

Package ‘retel’

January 22, 2024

Type Package

Title Regularized Exponentially Tilted Empirical Likelihood

Version 0.1.0

Description Implements the regularized exponentially tilted empirical likelihood method. Details of the methods are given in Kim, MacEachern, and Peruggia (2023) <[doi:10.48550/arXiv.2312.17015](https://doi.org/10.48550/arXiv.2312.17015)>. This work was supported by the U.S. National Science Foundation under Grants No. SES-1921523 and DMS-2015552.

License GPL (>= 3)

URL <https://github.com/markean/retel>

BugReports <https://github.com/markean/retel/issues>

Depends R (>= 4.1.0)

Suggests melt, spelling

Encoding UTF-8

Language en-US

LazyData true

RoxygenNote 7.3.0

Imports checkmate, Matrix, matrixcalc, nloptr

NeedsCompilation no

Author Eunseop Kim [aut, cph, cre],
Steven MacEachern [ctb, ths],
Mario Peruggia [ctb, ths]

Maintainer Eunseop Kim <markean@pm.me>

Repository CRAN

Date/Publication 2024-01-22 18:20:06 UTC

R topics documented:

etel	2
income	3
retel	3

etel	<i>Exponentially tilted empirical likelihood</i>
------	--

Description

Computes exponentially tilted empirical likelihood.

Usage

```
etel(fn, x, par, opts = NULL)
```

Arguments

fn	An estimating function that takes the data <code>x</code> and parameter value <code>par</code> as its arguments, returning a numeric matrix. Each row is the return value from the corresponding row in <code>x</code> .
x	A numeric matrix, or an object that can be coerced to a numeric matrix. Each row corresponds to an observation. The number of rows must be greater than the number of columns.
par	A numeric vector of parameter values to be tested.
opts	An optional list with optimization options for <code>nloptr()</code> . Defaults to <code>NULL</code> .

Value

A single numeric value representing the log-likelihood ratio. It contains the optimization results as the attribute `optim`.

References

Schennach, SM (2005). "Bayesian Exponentially Tilted Empirical Likelihood." *Biometrika*, 92, 31–46.

Examples

```
# Generate data
set.seed(63456)
x <- rnorm(100)

# Define an estimating function (ex. mean)
fn <- function(x, par) {
  x - par
}

# Set parameter value
par <- 0

# Call the etel function
etel(fn, x, par)
```

income	<i>Median Income for 4-Person Families in the USA</i>
--------	---

Description

A dataset of median income for 4-person families by state.

Usage

```
data("income")
```

Format

A data frame with 51 rows and 6 columns:

state States, including the District of Columbia.

mi_1979 Estimated median income for 4-person families in 1979 (standardized).

mi_1989 Estimated median income for 4-person families in 1989 (standardized).

pqi_1979 Per capita income in 1979.

pqi_1989 Per capita income in 1989.

ami Census median income in 1979, adjusted for per capita income (standardized).

Source

<https://www.census.gov/data/tables/time-series/demo/income-poverty/4-person.html>

Examples

```
data("income")
income
```

retel	<i>Regularized exponentially tilted empirical likelihood</i>
-------	--

Description

Computes regularized exponentially tilted empirical likelihood.

Usage

```
retel(fn, x, par, mu, Sigma, tau, type = "full", opts = NULL)
```

Arguments

fn	An estimating function that takes the data x and parameter value par as its arguments, returning a numeric matrix. Each row is the return value from the corresponding row in x .
x	A numeric matrix, or an object that can be coerced to a numeric matrix. Each row corresponds to an observation. The number of rows must be greater than the number of columns.
par	A numeric vector of parameter values to be tested.
mu	A numeric vector of parameters for regularization. See 'Details' for more information.
Sigma	A numeric matrix, or an object that can be coerced to a numeric matrix, of parameters for regularization. See 'Details' for more information.
tau	A single numeric parameter for regularization. See 'Details' for more information.
type	A single character indicating the type of regularized exponentially tilted empirical likelihood. It must be either "full" or "reduced". Defaults to "full". See 'Details' for more information.
opts	An optional list with optimization options for <code>nloptr()</code> . Defaults to NULL.

