Package 'catR'

June 13, 2010

Type Package

Title Procedures to generate IRT adaptive tests (CAT)

Version 1.0

Date 2010-06-13

Author David Magis (U Liege, Belgium), Gilles Raiche (UQAM, Canada)

Maintainer David Magis <david.magis@ulg.ac.be>

Depends R (>= 2.8.0)

Description The catR package allows the generation of response patterns under computerized adaptive testing (CAT) framework, with the choice of several starting rules, stopping rules and ability estimators.

License GPL (version 2 or later)

LazyLoad yes

R topics documented:

`createItemBank` *Item bank generation*

Description

This command creates an item bank from a matrix of item parameters. Item information functions are evaluated for all items and a fine grid of ability levels, to be supplied.

Usage

```
createItemBank(items=100, model="4PL", thMin=-4, thMax=4,
step=0.01, seed=1)
```

Arguments

`items`	either an integer value or a matrix of item parameters. See **Details**.
`model`	character: the name of the logistic IRT model, with possible values `"1PL"`, `"2PL"`, `"3PL"` or `"4PL"` (default). Ignored if `items` is a matrix.
`thMin`	numeric: the lower bound for the fine grid of ability levels (default is -4). See **Details**.
`thMax`	numeric: the upper bound for the fine grid of ability levels (default is 4). See **Details**.
`step`	numeric: the step value for the fine grid of ability levels (default is 0.01). See **Details**.
`seed`	numeric: the random seed number for the generation of item parameters (default is 1). See `set.seed` for further details.

Details

If `items` is a matrix, it has the following format: one row per item and four columns, with respectively the discrimination a_i, the difficulty b_i, the pseudo-guessing c_i and the inattention d_i parameters (Barton and Lord, 1981). If `items` is an integer, it corresponds to the number of items to be included in the item bank. Corresponding item parameters are (by default) randomly drawn from the following distributions: $a_i \sim N(1, 0.2^2)$, $b_i \sim N(0, 1)$, $c_i \sim U([0, 0.25])$ and $d_i \sim U([0.75, 1])$. Inattention parameters d_i are fixed to 1 if `model` is not `"4PL"`; pseudo-guessing parameters c_i are fixed to zero if `model` is either `"1PL"` or `"2PL"`; and discrimination parameters a_i are fixed to 1 if `model="1PL"`. The random generation of item parameters can be controlled by the `seed` argument.

The item bank consists of a matrix of item information functions (Baker, 1992), evaluated for each itelm in the bank and at each value of a sequence of ability levels. These abilities are ranging from `thMin` to `thMax` by steps of `step` units.

The returned list contains in addition the sequence of ability levels and the matrix of item parameters.

Value

A list of class "itBank" with three arguments:

`itemPar`	the matrix of item parameters, either provided by `items` or generated.
`theta`	a vector with the ability levels of the fine grid, defined by arguments `thMin`, `thMax` and `step`.
`infoTab`	a matrix of item information functions, evaluated for each ability level (one row per ability level) and each item (one column per item).

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

References

Baker, F.B. (1992). *Item response theory: parameter estimation techniques*. New York, NY: Marcel Dekker.

Barton, M.A., and Lord, F.M. (1981). An upper asymptote for the three-parameter logistic item-response model. Research Bulletin 81-20. Princeton, NJ: Educational Testing Service.

See Also

`Ii`

Examples

```
## Not run:

 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Item bank creation with 'tcals' item parameters
 createItemBank(tcals)

 # Changing the fine grid of ability levels
 createItemBank(tcals, thMin=-2, thMax=2, step=0.05)

 # Item bank creation with 500 items
 createItemBank(items=500)

## End(Not run)
```

`eapEst` *EAP ability estimation under the 4PL model*

Description

This command returns the EAP (expected a posteriori) ability estimate for a given matrix of item parameters of the 4PL model and a given response pattern.

Usage

```
eapEst(it, x, D=1, priorDist="norm", priorPar=c(0,1),
 lower=-4, upper=4, nqp=20)
```

Arguments

`it`	numeric: a matrix with one row per item and four columns, with the values of the discrimination, the difficulty, the pseudo-guessing and the inattention parameters (in this order).
`x`	numeric: a vector of dichotomous item responses.
`D`	numeric: the metric constant. Default is `D=1` (for logistic metric); `D=1.702` yields approximately the normal metric (Haley, 1952).
`priorDist`	character: specifies the prior distribution. Possible values are `"norm"` (default), `"unif"` and `"Jeffreys"`.
`priorPar`	numeric: vector of two components specifying the prior parameters (default is `c(0,1)`). Ignored if `priorDist="Jeffreys"`. See **Details**.
`lower`	numeric: the lower bound for numercal integration (default is -4).
`upper`	numeric: the upper bound for numercal integration (default is 4).
`nqp`	numeric: the number of quadrature points (default is 20).

Details

The EAP (expected a posteriori) ability estimator (Bock and Mislevy, 1982) is obtained by computing the average of the posterior distribution of ability, set as the prior distribution times the likelihood function.

Three prior distributions are available: the normal distribution, the uniform distribution and Jeffreys' prior distribution (Jeffreys, 1939, 1946). The prior distribution is specified by the argument `priorPar`, with values `"norm"`, `"unif"` and `"Jeffreys"`, respectively.

The argument `priorPar` determines either the prior mean and standard deviation of the normal prior distribution (if `priorDist="norm"`), or the range for defining the prior uniform distribution (if `priorDist="unif"`). This argument is ignored if `priorDist="Jeffreys"`.

