Statistics in Environmental Research (BUC Workshop Series) II
Problem sheet - WinBUGS - SOLUTIONS

1. (a) The posterior mean estimate of α is 14.27, and the posterior mean for the standard deviation of the random effects distribution $1/\sqrt{\alpha}$ is 0.27. When compared with the empirical Bayes estimates of 51.7 and 0.14 we see large differences, which is surprising given we have 88 areas. The explanation lies in the Ga(1,1) prior that was assumed. The 2.5%, 50% and 97.5% points of this distribution are 0.025, 0.69, 3.7, so that large values of α are strongly discouraged in the prior (the 2.5%, 50% and 97.5% points for the standard deviation are 0.52, 1.2, 6.3). Hence this prior will force between-area variability even when it is not present in the data – very dangerous!

   As an alternative we specify a lognormal prior for α. To decide on the cut-off points of this prior we assumed that residual relative risk standard deviations of 0.1 and 1 were in the left and right tails. Specifically we assumed that the 5% and 95% points of the prior for α were 1 and 100, giving a LogNormal(2.30,1.40) specification. The code for this model is given below (note the use of the `step` function).
   This gave a posterior mean for α of 41.2, and for $1/\sqrt{\alpha}$ of 0.16.

```
model
{
  for (i in 1 : N) {
      Obs[i]  ~ dpois(mu[i])
      mu[i] <- Exp[i]*exp(beta0)*theta[i]
      RR[i] <- exp(beta0)*theta[i]
      theta[i] ~ dgamma(alpha,alpha)
      thresh[i] <- step(RR[i]-1.2)
  }
  lalpha ~ dnorm(2.30,1.40)
  alpha <- exp(lalpha)
  beta0 ~ dflat()
# Functions of interest:
  sigma.theta <- sqrt(1/alpha)   # standard deviation of non-spatial
  base <- exp(beta0)
}
```

   (b) I copied all of the `WinBUGS` results into a single file whose first two lines are:

```
RR[1]   0.9229  0.1769  0.001729   0.6128  0.911   1.304   10000   10001
RR[2]   0.9027  0.108   0.001094   0.7035  0.8987  1.123   10000   10001
```

   The `R` commands below then produced the figures that follow.

```
out <- read.table("ex3q1out1.dat",sep="",header=F)
gamRRp1 <- out[1:88,2]
gamthreshp1 <- out[89:176,2]
gamRRp2 <- out[177:264,2]
gamthreshp2 <- out[265:352,2]
```

```
lnormRR <- out[353:440,2]
lnormthresh <- out[441:528,2]
OhioMap(gamRRp1,ncol=8,type="e",figmain="Ohio lung cancer Poisson gamma"
,lower=0,upper=max(SMR))
```

The empirical Bayes, Poisson-Gamma with a lognormal prior, and Poisson-Lognormal model give very similar estimates. The Poisson-Gamma model with the Ga(1,1) prior gives estimates that are much more variable, since the prior has discouraged a lack of global smoothing.

If we plot the posterior probabilities that the relative risk exceeds 1.2 we see that there is strong agreement between the maps, apart from the Ga(1,1) specification which shows higher probabilities in some areas.

(c) Under the Poisson-Lognormal model the standard deviation is estimated as 0.14, in agreement with the Poisson-Gamma model under the appropriate model.

A plot of the relative risk estimates under the Poisson-Gamma model (with the sensible prior), and the Poisson-Lognormal model shows the good agreement between the estimates. Under the lognormal model there is a narrower range, which is consistent with the standard deviation having a slightly smaller estimate.

We also defined the endpoints of a 95% interval for the relative risk using the code below in the `WinBUGS` model specification. The posterior means of these points were 0.76 and 1.32, again emphasizing that for these data there is little residual variability.

```
RRRlo <- exp(-1.96*sigma.V)
RRRhi <- exp(1.96*sigma.V)
```

2. (a)

```
> OhioMap(Obs/Exp,ncol=8,type="e",figmain="Ohio lung cancer SMRs")
> map.scale(x=-84.5,y=38.6,ratio=F)
> lnprior(50,200,.05,.95)
$mu
[1] 4.60517
$sigma
[1] 0.4214036
```

(b) Adjacency map was given on website.

(c) Posterior summaries under the two spatial models are given below, the proportion of the total variability on the log residual relative risk scale is estimated to be 0.56 and 0.41 under the joint and ICAR models respectively.

