

Skew-Normal likelihood

Parametrisation

The standardised Skew-Normal distribution is

$$f(z) = \frac{2}{\omega_\alpha} \phi\left(\frac{z - \xi_\alpha}{\omega_\alpha}\right) \Phi\left(\alpha \frac{z - \xi_\alpha}{\omega_\alpha}\right)$$

where ω_α and ξ_α are so that the mean is zero and variance is one, and they depends both on the skewness parameter α .

The skew-normal likelihood is defined as the density wrt y , where

$$z = (y - \eta)\sqrt{w\tau} \sim f(z)$$

and

η : is the the linear predictor

τ : is the precision

w : is a fixed scale or weight, $w > 0$,

Link-function

The mean equals the linear predictor

$$\mu = \eta$$

Hyperparameters

The precision is represented as

$$\theta_1 = \log \tau$$

and the prior is defined on θ_1 .

The (standardised) skewness γ , is represented as

$$\gamma = 0.988 \left(2 \frac{\exp(\theta_2)}{1 + \exp(\theta_2)} - 1 \right)$$

and the prior is defined on θ_2 . The standardised skewness depends on α as

$$\gamma = \frac{4 - \pi}{2} \frac{\left(\delta \sqrt{2/\pi}\right)^3}{(1 - 2\delta^2/\pi)^{3/2}}, \quad \delta = \frac{\alpha}{\sqrt{1 + \alpha^2}}$$

Specification

- `family="sn"`
- Required arguments: y and w (keyword `scale`, and $w = 1$ by default).

Hyperparameter specification and default values

doc The Skew-Normal likelihood

hyper

theta1

hyperid 74001
name log precision
short.name prec
initial 4
fixed FALSE
prior loggamma
param 1 5e-05
to.theta function(x) log(x)
from.theta function(x) exp(x)

theta2

hyperid 74002
name logit skew
short.name skew
initial 0.00123456789
fixed FALSE
prior pc.sn
param 10
to.theta function(x, skew.max = 0.988) log((1 + x / skew.max) / (1 - x / skew.max))
from.theta function(x, skew.max = 0.988) skew.max * (2 * exp(x) / (1 + exp(x)) - 1)

status experimental

survival FALSE

discrete FALSE

link default identity

pdf sn

Example

```
library(sn)
set.seed(246)
n = 300
x = rnorm(n, sd = 1)
eta = 1+x
skewness = 0.25
y = numeric(n)
prec <- 100
for(i in 1:n) {
  ## map moments to sn-parameters c(xi, omega, alpha)
  param = INLA::inla.sn.reparam(moments = c(eta[i], 1/prec, skewness))
  y[i] = rsn(1, xi=param$xi, omega = param$omega, alpha = param$alpha)
```

```

}

r = inla(y ~ 1+x,
        family = "sn",
        data = data.frame(y, x),
        control.family = list(
          hyper = list(prec = list(
                        prior = "pc.prec",
                        param = c(3, 0.01))))))
summary(r)

```

Notes

- This implementation replaces older ones ("sn" and "sn2") from 16th September 2020.
- A $N(a, 0)$ prior is interpreted as a constant prior with density equal to one.