

PC prior for precision

Parametrization

The PC prior for the precision τ has density

$$\pi(\tau) = \frac{\lambda}{2} \tau^{-3/2} \exp\left(-\lambda \tau^{-1/2}\right), \quad \tau > 0$$

for $\lambda > 0$ where

$$\lambda = -\frac{\ln(\alpha)}{u}$$

and (u, α) are the parameters to this prior. The interpretation of (u, α) is that

$$\text{Prob}(\sigma > u) = \alpha, \quad u > 0, \quad 0 < \alpha < 1,$$

where the standard deviation is $\sigma = 1/\sqrt{\tau}$. The density, cumulative distribution function, quantile function, and a random number generator for this distribution are implemented in the `inla.pc.{d,p,q,r}prec` functions.

Internally, R-INLA uses the log-precision rather than the precision and the corresponding PC prior for the log-precision x has density

$$\pi(x) = \frac{\lambda}{2} \exp\left(-\lambda \exp\left(-\frac{x}{2}\right) - \frac{x}{2}\right). \quad (1)$$

Specification

This prior for the hyperparameters is specified inside the `hyper`-specification, as

```
hyper = list(<theta> = list(prior="pc.prec", param=c(<u>,<alpha>)))
```

Example

Notes

See also functions `inla.pc.{d,p,q,r}prec`