Note the large uncertainty about $d_{1/2}$ (95% interval is 3.3 to 1128). The joint model ran very slowly (due to inversion of 88×88 matrix at every iteration), and there is very high dependence in the Markov chain (reflected in large Monte Carlo error for $d_{1/2}$).

```
JOINT MODEL
beta0   0.02179 0.1631  0.01444 -0.2179 -0.0243 0.3951  5000    15001
dhalf   172.5   450.4   28.02   3.26    53.53   1128.0  5000    15001
p       0.5613  0.241   0.017   0.09528 0.5902  0.935   5000    15001
phi     0.03436 0.0728  0.00460 6.14E-4 0.01295 0.2126  5000    15001
```

```
sigma.U 0.155   0.08047 0.00585 0.04501 0.1388  0.3547  5000    15001
sigma.V 0.1209  0.03095 0.00157 0.05566 0.1214  0.1807  5000    15001

ICAR MODEL
beta0   -0.0368 0.02401 6.08E-4 -0.0849 -0.0363 0.00892 10000   10001
sd.U    0.09954 0.02655 0.00146 0.05214 0.09791 0.1542  10000   10001
sigma.V 0.1204  0.0254  9.76E-4 0.07658 0.1183  0.1746  10000   10001
vratio  0.4108  0.1628  0.00877 .1219   0.4042  0.7318  10000   10001
```

3. A small random amount has to be added to each location, to avoid having points at the same location, which leads to a variance-covariance matrix that is not invertible.

   This analysis is for illustration only – we would really like to analyze the complete data (though the computation would be very, very slow). We would like to sample more controls (in a 3 to 1 ratio, say) to get more power.

   The output from a logistic regression analysis is given below – not surprisingly, given the reduced sample size, the exposure effect is not significant.

```
mod <- glm(cc~exposure,family="binomial")
summary(mod)
Coefficients:
            Estimate Std. Error z value Pr(>|z|)
(Intercept) -0.18753    0.19278  -0.973    0.331
exposure     0.09428    0.06602   1.428    0.153
```

   A Bayesian version of this model with flat priors (and no random effects) gave the summaries below. Not surprisingly these are very similar to asymptotic likelihood inference.

   We give results from various random effects models to illustrate the sensitivity of inference to the prior – not surprising given binary data.

```
Flat priors and no random effects
node     mean    sd      MC error 2.5%   median  97.5%   start   sample
RRx      1.104   0.07326 0.001662 0.968  1.101   1.254   5000    5001
beta0   -0.1923 0.1933  0.004245 -0.573 -0.191   0.184   5000    5001
beta1    0.09631 0.06629 0.001506 -0.032 0.09599 0.226   5000    5001
```

   When we add random effects we fit a much tighter prior to the random effects precision, $\tau_v \sim \text{Ga}(10, 1.10)$ – this prior corresponds to believing that the residual odds fall between 0.5 and 2 with probability 0.95 and follow a log Student t distribution with 20 degrees of freedom.

```
NON-SPATIAL RANDOM EFFECTS with flat priors and random effects and flat
priors on beta
node     mean    sd      MC error 2.5%    median  97.5%   start  sample
RRx      1.106   0.07612 0.001869 0.9692  1.103   1.271   5000   5001
beta0   -0.1953 0.1999  0.004562 -0.5935 -0.1955 0.1971  5000   5001
beta1    0.09845 0.0685  0.001685 -0.0313 0.09777 0.2395  5000   5001
sdV      0.3431  0.05795 0.003585 0.2509  0.3369  0.4768  5000   5001
```

We see little change in the estimate and standard error for β_1. There should be a slight movement away from 0 when we move to a random effects model, due to the interpretation as a conditional rather than a marginal parameter.

