Beta-binomial

Parametrisation

The Beta-binomial distribution arise from a hierarchical model where the probability p is Beta-distributed[1]

$$\pi(p) = \frac{1}{B(\alpha, \beta)} p^{\alpha-1}(1-p)^{\beta-1}, \quad \alpha > 0, \ \beta > 0$$

and the response y is Binomial

$$\pi(y \mid p) = \binom{n}{y} p^y (1-p)^{(n-y)}, \quad y = 0, 1, \ldots, n.$$

The marginal distribution for y is the Beta-binomial,

$$\pi(y) = \binom{n}{y} \frac{B(y + \alpha, n - y + \beta)}{B(\alpha, \beta)}.$$

Link-function

The mean and variance of y are given as

$$\mu = n\frac{\alpha}{\alpha + \beta} = n\mu_p, \quad \text{and} \quad \sigma^2 = n\mu_p(1 - \mu_p)\left(1 + (n-1)\rho\right),$$

where $\mu_p = \frac{\alpha}{\alpha+\beta}$ is the mean of the probability p from the Beta-distribution, and $\rho = \frac{1}{\alpha+\beta+1}$ ($0 < \rho < 1$) is the pairwise correlation between the n Bernoulli draws and is the over-dispersion parameter.

The mean probability μ_p is linked to the linear predictor by

$$\mu_p = \frac{\exp(\eta)}{1 + \exp(\eta)}$$

using the default logit-link.

Hyperparameter

The hyperparameter is the over-dispersion parameter ρ, which is represented as

$$\rho = \frac{\exp(\theta)}{1 + \exp(\theta)}$$

and the prior is defined on θ.

Specification

- family = `betabinomial`

- Required arguments: y and Ntrials $= n$ (default Ntrials $= 1$).

---

[1]The Beta-function is $B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$, where $\Gamma(x)$ is the Gamma-function.

Hyperparameter spesification and default values

doc The Beta-Binomial likelihood

hyper

    theta

        hyperid 62001
        name overdispersion
        short.name rho
        initial 0
        fixed FALSE
        prior gaussian
        param 0 0.4
        to.theta `function(x) log(x / (1 - x))`
        from.theta `function(x) exp(x) / (1 + exp(x))`

survival FALSE

discrete TRUE

link default logit loga cauchit probit cloglog loglog robit sn

pdf betabinomial

Example

In the following example we estimate the parameters in a simulated example.

```
## overdispersion parameter in the betabinomial
rho = 0.7

n = 1000
z = rnorm(n, sd=0.2)
Ntrials = sample(1:10, n, replace=TRUE)
eta = 1 + z
p.eta = exp(eta)/(1+exp(eta))
a = p.eta * (1-rho)/rho
b = (p.eta * rho - p.eta - rho + 1)/rho
p = rbeta(n, a, b)
y = rbinom(n, Ntrials, p)

formula = y ~ 1 + z
data = data.frame(y, z)
r = inla(formula, data = data,
        family = "betabinomial", Ntrials=Ntrials)
summary(r)
```

Notes

None.