

3.3 Mixed logistic regression; a comparison with WinBUGS

Model description Let $\mathbf{y} = (y_1, \dots, y_n)$ be a vector of dichotomous observations ($y_i \in \{0, 1\}$), and let $\mathbf{u} = (u_1, \dots, u_q)$ be a vector of independent random effects, with $u_i \sim N(0, \sigma^2)$. The following relationship between the success probability $\pi_i = \Pr(y_i = 1)$ and explanatory variables (contained in matrices \mathbf{X} and \mathbf{Z}) is assumed:

$$\log\left(\frac{\pi_i}{1 - \pi_i}\right) = \mathbf{X}_i\boldsymbol{\beta} + \mathbf{Z}_i\mathbf{u},$$

where \mathbf{X}_i and \mathbf{Z}_i are the i 'th rows of the known design matrices \mathbf{X} and \mathbf{Z} , respectively, and $\boldsymbol{\beta}$ is a vector of regression parameters. Thus, the hyper-parameters of the model are $\boldsymbol{\beta}$ and σ .

Results Our goal is to compare computation times with WinBUGS on a simulated data set. For this purpose we use $n = 200$, $p = 5$, $q = 30$, $\sigma = 0.1$ and $\beta_j = 0$ for all j . The covariate matrices \mathbf{X} and \mathbf{Z} are generated randomly with each element uniformly distributed on $[-2, 2]$. As start values for both ADMB and WinBUGS we use $\beta_j = -1$ and $\sigma = 4.5$. In WinBUGS we use a uniform $[-10, 10]$ prior for β_j and a standard (in the WinBUGS literature) noninformative gamma prior on $\tau = \sigma^{-2}$. In ADMB the parameter constraints $\beta_j \in [-10, 10]$ and $\log(\sigma) \in [-5, 3]$ are used in the optimization process.

On the simulated dataset ADMB used 15 seconds to converge to the optimum of likelihood surface. On the same dataset we first ran WinBUGS (Version 1.4) for 5,000 iterations. The recommended convergence diagnostic in WinBUGS is the Gelman-Rubin plot (see the help files available from the menus in WinBUGS), which requires that at least two Markov chains are run in parallel. From the Gelman-Rubin plots for σ and $\boldsymbol{\beta}$ it was apparent that convergence occurred after approximately 2,000 iterations. A few of the components of \mathbf{u} did not yet seem to have reached equilibrium. The time taken by WinBUGS to perform 2,000 iterations of two chains was approximately 700 seconds. The estimate of σ was 0.1692 as calculated by AD Model Builder, while the posterior mean estimate calculated by WinBUGS σ was 0.1862.

The robustness of the two approaches was investigated by decreasing the amount of information while holding the number of parameters (p and q) constant. For $n = 50$, WinBUGS came up with an error message before convergence was reached. ADMB converged without problems to the optimum of the likelihood function, although observed Fisher information matrix showed that the hyper-parameters were weakly determined.