We now experiment with putting a proper prior on β_0, β_1. If we assume that the 5% and 95% points of the prior for β_0 are 0.8 and 1.2 (perhaps more sensible if we center the exposure) then we obtain a prior mean of -0.02 and prior sd of 0.123 (giving a precision of 65.8). For β_1 50% and 95% points of 1 and 10 give a prior mean and sd of 0 and 1.40 (to give precision 0.51). The results were found to be very sensitive to the choice of standard deviation for β_1, as we see below.

```
NON-SPATIAL RANDOM EFFECTS with informative priors on beta
node     mean    sd      MC error 2.5%    median  97.5%   start   sample
RRx      1.069   0.06011 5.964E-4 0.956   1.067   1.192   5000    15001
beta0    -0.0547 0.1051  0.00104  -0.260 -0.0544 0.1515  5000    15001
beta1    0.06556 0.05613 5.567E-4 -0.044 0.06517 0.1754  5000    15001
sdV      0.3473  0.05828 0.00228 0.2552  0.3398  0.4794  5000    15001

SPATIAL AND NON-SPATIAL with flat priors.
node     mean    sd      MC error 2.5%    median  97.5%   start   sample
RRx      1.109   0.07632 8.538E-4 0.9674 1.105   1.268   5000    25001
beta0    -0.2025 0.2777  0.01489  -0.737 -0.2022 0.348   5000    25001
beta1    0.1007  0.06861 7.654E-4 -0.033 0.1001  0.2374  5000    25001
dhalf    11.84   5.174   0.2259  4.635   10.92   24.35   5000    25001
phi      0.07028 0.03217 0.00141 0.02847 0.0635  0.1495  5000    25001
sdU      0.3288  0.04991 0.00191 0.247   0.3231  0.4415  5000    25001
sdV      0.3493  0.05906 0.00178 0.2578  0.3415  0.4897  5000    25001
vratio   0.4725  0.1056  0.00375 0.2664  0.4726  0.6747  5000    25001

SPATIAL AND NON-SPATIAL with informative priors.

node     mean    sd      MC error 2.5%    median  97.5%  start   sample
RRx      1.092   0.06832 0.002265 0.9661 1.089   1.233   5000    25001
beta0    -0.0233 0.1124  0.00324 -0.2416 -0.0233 0.1974  5000    25001
beta1    0.08625 0.06233 0.00206 -0.0345 0.08529 0.2096  5000    25001
dhalf    11.63   5.36    0.23    4.481   10.6    25.04   5000    25001
phi      0.07196 0.03262 0.00134 0.02768 0.06541 0.1547  5000    25001
sdU      0.3273  0.05101 0.00187 0.2448  0.3212  0.4447  5000    25001
sdV      0.3489  0.06001 0.00183 0.2531  0.3411  0.4849  5000    25001
vratio   0.4708  0.1077  0.00355 0.2661  0.4686  0.6841  5000    25001

SPATIAL AND NON-SPATIAL with flat priors and random effects and
informative priors and tau.T prior
node     mean    sd      MC error 2.5%    median  97.5%   start   sample
RRx      1.106   0.07554 0.001611 0.9693 1.103   1.264   1000    7001
beta0    -0.1854 0.2143  0.009454 -0.611 -0.1796 0.2245  1000    7001
beta1    0.09836 0.06813 0.001458 -0.031 0.09799 0.2341  1000    7001
dhalf    11.59   5.553   0.3881   4.323  10.48   25.46   1000    7001
phi      0.07348 0.03461 0.00252 0.02722 0.06615 0.1603  1000    7001
```

```
sdU     0.19    0.09709 0.01007 0.03716 0.1941  0.3608  1000    7001
sdV     0.2432  0.1061  0.01063 0.03126 0.2651  0.4173  1000    7001
vratio  0.4112  0.3363  0.03542 0.01362 0.3177  0.9882  1000    7001

spatial

model
{
  for (i in 1:nind){
    Y[i] ~ dbern(p[i])
    logit(p[i]) <- beta0 + beta1*exposure[i] + V[i] + U[i]
    V[i] ~ dnorm(0,tau.V)
    mean[i] <- 0
  }
  U[1:nind] ~ spatial.exp(mean[],x[],y[],tau.U,phi,1)
  dhalf ~ dlnorm(2.30,5.63) # 50% chance that corr falls to half in less than
                            # 10km, 95% chance in less than 20km. Note 1/var!
  phi <- 0.6931/dhalf
  vratio <- sdU*sdU/(sdU*sdU+sdV*sdV)
  tau.T ~ dgamma(10,1.10)
  pn ~ dbeta(1,1)
  sdU <- sqrt(pn/tau.T)
  sdV <- sqrt((1-pn)/tau.T)
  tau.U <- 1/(sdU*sdU)
  tau.V <- 1/(sdV*sdV)
#  tau.V ~ dgamma(10,1.10)
#  tau.U ~ dgamma(10,1.10)
#  beta0 ~ dnorm(0,65.8)
#  beta1 ~ dnorm(0,0.51)
  beta0 ~ dflat()
  beta1 ~ dflat()
  RRx <- exp(beta1)
#  sdU <- 1/sqrt(tau.U)
#  sdV <- 1/sqrt(tau.V)
}
```