The required integrals are approximated by numerical adaptive quadarture using the `integrate` function. Arguments `lower`, `upper` and `nqp` define respectively the lower and upper bounds for numerical integration, and the number of quadrature points. By default, the numerical integration runs with 20 quadrature points on the range [-4; 4].

Value

The estimated EAP ability level.

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

References

Bock, R. D., and Mislevy, R. J. (1982). Adaptive EAP estimation of ability in a microcomputer environment. *Applied Psychological Measurement, 6*, 431-444.

Haley, D.C. (1952). Estimation of the dosage mortality relationship when the dose is subject to error. Technical report no 15. Palo Alto, CA: Applied Mathematics and Statistics Laboratory, Stanford University.

Jeffreys, H. (1939). *Theory of probability*. Oxford, UK: Oxford University Press.

Jeffreys, H. (1946). An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 186*, 453-461.

See Also

thetaEst

Examples

```
## Not run:

 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Creation of a response pattern (tcals item parameters,
 # true ability level 0)
 set.seed(1)
 x <- rbinom(85, 1, Pi(0, tcals)$Pi)

 # EAP estimation, standard normal prior distribution
 eapEst(tcals, x)

 # EAP estimation, uniform prior distribution upon range [-2,2]
 eapEst(tcals, x, priorDist="unif", priorPar=c(-2,2))

 # EAP estimation, Jeffreys' prior distribution
 eapEst(tcals, x, priorDist="Jeffreys")

 # Changing the integration settings
 eapEst(tcals, x, nqp=100)

## End(Not run)
```

`eapSem` *Standard error of EAP ability estimation under the 4PL model*

Description

This command returns the estimated standard error of the ability estimate, for a given matrix of item parameters of the 4PL model, an ability estimate and a specified estimator.

Usage

```
eapSem(thEst, it, x, D=1, priorDist="norm", priorPar=c(0,1),
 lower=-4, upper=4, nqp=20)
```

Arguments

`thEst`	numeric: the EAP ability estimate.
`it`	numeric: a matrix with one row per item and four columns, with the values of the discrimination, the difficulty, the pseudo-guessing and the inattention parameters (in this order).
`x`	numeric: a vector of dichotomous item responses.
`D`	numeric: the metric constant. Default is `D=1` (for logistic metric); `D=1.702` yields approximately the normal metric (Haley, 1952).
`priorDist`	character: specifies the prior distribution. Possible values are `"norm"` (default), `"unif"` and `"Jeffreys"`.
`priorPar`	numeric: vector of two components specifying the prior parameters (default is `c(0,1)`). Ignored if `priorDist="Jeffreys"`. See **Details**.
`lower`	numeric: the lower bound for numercal integration (default is -4).
`upper`	numeric: the upper bound for numercal integration (default is 4).
`nqp`	numeric: the number of quadrature points (default is 20).

Details

This command computes the standard error of the EAP (expected a posteriori) ability estimator (Bock and Mislevy, 1982).

Three prior distributions are available: the normal distribution, the uniform distribution and Jeffreys' prior distribution (Jeffreys, 1939, 1946). The prior distribution is specified by the argument `priorPar`, with values `"norm"`, `"unif"` and `"Jeffreys"`, respectively.

The argument `priorPar` determines either the prior mean and standard deviation of the normal prior distribution (if `priorDist="norm"`), or the range for defining the prior uniform distribution (if `priorDist="unif"`). This argument is ignored if `priorDist="Jeffreys"`.

The required integrals are approximated by numerical adaptive quadarture using the `integrate` function. Arguments `lower`, `upper` and `nqp` define respectively the lower and upper bounds for numerical integration, and the number of quadrature points. By default, the numerical integration runs with 20 quadrature points on the range [-4; 4].

Note that in the current version, the EAP ability estimate must be specified through the `thEst` argument.

Value

The estimated standard error of the EAP ability level.

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

References

Bock, R. D., and Mislevy, R. J. (1982). Adaptive EAP estimation of ability in a microcomputer environment. *Applied Psychological Measurement, 6*, 431-444.

Haley, D.C. (1952). Estimation of the dosage mortality relationship when the dose is subject to error. Technical report no 15. Palo Alto, CA: Applied Mathematics and Statistics Laboratory, Stanford University.

Jeffreys, H. (1939). *Theory of probability*. Oxford, UK: Oxford University Press.

Jeffreys, H. (1946). An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 186*, 453-461.

See Also

thetaEst

Examples

```
## Not run:

 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Creation of a response pattern (tcals item parameters,
 # true ability level 0)
 set.seed(1)
 x <- rbinom(85, 1, Pi(0, tcals)$Pi)

 # EAP estimation, standard normal prior distribution
 th <- eapEst(tcals, x)
 c(th, eapSem(th, tcals, x))

 # EAP estimation, uniform prior distribution upon range [-2,2]
 th <- eapEst(tcals, x, priorDist="unif", priorPar=c(-2,2))
 c(th, eapSem(th, tcals, x, priorDist="unif", priorPar=c(-2,2)))

 # EAP estimation, Jeffreys' prior distribution
 th <- eapEst(tcals, x, priorDist="Jeffreys")
 c(th, eapSem(th, tcals, x, priorDist="Jeffreys"))

## End(Not run)
```

Ii *Item information functions, first and second derivatives (4PL)*

Description

This command returns the item information functions for a given matrix of item parameters of the 4PL model and a given ability value. Numerical values of the first and second derivatives of the item information functions are also returned.

Usage

```
Ii(th, it, D=1)
```

Arguments

th	numeric: the ability value.
it	numeric: a matrix with one row per item and four columns, with the values of the discrimination, the difficulty, the pseudo-guessing and the inattention parameters (in this order).
D	numeric: the metric constant. Default is D=1 (for logistic metric); D=1.702 yields approximately the normal metric (Haley, 1952).

