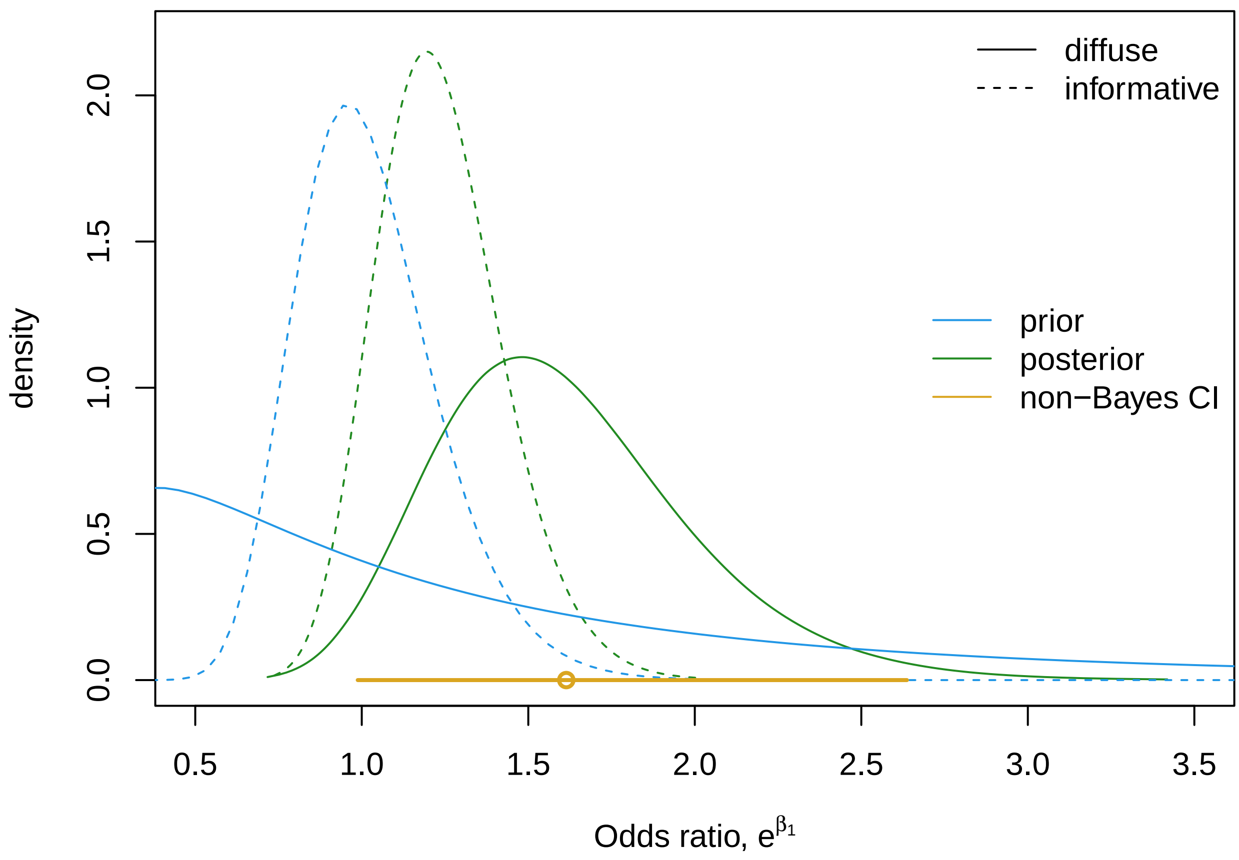
diffuse
informative
prior
posterior
non−Bayes CI
density
0.0
0.5
1.0
1.5
2.0
0.5
1.0
1.5
2.0
2.5
3.0
3.5
Odds ratio, e^{β_1}

```
#dev.off()
```