

Package ‘hIRT’

August 1, 2017

Type Package

Title Hierarchical Item Response Theory Models

Version 0.1.2

Description Implementation of a class of hierarchical item response theory (IRT) models where both the mean and the variance of latent preferences (ability parameters) may depend on observed covariates. The current implementation includes both the two-parameter latent trait model for binary data and the graded response model for ordinal data. Both are fitted via the Expectation-Maximization (EM) algorithm. Asymptotic standard errors are derived from the observed information matrix.

Depends R (>= 3.3.2), stats

Imports pryr (>= 0.1.2), rms (>= 5.1-1)

Suggests ggplot2 (>= 2.2.1)

License GPL (>= 3)

Encoding UTF-8

LazyData true

RoxygenNote 6.0.1

URL <http://github.com/xiangzhou09/hIRT>

BugReports <http://github.com/xiangzhou09/hIRT>

NeedsCompilation no

Author Xiang Zhou [aut, cre]

Maintainer Xiang Zhou <xiang_zhou@fas.harvard.edu>

Repository CRAN

Date/Publication 2017-08-01 19:00:01 UTC

R topics documented:

coef.hIRT	2
coef_item	3

hgrm	4
hltm	5
latent_scores	7
nes_econ2008	8
print.hIRT	8
summary.hIRT	9

Index	10
--------------	-----------

coef.hIRT	<i>Extracting Parameter Estimates from Hierarchical IRT Models.</i>
-----------	---

Description

Parameter estimates from either hltm or hgrm models. coef_mean reports results only for the mean equation. coef_var reports results only for the variance equation.

Usage

```
## S3 method for class 'hIRT'
coef(object, digits = 3, ...)
```

```
coef_mean(x, digits = 3)
```

```
coef_var(x, digits = 3)
```

Arguments

object	An object of class hIRT
digits	The number of significant digits to use when printing
...	further arguments passed to or from other methods
x	An object of class hIRT

Value

A data frame of parameter estimates, standard errors, z values and p values.

Examples

```
y <- nes_econ2008[, -(1:3)]
x <- model.matrix(~ party * educ, nes_econ2008)
z <- model.matrix(~ party, nes_econ2008)
nes_m1 <- hgrm(y, x, z)
coef(nes_m1)
coef_mean(nes_m1)
coef_var(nes_m1)
```

coef_item	<i>Extracting Estimates of Item Parameters from Hierarchical IRT Models.</i>
-----------	--

Description

Item parameter estimates from either hltm or hgrm models.

Usage

```
coef_item(x, by_item = TRUE, digits = 3)

## S3 method for class 'hgrm'
coef_item(x, by_item = TRUE, digits = 3)

## S3 method for class 'hltm'
coef_item(x, by_item = TRUE, digits = 3)
```

Arguments

x	An object of class hIRT
by_item	Logical. Should item parameters be stored item by item (if TRUE) or put together in a data frame (if FALSE)?
digits	the number of significant digits to use when printing

Value

Item parameter estimates, standard errors, z values, and p values organized as a data frame (if `by_item = TRUE`) or a list (if `by_item = FALSE`).

Examples

```
y <- nes_econ2008[, -(1:3)]
x <- model.matrix(~ party * educ, nes_econ2008)
z <- model.matrix(~ party, nes_econ2008)
nes_m1 <- hgrm(y, x, z)
coef_item(nes_m1)
```

hgrm

*Fitting Hierarchical Graded Response Models***Description**

hgrm fits a hierarchical graded response model in which both the mean and the variance of the latent preference (ability parameter) may depend on person-specific covariates (x and z). Specifically, the mean is specified as a linear combination of x and the log of the variance is specified as a linear combination of z . Nonresponses are treated as missing at random.