Details

The first and second derivatives are computed algebraically from the four-parameter logistic (4PL) model (Barton and Lord, 1981). These derivatives are necessary for both the estimation of ability and the computation of related standard errors.

Value

A list with three arguments:

Ii	the vector with item informations (one value per item)
dIi	the vector with first derivatives of the item information functions (one value per item)
d2Ii	the vector with second derivatives of the item information functions (one value per item)

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

References

Barton, M.A., and Lord, F.M. (1981). An upper asymptote for the three-parameter logistic item-response model. Research Bulletin 81-20. Princeton, NJ: Educational Testing Service.

Haley, D.C. (1952). Estimation of the dosage mortality relationship when the dose is subject to error. Technical report no 15. Palo Alto, CA: Applied Mathematics and Statistics Laboratory, Stanford University.

See Also

Pi, thetaEst

Examples

```
## Not run:

 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Response probabilities and derivatives (various th and D values)
 Ii(th=0, tcals)
 Ii(th=0, tcals, D=1.702)
 Ii(th=1, tcals)

## End(Not run)
```

`nextItem` *Selection of the next item*

Description

This command selects the next item to be administered, given the list of previously administered items and the current ability estimate, with the 'maximum information' criterion.

Usage

```
nextItem(itemBank, theta, out=NULL)
```

Arguments

`itemBank`	an item bank of class `itBank` sa output of the function `createItemBank`.
`theta`	numeric: the current value of the ability estimate.
`out`	either a vector of integer values specifying the items previously administered, or `NULL` (default).

Details

Currently only one ethod is available for selecting the next item to be administered in the adaptive test. For a given current ability estimate, the next item is selected (among the available items) as the one which maximizes the item information function (Baker, 1992). The most informative item is selected from the table of item informations provided by the bank of items specified with `itemBank`.

The available items are those that are not specified in the `out` argument. By default, `out` is `NULL`, which means that all items are available.

Value

A list with three arguments:

`item`	the selected item (identified by its number in the item bank).
`par`	the vector of item parameters of the selected item.
`info`	the value of the information function for the selected item and the current ability estimate.

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

References

Baker, F.B. (1992). *Item response theory: parameter estimation techniques*. New York, NY: Marcel Dekker.

See Also

`createItemBank`

Examples

```
## Not run:

 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Item bank creation with 'tcals' item parameters
 bank <- createItemBank(tcals)

 # Selecting the next item, current ability estimate is 0
 nextItem(bank, 0) # item 63 is selected

 # Selecting the next item, current ability estimate is 0
 # and item 63 is removed
 nextItem(bank, 0, out=63) # item 10 is selected

 # Selecting the next item, current ability estimate is 0
 # and items 63 and 10 are removed
```

```
 nextItem(bank, 0, out=c(63,10)) # item 62 is selected

## End(Not run)
```

`Pi` *Item response probabilities, first, second and third derivatives (4PL)*

Description

This command returns the item reponse probabilities for a given matrix of item parameters of the 4PL model and a given ability value. Numerical values of the first, second and third derivatives of the response probabilities are also returned.

Usage

```
Pi(th, it, D=1)
```

Arguments

`th`	numeric: the ability value.
`it`	numeric: a matrix with one row per item and four columns, with the values of the discrimination, the difficulty, the pseudo-guessing and the inattention parameters (in this order).
`D`	numeric: the metric constant. Default is `D=1` (for logistic metric); `D=1.702` yields approximately the normal metric (Haley, 1952).

Details

The first, second and third derivatives are computed algebraically from the four-parameter logistic (4PL) model (Barton and Lord, 1981). These derivatives are necessary for both the estimation of ability and the computation of related standard errors.

Value

A list with four arguments:

`Pi`	the vector with response probabilities (one value per item)
`dPi`	the vector with first derivatives of the response probabilities (one value per item)
`d2Pi`	the vector with second derivatives of the response probabilities (one value per item)
`d3Pi`	the vector with third derivatives of the response probabilities (one value per item)

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

References

Barton, M.A., and Lord, F.M. (1981). An upper asymptote for the three-parameter logistic item-response model. Research Bulletin 81-20. Princeton, NJ: Educational Testing Service.

Haley, D.C. (1952). Estimation of the dosage mortality relationship when the dose is subject to error. Technical report no 15. Palo Alto, CA: Applied Mathematics and Statistics Laboratory, Stanford University.

See Also

Ii, thetaEst

Examples

```
## Not run:

 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Response probabilities and derivatives (various th and D values)
 Pi(th=0, tcals)
 Pi(th=0, tcals, D=1.702)
 Pi(th=1, tcals)

## End(Not run)
```

`randomCAT` *Random generation of adaptive tests*

Description

This command generates a response pattern to an adaptive test, for agiven item bank, a true ability level, and several lists of CAT parameters (starting items, stopping rule, provisional and final ability estimators).

