

Tweedie

Parametrisation

The Tweedie distribution¹ is a compound Poisson-Gamma model, where

$$Y = \sum_{i=1}^N X_i,$$

and $\{X_i\}$ are iid Gamma variables with parameter (α, γ) so that the mean of X_i is $\alpha\gamma$ and variance $\alpha\gamma^2$, and N is (independent) Poisson with mean λ . Since N can be 0 with a positive probability, then the Tweedie distribution have a singleton in zero and is continuous for $y > 0$.

We will use the following reparametersation

$$\mu = \lambda\alpha\gamma, \quad p = \frac{\alpha + 2}{\alpha + 1}, \quad \frac{\phi}{w} = \frac{\lambda^{1-p}(\alpha\gamma)^{2-p}}{2-p}$$

where $w > 0$ is a fixed scaling, so the mean of Y is $\mu > 0$, variance is $\frac{\phi}{w}\mu^p$ where $1 < p < 2$, and ϕ is a dispersion parameter.

Link-function

The linkfunction is given as

$$\log(\mu) = \eta$$

where η is the linear predictor.

Hyperparameters

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

$$p = 1 + \frac{\exp(\theta_1)}{1 + \exp(\theta_1)}, \quad 1 < p < 2.$$

and

$$\phi = \exp(\theta_2), \quad \phi > 0$$

The priors are given on (θ_1, θ_2) .

Specification

- `family="tweedie"`
- Required arguments: y (and optional w through option `scale`)

Hyperparameter spesification and default values

`doc` Tweedie distribution

`hyper`

`thetal`

`hyperid` 102101

`name` p

¹This documentation follows the notation in *Likelihood-based and Bayesian methods for Tweedie compound Poisson linear mixed models*, by Yanwei Zhang, *Stat Comput* (2013) 23:743–757, DOI 10.1007/s11222-012-9343-7

```

short.name p
output.name p parameter for Tweedie
output.name.intern p_intern parameter for Tweedie
initial 0
fixed FALSE
prior normal
param 0 100
to.theta function(x, interval = c(1.0, 2.0)) log(-(interval[1] - x) / (interval[2] - x))
from.theta function(x, interval = c(1.0, 2.0)) interval[1] + (interval[2] - interval[1]) * x

theta2
  hyperid 102201
  name dispersion
  short.name phi
  output.name Dispersion parameter for Tweedie
  output.name.intern Log dispersion parameter for Tweedie
  initial -4
  fixed FALSE
  prior loggamma
  param 100 100
  to.theta function(x) log(x)
  from.theta function(x) exp(x)

survival FALSE

discrete FALSE

link default log

pdf tweedie

```

Example

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

```

library(tweedie)
library(INLA)

n <- 300
x <- rnorm(n, sd = 0.2)
eta <- 1 + x
mu <- exp(eta)

p <- 1.32
phi <- 2.0
y <- numeric(n)
for(i in 1:n) {
  y[i] <- rtweedie(1, xi = p, mu = mu[i], phi = phi)
}

r <- inla(y ~ 1 + x,

```

```

data = data.frame(y, x),
## offset = rep(log(mean(y)), n),
family = "tweedie",
control.family = list(hyper = list(
  theta1 = list(initial = 0),
  theta2 = list(initial = -4,
    prior = "loggamma",
    param = c(100, 100))),
## control.fixed = list(prec = 0, prec.intercept = 1),
## control.inla = list(cmin = 0, b.strategy = "skip"),
## inla.mode = "experimental",
num.threads = "4:1",
verbose = T)
summary(r)

```

Notes

This distribution is experimental, and changes will occur.