Usage

```
hgrm(y, x = matrix(1, nrow(y), 1), z = matrix(1, nrow(y), 1),
     beta_set = 1, sign_set = TRUE, control = list())
```

Arguments

<code>y</code>	A data frame or matrix of item responses.
<code>x</code>	A matrix of covariates, including intercept, that predict the mean of the latent preference.
<code>z</code>	A matrix of covariates, including intercept, that predict the variance of the latent preference.
<code>beta_set</code>	The index of the item whose discrimination parameter is restricted to be positive (or negative). It may take an integer value from 1 to <code>ncol(y)</code> .
<code>sign_set</code>	Logical. Should the discrimination parameter of the corresponding item (indexed by <code>beta_set</code>) be positive (if TRUE) or negative (if FALSE)?
<code>control</code>	A list of control values

max_iter The maximum number of iterations of the EM algorithm. Default 150.

eps Tolerance parameter used to determine convergence of the EM algorithm. Specifically, iterations continue until the Euclidean distance between β_n and β_{n-1} falls under `eps`, where β is the vector of item discrimination parameters. `eps=1e-4` by default.

max_iter2 The maximum number of iterations of the conditional maximization procedures for updating γ and λ . Default 15.

eps2 Tolerance parameter used to determine convergence of the conditional maximization procedures for updating γ and λ . Specifically, iterations continue until the Euclidean distance between two consecutive log likelihoods falls under `eps2`. `eps2=1e-3` by default.

K Number of Gauss-Legendre quadrature points for the E-step. Default 21.

C [-C, C] sets the range of integral in the E-step. `C=5` by default.

Value

An object of class hgrm.

coefficients	A data frame of parameter estimates, standard errors, z values and p values.
scores	A data frame of EAP estimates of latent preferences and their approximate standard errors.
vcov	Variance-covariance matrix of parameter estimates.
log_Lik	The log-likelihood value at convergence.
H	A vector denoting the number of response categories for each item.
p	The number of predictors for the mean equation.
q	The number of predictors for the variance equation.
item_names	Names of items.
call	The matched call.

References

Zhou, Xiang. 2017. "Hierarchical Item Response Models for Analyzing Public Opinion." Working Paper.

Examples

```
y <- nes_econ2008[, -(1:3)]
x <- model.matrix( ~ party * educ, nes_econ2008)
z <- model.matrix( ~ party, nes_econ2008)
nes_m1 <- hgrm(y, x, z)
print(nes_m1)
```

hltm

Fitting Hierarchical Latent Trait Models (for Binary Data).

Description

hltm fits a hierarchical latent trait model in which both the mean and the variance of the latent preference (ability parameter) may depend on person-specific covariates (x and z). Specifically, the mean is specified as a linear combination of x and the log of the variance is specified as a linear combination of z.

Usage

```
hltm(y, x = matrix(1, nrow(y), 1), z = matrix(1, nrow(y), 1),
     beta_set = 1, sign_set = TRUE, control = list())
```

Arguments

<code>y</code>	A data frame or matrix of item responses.
<code>x</code>	A matrix of covariates, including intercept, that predict the mean of the latent preference.
<code>z</code>	A matrix of covariates, including intercept, that predict the variance of the latent preference.
<code>beta_set</code>	The index of the item whose discrimination parameter is restricted to be positive (or negative). It may take an integer value from 1 to <code>ncol(y)</code> .
<code>sign_set</code>	Logical. Should the discrimination parameter of the corresponding item (indexed by <code>beta_set</code>) be positive (if TRUE) or negative (if FALSE)?
<code>control</code>	A list of control values
	max_iter The maximum number of iterations of the EM algorithm. Default 150.
	eps Tolerance parameter used to determine convergence of the EM algorithm. Specifically, iterations continue until the Euclidean distance between β_n and β_{n-1} falls under <code>eps</code> , where β is the vector of item discrimination parameters. <code>eps=1e-4</code> by default.
	max_iter2 The maximum number of iterations of the conditional maximization procedures for updating γ and λ . Default 15.
	eps2 Tolerance parameter used to determine convergence of the conditional maximization procedures for updating γ and λ . Specifically, iterations continue until the Euclidean distance between two consecutive log likelihoods falls under <code>eps2</code> . <code>eps2=1e-3</code> by default.
	K Number of Gauss-Legendre quadrature points for the E-step. Default 21.
	C [-C, C] sets the range of integral in the E-step. <code>C=5</code> by default.

Value

An object of class `hltm`.

<code>coefficients</code>	A data frame of parameter estimates, standard errors, z values and p values.
<code>scores</code>	A data frame of EAP estimates of latent preferences and their approximate standard errors.
<code>vcov</code>	Variance-covariance matrix of parameter estimates.
<code>log_Lik</code>	The log-likelihood value at convergence.
<code>J</code>	The number of items.
<code>p</code>	The number of predictors for the mean equation.
<code>q</code>	The number of predictors for the variance equation.
<code>item_names</code>	Names of items.
<code>call</code>	The matched call.