Usage

```
 randomCAT(trueTheta, itemBank, maxItems=50,
  start=list(fixItems=NULL, seed=NULL, nrItems=1, theta=0,
  bw=4, range=c(-4,4)), test=list(method="BM",
  priorDist="norm", priorPar=c(0,1), range=c(-4,4), D=1,
  eapPar=c(-4,4,20)), stop=list(rule="length", thr=20,
alpha=0.05), final=list(method="BM", priorDist="norm",
priorPar=c(0,1), range=c(-4,4),D=1, eapPar=c(-4,4,20),
alpha=0.05))
 ## S3 method for class 'cat':
print(x, ...)
 ## S3 method for class 'cat':
plot(x, ci=TRUE, alpha=0.05, thr=NULL, ...)
```

Arguments

trueTheta	numeric: the value of the true ability level.
itemBank	an item bank, i.e. a list of class itBank, typically an output of the function createItemBank.
maxItems	numeric: the maximal number of items to be administered (default is 50).
start	a list with the options for starting the adaptive test. See **Details**.
test	a list with the options for provisional ability estimation. See **Details**.
stop	a list with the options of the stopping rule. See **Details**.
final	a list with the options for final ability estimation. See **Details**.
x	an object of class "cat", typically an output of randomCAT function.
ci	logical: should the confidence intervals be plotted for each provisional ability estimate? (default is TRUE).
alpha	numeric: the significance level for provisional confidence intervals (default is 0.05). Ignored if ci is FALSE.
thr	either a numeric value specifying the threshold to be plotted, or NULL (default) for not displaying any threshold.
...	other generic arguments to be passed to print and plot functions.

Details

The randomCAT function generates an adaptive test using an item bank specified by argument itemBank, and for a given true ability level specified by argument trueTheta. The maximal length of the test can be fixed through the maxItems argument, with a default value of 50 items.

The test specification is made by means of four lists of options: one list for the selection of the starting items, one list with the options for provisional ability estimation, one list to define the stopping rule, and one list with the options for final ability estimation. These lists are specified respectively by the arguments start, test, stop and final.

The start list can contain one or several of the following arguments:

- fixItems: either a vector of integer values, setting the items to be administered as first items, or NULL (default) to let the function select the items.
- seed: either a numeric value to fix the random seed for item selection, or NULL (default) to select the items on the basis of their difficulty level. Ignored if fixItems is not NULL.
- nrItems: the number of first items to be selected (default is 1). Ignored if fixItems is not NULL.
- theta: the central initial ability value, used to define the range of ability levels for selecting the first items (default is 0). Ignored if either fixItems or seed is not NULL. See startItems for further details.
- bw: the bandwidth value, used to define the range of ability levels for selecting the first items (default is 4). Ignored if either fixItems or seed is not NULL. See startItems for further details.
- range: the maximal range of difficulty levels, specified as a vector of two numeric values (default is c(-4,4)). Ignored if either fixItems or seed is not NULL. See startItems for further details.

These arguments are passed to the function startItems to select the first items of the test.

The test list can contain one or several of the following arguments:

- method: a character string to specify the method for ability estimation. Possible values are: "BM" (default) for Bayesian modal estimation (Birnbaum, 1969), "ML" for maximum likelihood estimation (lord, 1980), "EAP" for expected a posteriori (EAP) estimation (Bock and Mislevy, 1982), and "WL" for weighted likelihood estimation (Warm, 1989).
- priorDist: a character string which sets the prior distribution. Possible values are: "norm" (default) for normal distribution, "unif" for uniform distribution, and "Jeffreys" for Jeffreys' noninformative prior distribution (Jeffreys, 1939, 1946). ignored if method is neither "BM" nor "EAP".
- priorPar: a vector of two numeric components, which sets the parameters of the prior distribution. If (method="BM" or method=="EAP") and priorDist="norm", the components of priorPar are respectively the mean and the standard deviation of the prior normal density. If (method="BM" or method=="EAP") and priorDist="unif", the components of priorPar are respectively the lower and upper bound of the prior uniform density. Ignored in all other cases. By default, priorPar takes the parameters of the prior standard normal distribution (i.e., priorPar=c(0,1)).
- range: the maximal range of ability levels, set as a vector of two numeric components. The ability estimate will always lie to this interval (set by default to [-4, 4]). Ignored if method=="EAP".
- D: the value of the metric constant. Default is D=1 for logistic metric. Setting D=1.702 yields approximately the normal metric (Haley, 1952).
- eapParnumeric: vector of three components, holding respectively the values of the arguments lower, upper and nqp of the eapEst command. Default vector is (-4, 4, 20). Ignored if method is not "EAP".

These arguments are passed to the functions thetaEst and semTheta to estimate the ability level and the standard error of this estimate.

The stop list can contain one or several of the following arguments:

- rule: a character string specifying the type of stopping rule. Possible values are: "length" (default), to stop the test after a pre-specified number of items administered; "precision", to stop the test when the provisional standard error of ability becomes less than or equal to the pre-specified value; and "classification", for which the test ends whenever the provisional confidence interval (set by the alpha argument) does not hold the classification threshold anymore.
- thr: a numeric value fixing the threshold of the stopping rule. If rule="length", thr is the maximal number of items to be administered (in practice, it is replaced by the value of the maxItems argument if the latter is smaller than thr). If rule="precision", thr is the precision level (i.e. the standard error) to be reached before stopping. Finally, if rule="classification", thr corresponds to the ability level which serves as a classification rule (i.e. which must not be covered by the provisional confidence interval).
- alpha: the significance (or α) level for computing the priovisional confidence interval of ability. Ignored if rule is not "classification".

Eventually, the final list can contain one or several arguments of the test list (with possiblly different values), as well as the additional alpha argument. The latter specifies the α level of the final confidence interval of ability, which is computed as

$$[\hat{\theta} - z_{1-\alpha/2}\ se(\hat{\theta}); \hat{\theta} + z_{1-\alpha/2}\ se(\hat{\theta})]$$

where $\hat{\theta}$ and $se(\hat{\theta})$ are respectively the ability estimate and its standard error.

Note that if some arguments of these lists are missing, they are automatically set to their default value. The contents of the lists is checked with the testList function, and the adaptive test is generated only if the lists are adequately defined.

