

Likelihood fmri

Parametrisation

This is special parameterisation of the non-central χ_ν distribution. Let $\{x_i\}$ are iid Gaussians with mean μ and precision τ , then

$$z = \sqrt{\sum_{i=1}^{\nu} \tau x_i^2}$$

is non-central χ -distribution with (integer and fixed by design) $\nu > 0$ degrees of freedom, and non-centrality parameter

$$\rho = \sqrt{\nu\tau\mu^2}.$$

The observation y is $y = z/\sqrt{\tau}$ and we're interested in the underlying true signal

$$\lambda = \rho/\sqrt{\tau} = \sqrt{\nu\mu^2}$$

Link-function

The linkfunction is given as

$$\log(\lambda) = \eta$$

where η is the linear predictor.

Hyperparameters

The hyperparameters are $\theta = (\theta_1, \theta_2)$, where

$$\tau = \exp(\theta_1)$$

is the precision, and

$$\nu = \theta_2$$

For technical reasons, ν is implemented as a hyper-parameter, but is required to be fixed. Hence, the initial value for θ_2 defines the (fixed) value for ν .

The prior is given on θ_1 .

Specification

- `family="fmri"` or `family="fmrismr"`
- Required arguments: `y` (and optional `scale` for `fmri` to scale τ)

Hyperparameter specification and default values

doc fmri distribution (special nc-chi)

hyper

theta1

hyperid 103101

name precision

short.name prec

initial 0

fixed FALSE

```

    prior loggamma
    param 10 10
    to.theta function(x) log(x)
    from.theta function(x) exp(x)
theta2
    hyperid 103202
    name dof
    short.name df
    initial 4
    fixed TRUE
    prior normal
    param 0 1
    to.theta function(x) x
    from.theta function(x) x

status experimental

survival FALSE

discrete FALSE

link default log

pdf fmri

doc fmri distribution (special nc-chi)

hyper
    theta1
        hyperid 104101
        name precision
        short.name prec
        initial 0
        fixed FALSE
        prior loggamma
        param 10 10
        to.theta function(x) log(x)
        from.theta function(x) exp(x)
    theta2
        hyperid 104201
        name dof
        short.name df
        initial 4
        fixed TRUE
        prior normal
        param 0 1
        to.theta function(x) x

```

```

from.theta function(x) x
status experimental
survival TRUE
discrete FALSE
link default log
pdf fmri

```

Example

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

```

n <- 300
x <- rnorm(n, sd = 0.3)
df <- 1
prec <- 3
eta <- 1 + x
lambda <- exp(eta)
y <- sqrt(rchisq(n, df = df, ncp = prec * lambda^2) /prec)

r <- inla(y ~ 1 + x,
          data = data.frame(y, x),
          family = "fmri",
          control.family = list(hyper = list(df = list(initial = df))),
          control.inla = list(cmin = 0,
                              int.strategy = "eb",
                              strategy = "adaptive"),
          verbose = TRUE)
summary(r)

## 'cmin=0' seems to be required only for initial values that can give
## 'crazy' values. We can rerun without this re-starting at the prev fit,
## to validate
r$.args$control.inla$cmin <- -Inf
r$.args$control.inla$int.strategy <- "auto"
rr <- inla.rerun(r)
summary(rr)

```

Notes

Thanks to LS for providing all the details and a robust implementation of this likelihood.