References

Zhou, Xiang. 2017. "Hierarchical Item Response Models for Analyzing Public Opinion." Working Paper.

Examples

```

y <- nes_econ2008[, -(1:3)]
x <- model.matrix( ~ party * educ, nes_econ2008)
z <- model.matrix( ~ party, nes_econ2008)

# don't run
# nes_m1 <- hgrm(y, x, z)

dichotomize <- function(x) findInterval(x, c(mean(x, na.rm = TRUE)))
y_bin <- as.data.frame(lapply(y, dichotomize))
nes_m1 <- hltm(y_bin, x, z)
print(nes_m1)

```

latent_scores	<i>Estimates of Latent Preferences/Abilities</i>
---------------	--

Description

EAP estimates of latent preferences for either hltm or hgrm models.

Usage

```
latent_scores(x, digits = 3)
```

Arguments

`x` An object of class hIRT
`digits` the number of significant digits to use when printing

Value

A data frame of EAP estimates of latent preferences and their approximate standard errors.

Examples

```

y <- nes_econ2008[, -(1:3)]
x <- model.matrix( ~ party * educ, nes_econ2008)
z <- model.matrix( ~ party, nes_econ2008)
nes_m1 <- hgrm(y, x, z)
pref <- latent_scores(nes_m1)
require(ggplot2)
ggplot(data = nes_econ2008) +
  geom_density(aes(x = pref$est, col = party))

```

nes_econ2008

Public Attitudes on Economic Issues in ANES 2008

Description

A dataset containing gender, party ID, education, and responses to 10 survey items on economic issues from the American National Election Studies, 2008.

Usage

nes_econ2008

Format

A data frame with 2268 rows and 13 variables:

gender gender. 1: male; 2: female

party party identification: Democrat, independent, or Republican

educ education. 1: high school or less; 2: some college or above

health_ins7 Support for government or private health insurance, 7 categories

jobs_guar7 Support for government guarantee jobs and income, 7 categories

gov_services7 Should government reduce or increase spending on services?, 7 categories

FS_poor3 Federal spending on the poor, 3 categories

FS_childcare3 Federal spending on child care, 3 categories

FS_crime3 Federal spending on crime, 3 categories

FS_publicschools3 Federal spending on public schools, 3 categories

FS_welfare3 Federal spending on welfare, 3 categories

FS_envir3 Federal spending on environment, 3 categories

FS_socsec3 Federal spending on Social Security, 3 categories

print.hIRT*Printing an object of class hIRT*

Description

Printing an object of class hIRT

Usage

```
## S3 method for class 'hIRT'  
print(x, digits = 3, ...)
```


Arguments

x	An object of class hIRT
digits	the number of significant digits to use when printing
...	further arguments passed to or from other methods

summary.hIRT

*Summarizing Hierarchical Item Response Theory Models***Description**

Summarizing the fit of either h1tm or hgrm.

Usage

```
## S3 method for class 'hIRT'
summary(object, by_item = FALSE, digits = 3, ...)

## S3 method for class 'summary_hIRT'
print(x, digits = 3, ...)
```

Arguments

object	An object of class hIRT.
by_item	Logical. Should item parameters be stored item by item (if TRUE) or put together in a data frame (if FALSE)?
digits	the number of significant digits to use when printing.
...	further arguments passed to or from other methods
x	An object of class hIRT

Value

An object of class summary_hIRT.

call	The matched call.
model	Model fit statistics: Log likelihood, AIC, and BIC.
item_coefs	Item parameter estimates, standard errors, z values, and p values.
mean_coefs	Parameter estimates for the mean equation.
var_coefs	Parameter estimates for the variance equation.

Examples

```
y <- nes_econ2008[, -(1:3)]
x <- model.matrix(~ party * educ, nes_econ2008)
z <- model.matrix(~ party, nes_econ2008)
nes_m1 <- hgrm(y, x, z)
summary(nes_m1, by_item = TRUE)
```

Index

*Topic **datasets**

nes_econ2008, 8

coef.hIRT, 2

coef_item, 3

coef_mean (coef.hIRT), 2

coef_var (coef.hIRT), 2

hgrm, 4

hltm, 5

latent_scores, 7

nes_econ2008, 8

print.hIRT, 8

print.summary_hIRT (summary.hIRT), 9

summary.hIRT, 9