The function plot.cat represents the set of provisional and final ability estimates throughout the test. Corresponding confidence intervals (with confidence level defined by the argument alpha) are also drawn if ci=TRUE (which is the default value). In addition, a classification threshold can be displayed by specifying a numeric value to the argument thr. The default value NULL does not represent any threshold.

Value

The function randomCAT returns a list of class "cat" with the following arguments:

trueTheta	the value of the trueTheta argument.
maxItems	the value of the maxItems argument.
testItems	a vector with the items that were administered during the test.
itemPar	a matrix with the parameters of the items administered during the test.
pattern	the generated response pattern (as vector of 0 and 1 entries).
thetaProv	a vector with the provisional ability estimates.
seprov	a vector with the standard errors of the provisional ability estimates.
thFinal	the final ability estimate.
seFinal	the standrad error of the final ability estimate.
ciFinal	the confidence interval of the final ability estimate.
startFixItems	the value of the start$fixItems argument (or its default value if missing).
startSeed	the value of the start$seed argument (or its default value if missing).
startNrItems	the value of the start$nrItems argument (or its default value if missing).
startTheta	the value of the start$theta argument (or its default value if missing).
startBw	the value of the start$bw argument (or its default value if missing).
startbOpt	the value of the start$bOpt argument (or its default value if missing).
startRange	the value of the start$range argument (or its default value if missing).
provMethod	the value of the test$method argument (or its default value if missing).
provDist	the value of the test$priorDist argument (or its default value if missing).
provPar	the value of the test$priorPar argument (or its default value if missing).
provRange	the value of the test$range argument (or its default value if missing).
provD	the value of the test$D argument (or its default value if missing).
stopRule	the value of the stop$rule argument (or its default value if missing).
stopThr	the value of the stop$thr argument (or its default value if missing).
stopAlpha	the value of the stop$alpha argument (or its default value if missing).
endWarning	a logical indactor indicating whether the adaptive test stopped because the stopping rule was satisfied or not.
finalMethod	the value of the final$method argument (or its default value if missing).
finalDist	the value of the final$priorDist argument (or its default value if missing).
finalPar	the value of the final$priorPar argument (or its default value if missing).

finalRange	the value of the final$range argument (or its default value if missing).
finalD	the value of the final$D argument (or its default value if missing).
finalAlpha	the value of the final$alpha argument (or its default value if missing).

The function print.cat returns similar (but differently organized) results.

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

References

Birnbaum, A. (1969). Statistical theory for logistic mental test models with a prior distribution of ability. *Journal of Mathematical Psychology, 6*, 258-276.

Bock, R. D., and Mislevy, R. J. (1982). Adaptive EAP estimation of ability in a microcomputer environment. *Applied Psychological Measurement, 6*, 431-444.

Haley, D.C. (1952). Estimation of the dosage mortality relationship when the dose is subject to error. Technical report no 15. Palo Alto, CA: Applied Mathematics and Statistics Laboratory, Stanford University.

Jeffreys, H. (1939). *Theory of probability*. Oxford, UK: Oxford University Press.

Jeffreys, H. (1946). An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 186*, 453-461.

Lord, F.M. (1980). *Applications of item response theory to practical testing problems*. Hillsdale, NJ: Lawrence Erlbaum.

Warm, T.A. (1989). Weighted likelihood estimation of ability in item response models. *Psychometrika, 54*, 427-450.

See Also

testList, startItems, nextItem, thetaEst, semTheta

Examples

```
## Not run:
 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Item bank creation with 'tcals' item parameters
 bank <- createItemBank(tcals)

 # Creation of a starting list: 5 items, initial theta 0, bw 2
 start<-list(nrItems=5, theta=0, bw=2)

 # Creation of 'test' and 'final' lists: weighted likelihood
 # estimation of ability (both provisional and final)
 test<-list(method="WL")
 final<-test

 # Creation of a stopping rule: precision criterion, standard
```

```
# error to be reached 0.3
stop<-list(rule="precision", thr=0.3)

# CAT test
res<-randomCAT(0, bank, start=start, test=test, stop=stop,
 final=final)

# New 'test' and 'final' rules (BM and EAP estimation
# with Jeffreys' prior)
test2<-list(method="BM", priorDist="Jeffreys")
final2<-list(method="EAP", priorDist="Jeffreys")

# New stopping rule: classification criterion, with
# classification threshold 0 and alpha level 0.05
stop2<-list(rule="classification", thr=0, alpha=0.05)

# CAT test with new 'test', 'stop' and 'final' rules
res2<-randomCAT(0, bank, start=start, test=test2, stop=stop2,
 final=final2)

# New stopping rule: classification criterion, with
# classification threshold 0.5 and alpha level 0.05
stop3<-list(rule="classification", thr=0.5, alpha=0.05)

# CAT test with new 'stop' rule
res3<-randomCAT(0, bank, start=start, test=test2, stop=stop3,
 final=final2)

# Plotting results
plot(res)
plot(res, ci=FALSE)
plot(res, thr=0)
plot(res, thr=0.5)

plot(res2,thr=0)

plot(res3, thr=0.5)

## End(Not run)
```

`semTheta` *Standard error of ability estimation under the 4PL model*

Description

This command returns the estimated standard error of the ability estimate, for a given matrix of item parameters of the 4PL model, an ability estimate and a specified estimator.

