Notes: mixed models

Ben Bolker

November 23, 2011

Version: 2011-11-23 19:40:12

1 Random effects

Sometimes it's useful to treat levels of a factor as random variables rather than as fixed parameters to estimate. When should you use a random variable? It's hard to be precise

- levels chosen from a larger population
- levels chosen randomly
- levels non-exhaustive
- want to make predictions about new levels in the future
- sufficient number of levels to estimate a variance (> 5, preferably > 10)
- not interested in hypotheses about specific values
- small/variable amounts of information available from many levels
- parameters estimated with *shrinkage*

2 Model definition

Instead of our usual $\eta = X\beta$, $Y \sim D(g^{-1}(\eta))$ we add another level: now $\eta = X\beta + Zb$, where b is a vector of *random variables* drawn from (typically) a normal distribution; $b \sim \text{MVN}(0, \Sigma)$. Typically the b values are independent.

Simplest example: intercept-only random effect. Z is an indicator matrix describing which unit an observation is in. Equivalent to $\eta_i = X_{i\cdot}\beta + b_{u(i)}$, where $u(i)$ defines the unit the i^{th} observation belongs to. Compare with random-slopes model.

3 Estimation

Typically based on *marginal likelihood*:

- likelihood of i^{th} obs. in block j is $L(x_{ij}|\theta_i, \sigma_w^2)$
- likelihood of a particular block mean θ_j is $L(\theta_j|0, \sigma_b^2)$
- marginal likelihood is $\int L(x_{ij}|\theta_j, \sigma_w^2) L(\theta_j|0, \sigma_b^2)\, d\theta_j$

Balance (dispersion of RE around 0) with (dispersion of data conditional on RE)

Shrinkage: estimated values get "shrunk" toward the overall mean, especially in small-sample/extreme units

How do we do it?

Penalized quasi-likelihood (PQL) Alternate steps of estimating GLM using known RE variances to calculate weights; estimate LMMs given GLM fit. Flexible (allows spatial/temporal correlations, crossed REs) *biased* for small unit samples (e.g. counts < 5, binary or low-survival data). Widely used: SAS `PROC GLIMMIX`, R glmmPQL

Laplace approximation approximate marginal likelihood. for given $\boldsymbol{\beta}$, $\boldsymbol{\theta}$ (RE parameters), find conditional modes by penalized, iterated reweighted least squares; then use second-order Taylor expansion around the conditional modes. More accurate than PQL; reasonably fast and flexible (glmer in the lme4 package)

(adaptive) Gauss-Hermite quadrature (AGQ) As above, but compute additional terms in the integral (typically 8, but often up to 20). Most accurate, slowest, hence not flexible (2–3 RE at most, maybe only 1)