Details

Let $\{\mathbf{X}_i\}_{i=1}^n$ denote independent d_x -dimensional observations from a complete probability space $(\mathcal{X}, \mathcal{F}, P)$ satisfying the moment condition:

$$E_P[\mathbf{g}(\mathbf{X}_i, \boldsymbol{\theta})] = \mathbf{0},$$

where $\mathbf{g} : \mathbb{R}^{d_x} \times \Theta \mapsto \mathbb{R}^p$ is an estimating function with the true parameter value $\boldsymbol{\theta}_0 \in \Theta \subset \mathbb{R}^p$.

For a given parameter value $\boldsymbol{\theta}$, regularized exponentially tilted empirical likelihood solves the following optimization problem:

$$\min_{\boldsymbol{\lambda} \in \mathbb{R}^p} \{d_n(\boldsymbol{\theta}, \boldsymbol{\lambda}) + p_n(\boldsymbol{\theta}, \boldsymbol{\lambda})\},$$

where

$$d_n(\boldsymbol{\theta}, \boldsymbol{\lambda}) = \frac{1}{n + \tau_n} \sum_{i=1}^n \exp(\boldsymbol{\lambda}^\top \mathbf{g}(\mathbf{X}_i, \boldsymbol{\theta}))$$

and

$$p_n(\boldsymbol{\theta}, \boldsymbol{\lambda}) = \frac{\tau_n}{n + \tau_n} \exp\left(\boldsymbol{\lambda}^\top \boldsymbol{\mu}_{n,\boldsymbol{\theta}} + \frac{1}{2} \boldsymbol{\lambda}^\top \boldsymbol{\Sigma}_{n,\boldsymbol{\theta}} \boldsymbol{\lambda}\right).$$

Here, $\tau_n > 0$, $\boldsymbol{\mu}_{n,\boldsymbol{\theta}}$, $\boldsymbol{\Sigma}_{n,\boldsymbol{\theta}}$ are all tuning parameters that control the strength of $p_n(\boldsymbol{\theta}, \boldsymbol{\lambda})$ as a penalty.

Once we have determined the solution $\boldsymbol{\lambda}_{RET}$, we define the likelihood ratio function as follows:

$$R_{RET}(\boldsymbol{\theta}) = \left(\frac{n + \tau_n}{\tau_n} p_c(\boldsymbol{\theta})\right) \prod_{i=1}^n (n + \tau_n) p_i(\boldsymbol{\theta}),$$

where

$$p_i(\boldsymbol{\theta}) = \frac{\exp\left(\boldsymbol{\lambda}_{RET}^\top \mathbf{g}(\mathbf{X}_i, \boldsymbol{\theta})\right)}{c_n(\boldsymbol{\theta}, \boldsymbol{\lambda}_{RET})} \quad (i = 1, \dots, n), \quad p_c(\boldsymbol{\theta}) = \frac{p_n(\boldsymbol{\theta}, \boldsymbol{\lambda}_{RET})}{c_n(\boldsymbol{\theta}, \boldsymbol{\lambda}_{RET})},$$

and $c_n(\boldsymbol{\theta}, \boldsymbol{\lambda}_{RET}) = d_n(\boldsymbol{\theta}, \boldsymbol{\lambda}_{RET}) + p_n(\boldsymbol{\theta}, \boldsymbol{\lambda}_{RET})$. The reduced version of the likelihood ratio function is defined as:

$$\tilde{R}_{RET}(\boldsymbol{\theta}) = \prod_{i=1}^n (n + \tau_n) p_i(\boldsymbol{\theta}).$$

See the references below for more details on derivation, interpretation, and properties.

Value

A single numeric value representing the log-likelihood ratio. It contains the optimization results as the attribute `optim`.

References

Kim E, MacEachern SN, Peruggia M (2023). "Regularized Exponentially Tilted Empirical Likelihood for Bayesian Inference." [doi:10.48550/arXiv.2312.17015](https://doi.org/10.48550/arXiv.2312.17015).

Examples

```
# Generate data
set.seed(63456)
x <- rnorm(100)

# Define an estimating function (ex. mean)
fn <- function(x, par) {
  x - par
}

# Set parameter value
par <- 0

# Set regularization parameters
mu <- 0
Sigma <- 1
tau <- 1

# Call the retel function
retel(fn, x, par, mu, Sigma, tau)
```

Index

* **datasets**

income, 3

etel, 2

income, 3

nloptr(), 2, 4

retel, 3