Usage

```
semTheta(thEst, it, x=NULL, D=1, method="BM", priorDist="norm",
 priorPar=c(0,1), eapPar=c(-4,4,20))
```

Arguments

thEst	numeric: the ability estimate.
it	numeric: a matrix with one row per item and four columns, with the values of the discrimination, the difficulty, the pseudo-guessing and the inattention parameters (in this order).
x	numeric: a vector of dichotomous item responses (default is NULL). Ignored if method is not "EAP".
D	numeric: the metric constant. Default is D=1 (for logistic metric); D=1.702 yields approximately the normal metric (Haley, 1952).
method	character: the ability estimator. Possible values are "BM" (default), "ML" and "WL". See **Details**.
priorDist	character: specifies the prior distribution. Possible values are "norm" (default), "unif" and "Jeffreys". Ignored if method is neither "BM" nor "EAP". See **Details**.
priorPar	numeric: vector of two components specifying the prior parameters (default is c(0,1)). Ignored if method is neither "BM" nor "EAP", or if priorDist="Jeffreys". See **Details**.
eapPar	numeric: vector of three components, holding respectively the values of the arguments lower, upper and nqp of the eapEst command. Default vector is (-4, 4, 20). Ignored if method is not "EAP".

Details

Four ability estimators are available: the maximum likelihood (ML) estimator (Lord, 1980), the Bayes modal (BM) estimator (Birnbaum, 1969), the expected a posteriori (EAP) estimator (Bock and Mislevy, 1982)and the weighted likelihood (WL) estimator (Warm, 1989). The selected estimator is specified by the method argument, with values "ML", "BM", "EAP" and "WL" respectively.

For the BM and EAP estimators, three prior distributions are available: the normal distribution, the uniform distribution and the Jeffreys' prior distribution (Jeffreys, 1939, 1946). The prior distribution is specified by the argument priorPar, with values "norm", "unif" and "Jeffreys", respectively. The priorPar argument is ignored if method="ML" or method="WL".

The argument priorPar determines either: the prior mean and standard deviation of the normal prior distribution (if priorDist="norm"), or the range for defining the prior uniform distribution (if priorDist="unif"). This argument is ignored if priorDist="Jeffreys".

The eapPar argument sets the range and the number of quadrature points for numerical integration in the EAP process. By default, it takes the vector value (-4, 4, 20), that is, 20 quadrature points on the range [-4; 4]. See eapEst for further details.

Note that in the current version, the ability estimate must be specified through the thEst argument. Moreover, the response pattern must be specified through the x argument to compute the standard error of the EAP estimate. For the other estimation methods, thi is not necessary, and x is set to NULL by default for this purpose.

Value

The estimated standard error of the ability level.

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

References

Barton, M.A., and Lord, F.M. (1981). An upper asymptote for the three-parameter logistic item-response model. Research Bulletin 81-20. Princeton, NJ: Educational Testing Service.

Birnbaum, A. (1969). Statistical theory for logistic mental test models with a prior distribution of ability. *Journal of Mathematical Psychology, 6*, 258-276.

Bock, R. D., and Mislevy, R. J. (1982). Adaptive EAP estimation of ability in a microcomputer environment. *Applied Psychological Measurement, 6*, 431-444.

Haley, D.C. (1952). Estimation of the dosage mortality relationship when the dose is subject to error. Technical report no 15. Palo Alto, CA: Applied Mathematics and Statistics Laboratory, Stanford University.

Jeffreys, H. (1939). *Theory of probability*. Oxford, UK: Oxford University Press.

Jeffreys, H. (1946). An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 186*, 453-461.

Lord, F.M. (1980). *Applications of item response theory to practical testing problems*. Hillsdale, NJ: Lawrence Erlbaum.

Warm, T.A. (1989). Weighted likelihood estimation of ability in item response models. *Psychometrika, 54*, 427-450.

See Also

eapSem, thetaEst

Examples

```
## Not run:

 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Creation of a response pattern (tcals item parameters,
 # true ability level 0)
 set.seed(1)
 x <- rbinom(85, 1, Pi(0, tcals)$Pi)

 # ML estimation
 th <- thetaEst(tcals, x, method="ML")
 c(th, semTheta(th, tcals, method="ML"))

 # BM estimation, standard normal prior distribution
 th <- thetaEst(tcals, x)
 c(th, semTheta(th, tcals))

 # BM estimation, uniform prior distribution upon range [-2,2]
 th <- thetaEst(tcals, x, method="BM", priorDist="unif",
```

```
                priorPar=c(-2,2))
c(th, semTheta(th, tcals, method="BM", priorDist="unif",
   priorPar=c(-2,2)))

# BM estimation, Jeffreys' prior distribution
th <- thetaEst(tcals, x, method="BM", priorDist="Jeffreys")
c(th, semTheta(th, tcals, method="BM", priorDist="Jeffreys"))

# EAP estimation, standard normal prior distribution
th <- thetaEst(tcals, x, method="EAP")
c(th, semTheta(th, tcals, x, method="EAP"))

# EAP estimation, uniform prior distribution upon range [-2,2]
th <- thetaEst(tcals, x, method="EAP", priorDist="unif",
                priorPar=c(-2,2))
c(th, semTheta(th, tcals, x, method="EAP", priorDist="unif",
   priorPar=c(-2,2)))

# EAP estimation, Jeffreys' prior distribution
th <- thetaEst(tcals, x, method="EAP", priorDist="Jeffreys")
c(th, semTheta(th, tcals, x, method="EAP", priorDist="Jeffreys"))

# WL estimation
th <- thetaEst(tcals, x, method="WL")
c(th, semTheta(th, tcals, method="WL"))

## End(Not run)
```

`startItems` *Selection of the first items*

Description

This command selects the first items of the adaptive test, either randomly or on the basis of their difficulty level.

