

Generalized Pareto distribution

Parametrisation

The generalized Pareto (GP) distribution with positive shape parameter has cumulative distribution function

$$F(y; \sigma, \xi) = 1 - \left(1 + \xi \frac{y}{\sigma}\right)^{-1/\xi}, \quad y > 0,$$

for a continuous response y where

ξ : is the tail parameter, $\xi > 0$

σ : is the scale parameter, $\sigma > 0$

The limit for $\xi \downarrow 0$ is $F(y; \sigma, 0) = 1 - \exp(-y/\sigma)$.

Link function

The linear predictor η controls the α quantile of the GP

$$P(y \leq q_\alpha) = \alpha$$

and $q_\alpha = \exp(\eta)$. The scaling σ , is then a function of (q_α, ξ) , as

$$\sigma = \frac{\xi \exp(\eta)}{(1 - \alpha)^{-\xi} - 1}$$

Hyperparameters

The GP model has one hyperparameter. The tail $\xi > 0$ is represented as

$$\xi = \xi_{\text{low}} + (\xi_{\text{high}} - \xi_{\text{low}}) \frac{\exp(\theta)}{1 + \exp(\theta)}$$

and the prior is defined on θ , with constant low and high values. The prior is FIXED to `pc.gevtail`, see `inla.doc("pc.gevtail")` for more info.

Specification

- `family="gp"`
- Required arguments: y and the quantile α .

The quantile is given as `control.family=list(control.link=list(quantile= $\alpha$ ))`.

Hyperparameter specification and default values

doc Generalized Pareto likelihood

hyper

theta

hyperid 101201

name tail

short.name xi

output.name Tail parameter for the gp observations

output.name.intern Intern tail parameter for the gp observations

initial -4

fixed FALSE

prior pc.gevtail

param 7 0 0.5

to.theta function(x, interval = c(REPLACE.ME.low, REPLACE.ME.high)) log(-(interval

from.theta function(x, interval = c(REPLACE.ME.low, REPLACE.ME.high)) interval[1]

survival FALSE

discrete TRUE

link default quantile

pdf genPareto

Example

```
rgp = function(n, sigma, eta, alpha, xi = 0.001)
{
  if (missing(sigma)) {
    stopifnot(!missing(eta) && !missing(alpha))
    sigma = exp(eta) * xi / ((1.0 - alpha)^(-xi) - 1.0)
  }
  return (sigma / xi * (runif(n)^(-xi) - 1.0))
}
```

```
n = 300
x = runif(n)-0.5
eta = 1+x
alpha = 0.95
xi = 0.3
y = rgp(n, eta = eta, alpha = alpha, xi=xi)
```

```

r = inla(y ~ 1+x,
        data = data.frame(y, x),
        family = "gp",
        control.family = list(control.link = list(quantile = alpha)),
        control.predictor = list(compute=TRUE),
        verbose=TRUE)

rx = range(c(r$summary.fitted.values$mean, exp(eta)))
plot(r$summary.fitted.values$mean, exp(eta),
     xlim = rx, ylim = rx)
abline(a=0,b=1)

```

Notes

The prior for ξ is fixed to `pc.gevtail` or one that respect an given interval similarly. Note that the default prior (and the internal representation of it and its name) changed 2019/10/12.