Bayesian Generalized Linear Mixed Models

Random Intercepts Model

Suppose we have response measurements, Y_{ij}, at time j for subject i. For $\mu_{ij} = E(Y_{ij}|t_{ij}, b_i)$, we can fit a model with random intercepts:

$$g(\mu_{ij}) = \beta_0 + \beta_1 \cdot t_{ij} + b_{0,i},$$

where $g(\cdot)$ can be any of the usual link functions (identity, log, logit, $\cdots$).

Random Intercepts Model

In a Bayesian framework, we will need to assign prior distributions to b_i, β_0, and β_1 (and any other hyperparameters that may arise). Usually, we will assign flat priors to β_0 and β_1:

$$\begin{aligned} \pi(\beta_0) &= N(0, 10^2) \\ \pi(\beta_1) &= N(0, 10^2). \end{aligned}$$

We could assign b_i a $N(0, \tau^2)$ prior, in which case we would need to specify a prior for τ^2 (maybe an $IG(1, 0.0260)$). Note that our likelihood is given by

$$\prod_{i,j} p(Y_{ij} | t_{ij}, b_i)$$

Random Intercepts Model

If we assume that $p(Y_{ij}|t_{ij}, b_i) = N(\mu_{ij}, \sigma^2)$, where

$$\mu_{ij} = \beta_0 + \beta_1 \cdot t_{ij} + b_{0,i},$$

then we would need to assign prior distributions to b_i, β_0, β_1, and σ^2. If we used the above prior for b_i, our posterior distribution would be proportional to:

$$\prod_{i,j} p(Y_{ij}|t_{ij}, b_i, \sigma^2, \tau^2) \prod_i \pi(b_i|\tau^2)\pi(\beta_0)\pi(\beta_1)\pi(\sigma^2)\pi(\tau^2).$$

WinBUGS Note

Recall that in WinBUGS, we need to parametrize the Normal distribution in terms of the precision, rather than the variance or standard deviation. So in that case, we would usually assign Gamma(1, 0.0260) priors to the precision variables.

Xerophthalmia Example

- A study conducted in Indonesia to determine the causes and effects of vitamin A deficiency in pre-school children.
- Xerophthalmia is an ocular manifestation of vitamin A deficiency.
- Question of interest: Are vitamin A deficient children at an increased risk of respiratory infection?

Xerophthalmia Example

Along with the presence of respiratory infection, the following variables of interest were measured for 275 children for up to six consecutive quarterly visits:

- age (measured in months, centered at 36)
- presence of xerophthalmia
- cosine and sine terms for the annual cycle (seasons)
- gender
- height for age (measured as a percent of the NCHS standard (centered at 90%); indicates longterm nutritional status)
- presence of stunting (being below 85% in height for age)

Xerophthalmia Example

	ID	respinf	age	xerop	ctime	stime	sex	hage	stunt
1	121013	0	31	0	-1	0	0	-3	0
2	121013	0	34	0	0	-1	0	-3	0
3	121013	0	37	0	1	0	0	-2	0
4	121013	0	40	0	0	1	0	-2	0
5	121013	1	43	0	-1	0	0	-2	0
6	121013	0	46	0	0	-1	0	-3	0
7	121113	0	-9	0	-1	0	1	2	0
8	121113	0	-6	0	0	-1	1	0	0
9	121113	0	-3	0	1	0	1	-1	0
10	121113	0	0	0	0	1	1	-2	0
11	121113	1	3	0	-1	0	1	-3	0
12	121113	0	6	0	0	-1	1	-3	0
13	121114	0	-26	0	-1	0	0	8	0
14	121114	0	-23	0	0	-1	0	5	0
15	121114	0	-20	0	1	0	0	3	0

Xerophthalmia Example

- Since we are interested in inference at the individual level (rather than the population level), we want to fit a generalized linear mixed effects model (ie. a **conditional** model).
- Fitting a random effects model to the data allows us to address the question of how an individual child's risk for respiratory infection would change if their vitamin A status were to change.
- How will we do this? By allowing each child to have a distinct intercept, which represents their propensity for infection.

Xerophthalmia Example

Conditional Regression Model:

$$
\begin{aligned}
Y_{ij} &= \text{presence of a respiratory infection at visit } j \text{ for child } i \\
\mu_{ij} &= E(Y_{ij}|\mathbf{X}_{ij}, b_i)
\end{aligned}
$$

$$
\begin{aligned}
\text{logit}(\mu_{ij}) &= b_i + \beta_0 + \beta_1 \cdot age_{ij} + \beta_2 \cdot xerop_{ij} + \beta_3 \cdot sex_{ij} + \\
&\quad \beta_4 \cdot cos(time)_{ij} + \beta_5 \cdot sin(time)_{ij} + \beta_6 \cdot height(age)_{ij} + \\
&\quad \beta_7 \cdot stunted_{ij}.
\end{aligned}
$$

Xerophthalmia Example

Prior specification:

$$
\begin{aligned}
b_i &\sim N(0, \tau_b) \quad \text{where } \tau_b \text{ is the } \textbf{precision} \\
\beta_j &\sim N(0, 10^{-4}) \quad j = 0, \cdots, 7 \\
\tau_b &\sim Gamma(1, 0.0260)
\end{aligned}
$$

Xerophthalmia Example - Results

node	mean	sd	MC error	2.5%	median	97.5%
β_0	-2.612	0.2223	0.006759	-3.089	-2.598	-2.22
β_1	-0.034	0.0071	7.431E-5	-0.0482	-0.034	-0.0204
β_2	0.6214	0.4814	0.003674	-0.3736	0.6402	1.509
β_3	-0.4283	0.2469	0.002228	-0.9212	-0.4247	0.0447
β_4	-0.5993	0.1734	0.001316	-0.947	-0.597	-0.267
β_5	-0.1729	0.1732	0.001014	-0.5157	-0.171	0.1636
β_6	-0.0481	0.02619	2.593E-4	-0.1011	-0.0475	0.00172
β_7	0.158	0.4326	0.003831	-0.7069	0.1613	1.002
σ_b	0.5528	0.3007	0.01699	0.1026	0.5437	1.147

Closing Remarks

- Bayesian methods are still likelihood based.
- Inference obtained using Bayesian GLMMs will be at the **individual** level.
- If you hate WinBUGS, that's ok! These models aren't too bad to fit using R. You can come see me or talk to me for help with Bayesian models in R.