Usage

```
startItems(it, fixItems=NULL, seed=NULL, nrItems=1,
  theta=0, bw=4, range=c(-4,4))
```

Arguments

`it`	numeric: a matrix with one row per item and four columns, with the values of the discrimination, the difficulty, the pseudo-guessing and the inattention parameters (in this order).
`fixItems`	either a vector of integer values or `NULL` (default). See **Details**.
`seed`	either a numeric value or `NULL` (default). See **Details**.
`nrItems`	numeric: the number of starting items to be selected (default is 1). Ignored if `fixItems` is not `NULL`.

theta	numeric: the initial ability level for selecting the first items (default is 0). Ignored if either fixItems or seed is not NULL. See **Details**.
bw	numeric: the length (bandwidth) of the difficulty interval for selecting the first items (default is 4). Ignored if either fixItems or seed is not NULL. See **Details**.
range	numeric: the range of diffiulty levels for selecting the first items (default is c(-4,4)). Ignored if either fixItems or seed is not NULL. See **Details**.

Details

The first items of the adaptive test can be selected either by fixing these items or by some selection into the item bank. The first items can be fixed by specifying fixItems as a vector of integer values, each value referring to the corresponding item in the item bank.

If fixItems is kep to NULL (the default choice), then nrItems will be selected, either randomly or according to their difficulty level. In the former, one has to assign a numeric value to the seed argument. In the latter, seed is kept to NULL (the default value) and the items are selected as follows. The arguments theta and bw permit to define a sequence of value, ranging from theta - bw to theta + bw and of length nrItems (if nrItems equals 1, the range reduces to theta). Then, for each value in the sequence, the function selects the item in the item bank with difficulty level as close as possible to the sequential value. By default, theta equals 0 and bw equals 4, meaning that the sequence ranges from -4 to 4 and of length nrItems.

The argument range permits to constraint the range of difficulty levels, by default to the range [-4, 4].

Value

A list with three arguments:

items	the selected items (identified by their number in the item bank).
par	the matrix of item parameters of the selected items (one row per item).
bOpt	the values of the item difficulties chosen for selecting the first items.

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

Examples

```
## Not run:

 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Selecting 1 starting item, initial ability estimate is 0
 startItems(tcals)

 # Selecting 3 starting items, initial ability estimate is 1
 # and bandwidth is 2
 startItems(tcals, nrItems=3, theta=1, bw=2)
```

```
# Selecting 5 starting items, initial ability estimate is 2
# and bandwidth is 3
startItems(tcals, nrItems=5, theta=2, bw=3)

# Random selection of 4 starting items
startItems(tcals, seed=1, nrItems=4)

# Selection of the first 5 starting items
startItems(tcals, fixItems=1:5)

## End(Not run)
```

`tcals` *Items parameters of the TCALS 1998 data set*

Description

The TCALS (*Test d'Aptitude en Anglais Langue Seconde*) is an aptitude test of English language as a second language in the French speaking college of Outaouais (Gatineau, QC, Canada). The test consists of 85 items and is administered every year to newly incoming students. The item parameters of the year 1998 have been estimated under the 3PL model. Inattention parameters are therefore fixed to one.

Format

A matrix with 85 rows and four columns, respectively holding the discrimination, difficulty, pseudo-guessing a,nd inattention parameters as calibrated on the results of the 1998 application of the TCALS questionnaire.

Source

The TCALS test was originally developed by Laurier, Froio, Pearo and Fournier (1998) and item parameters were obtained from Raiche (2002)..

References

Laurier, M., Froio, L., Pearo C., and Fournier, M. (1998). Test de classement d'anglais langue seconde au collegial. Montreal, Canada: College de Maisonneuve.

Raiche, G. (2002). Le depistage du sous-classement aux tests de classement en anglais, langue seconde, au collegial [The detection of under classification to the collegial English as a second language placement tests]. Gatineau, QC: College de l'Outaouais.

testList *Testing the format of the input lists*

Description

This command tests whether format of the input lists for the random generation of adaptive tests is convenient, and returns a warning message otherwise.

Usage

```
testList(list, type="start")
```

Arguments

list	a list of arguments to be tested. See **Details**.
type	character: the type of list for checking. Possible values are "start" (default), "test", "stop" and "final". See **Details**.

Details

The testList function checks whether the list provided in the list argument is accurate for the selected type. It mainly serves as an initial check for the randomCAT function.

The four types of lists are: "start" with the parameters for selecting the first items; "test" with the options of the adaptive test (i.e. provisional ability estimator and related information); "stop" with the options setting the stopping rule; and "final" with the options for final ability estimation. See the help file of randomCAT for further details about the different lists, their allowed arguments and their contents.

The function returns an "ok" message if the arguments of list match the requirement of the corresponding type. Otherwise, a message is returned with information about list - type mismatch. This will be the case if:

- list is not a list, or has no argument names,
- list has too many arguments for the type specified,
- one of the argument names is incorrect,
- the contents of an argument is not adequate (e.g. character instead of numeric).

Each mismatch yields a different output message to help in debugging the problem.

Value

A list with two arguments:

test	a logical value indicating whether the format of the list is accurate (TRUE) or not (FALSE).
message	either a message to indicate the type of misspecification, or "ok" if the format is accurate.

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

See Also

randomCAT

Examples

```
## Not run:

 # Creation and test of a 'start' list
 start <- list(nrItems=3, theta=0, bw=2)
 testList(start, type="start")

 # Modification of the list to introduce a mistake
 names(start)[1] <- "nrItem"
 testList(start, type="start")

 # Creation and test of a 'test' list
 test <- list(method="WL")
 testList(test, type="test")

# Creation and test of a 'stop' list
 stop <- list(method="WL")
 testList(stop, type="test")

# Creation and test of a 'final' list (with mistake)
 final <- list(method="MAP")
 testList(final, type="final")

## End(Not run)
```

---

thetaEst *Ability estimation under the 4PL model*

Description

This command returns the ability estimate for a given matrix of item parameters of the 4PL model and a given response pattern. Available estimators are maximum likelihood, Bayes modal, expected a posteriori (EAP) and weighted likelihood.

Usage

```
thetaEst(it, x, D=1, method="BM", priorDist="norm",
 priorPar=c(0,1), range=c(-4,4),
 eapPar=c(-4,4,20))
```

Arguments

it	numeric: a matrix with one row per item and four columns, with the values of the discrimination, the difficulty, the pseudo-guessing and the inattention parameters (in this order).
x	numeric: a vector of dichotomous item responses.
D	numeric: the metric constant. Default is D=1 (for logistic metric); D=1.702 yields approximately the normal metric (Haley, 1952).
method	character: the ability estimator. Possible values are "BM" (default), "ML" and "WL". See **Details**.
priorDist	character: specifies the prior distribution. Possible values are "norm" (default), "unif" and "Jeffreys". Ignored if method is neither "BM" nor "EAP". See **Details**.
priorPar	numeric: vector of two components specifying the prior parameters (default is c(0,1)). Ignored if method is neither "BM" nor "EAP", or if priorDist="Jeffreys". See **Details**.
range	numeric: vector of two components specifying the range wherein the ability estimate must be looked for (default is c(-4,4)). Ignored if method=="EAP".
eapPar	numeric: vector of three components, holding respectively the values of the arguments lower, upper and nqp of the eapEst command. Default vector is (-4, 4, 20). Ignored if method is not "EAP".

Details

Four ability estimators are available: the maximum likelihood (ML) estimator (Lord, 1980), the Bayes modal (BM) estimator (Birnbaum, 1969), the expected a posteriori (EAP) estimator (Bock and Mislevy, 1982) and the weighted likelihood (WL) estimator (Warm, 1989). The selected estimator is specified by the method argument, with values "ML", "BM", "EAP" and "WL" respectively.

For the BM and EAP estimators, three prior distributions are available: the normal distribution, the uniform distribution and Jeffreys' prior distribution (Jeffreys, 1939, 1946). The prior distribution is specified by the argument priorPar, with values "norm", "unif" and "Jeffreys", respectively. The priorPar argument is ignored if method="ML" or method="WL".

The argument priorPar determines either the prior mean and standard deviation of the normal prior distribution (if priorDist="norm"), or the range for defining the prior uniform distribution (if priorDist="unif"). This argument is ignored if priorDist="Jeffreys".

The eapPar argument sets the range and the number of quadrature points for numerical integration in the EAP process. By default, it takes the vector value (-4, 4, 20), that is, 20 quadrature points on the range [-4; 4]. See eapEst for further details.

The range argument permits to limit the interval of investigation for the ML, BM and WL ability estimates (in particular, to avoid infinite ability estimates). The default range is [-4, 4].

Value

The estimated ability level.

Author(s)

David Magis
Post-doctoral researcher, FNRS (Fonds National de la Recherche Scientifique)
Department of Mathematics, University of Liege, Belgium
<david.magis@ulg.ac.be>

References

Birnbaum, A. (1969). Statistical theory for logistic mental test models with a prior distribution of ability. *Journal of Mathematical Psychology, 6*, 258-276.

Bock, R. D., and Mislevy, R. J. (1982). Adaptive EAP estimation of ability in a microcomputer environment. *Applied Psychological Measurement, 6*, 431-444.

Haley, D.C. (1952). Estimation of the dosage mortality relationship when the dose is subject to error. Technical report no 15. Palo Alto, CA: Applied Mathematics and Statistics Laboratory, Stanford University.

Jeffreys, H. (1939). *Theory of probability*. Oxford, UK: Oxford University Press.

Jeffreys, H. (1946). An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 186*, 453-461.

Lord, F.M. (1980). *Applications of item response theory to practical testing problems*. Hillsdale, NJ: Lawrence Erlbaum.

Warm, T.A. (1989). Weighted likelihood estimation of ability in item response models. *Psychometrika, 54*, 427-450.

See Also

eapEst, semTheta

Examples

```
## Not run:

 # Loading the 'tcals' parameters
 data(tcals)
 tcals <- as.matrix(tcals)

 # Creation of a response pattern (tcals item parameters,
 # true ability level 0)
 set.seed(1)
 x <- rbinom(85, 1, Pi(0, tcals)$Pi)

 # ML estimation
 thetaEst(tcals, x, method="ML")

 # BM estimation, standard normal prior distribution
 thetaEst(tcals, x)

 # BM estimation, uniform prior distribution upon range [-2,2]
 thetaEst(tcals, x, method="BM", priorDist="unif", priorPar=c(-2,2))

 # BM estimation, Jeffreys' prior distribution
 thetaEst(tcals, x, method="BM", priorDist="Jeffreys")

 # EAP estimation, standard normal prior distribution
 thetaEst(tcals, x, method="EAP")

 # EAP estimation, uniform prior distribution upon range [-2,2]
 thetaEst(tcals, x, method="EAP", priorDist="unif", priorPar=c(-2,2))

 # EAP estimation, Jeffreys' prior distribution
 thetaEst(tcals, x, method="EAP", priorDist="Jeffreys")
```

```
 # WL estimation
 thetaEst(tcals, x, method="WL")

## End(Not run)
```

Index