

Package ‘HEMDAG’

August 11, 2017

Title Hierarchical Ensemble Methods for Directed Acyclic Graphs

Version 1.0.0

Author Marco Notaro [aut, cre] and Giorgio Valentini [aut]
(AnacletoLab, Dipartimento di Informatica, Universita' degli Studi di Milano)

Maintainer Marco Notaro <marco.notaro@unimi.it>

Description An implementation of Hierarchical Ensemble Methods for DAGs: 'HTD-DAG' (Hierarchical Top Down) and 'TPR-DAG' (True Path Rule). 'HEMDAG' can be used to enhance the predictions of virtually any flat learning method, by taking into account the hierarchical nature of the classes of a bio-ontology. 'HEMDAG' is specifically designed for exploiting the hierarchical relationships of DAG-structured taxonomies, such as the Human Phenotype Ontology (HPO) or the Gene Ontology (GO), but it can be also safely applied to tree-structured taxonomies (as FunCat), since trees are DAGs. 'HEMDAG' scale nicely both in terms of the complexity of the taxonomy and in the cardinality of the examples. (Marco Notaro, Max Schubach, Peter N. Robinson and Giorgio Valentini, Prediction of Human Phenotype Ontology terms by means of Hierarchical Ensemble methods, BMC Bioinformatics 2017).

Depends R (>= 3.4.1)

License GPL (>= 3)

Encoding UTF-8

LazyData true

Repository CRAN

NeedsCompilation no

Imports graph,
RBGL,
PerfMeas,
precrec,
preprocessCore,
methods

Suggests Rgraphviz

RoxygenNote 6.0.1

R topics documented:

HEMDAG-package	3
ancestors	3
AUPRC.single.class	4
AUPRC.single.over.classes	5
AUROC.single.class	6
AUROC.single.over.classes	7
check.annotation.matrix.integrity	8
check.DAG.integrity	9
children	9
compute.flipped.graph	10
constraints.matrix	11
descendants	11
distances.from.leaves	12
do.edges.from.HPO.obo	13
Do.FLAT.scores.normalization	13
Do.full.annotation.matrix	15
Do.HTD	16
Do.HTD.holdout	18
do.subgraph	20
do.submatrix	21
Do.tpr.threshold.free	21
Do.tpr.threshold.free.holdout	23
do.unstratified.cv.data	25
example.datasets	26
find.best.f	27
find.leaves	28
full.annotation.matrix	29
graph.levels	30
hierarchical.checkers	31
htd	32
Multilabel.F.measure	33
normalize.max	34
parents	35
read.graph	36
read.undirected.graph	36
root.node	37
specific.annotation.list	38
specific.annotation.matrix	38
stratified.cross.validation	39
TPR-DAG	40
TPR.DAG.CV	42
TPR.DAG.HOLDOUT	45
transitive.closure.annotations	47
weighted.adjacency.matrix	48
write.graph	49

HEMDAG-package	<i>HEMDAG: Hierarchical Ensemble Methods for Directed Acyclic Graphs</i>
----------------	--

Description

Implementation of Hierarchical Ensemble Methods for DAGs: HTD-DAG (Hierarchical Top Down) and TPR-DAG (True Path Rule).

Details

HEMDAG can be used to enhance the predictions of virtually any off-the-shelf flat learning method, merely by taking into account the hierarchical nature of the classes of a bio-ontology. HEMDAG is specifically designed for exploiting the hierarchical relationships of DAG-structured taxonomies, such as the Human Phenotype Ontology (HPO) or the Gene Ontology (GO), but it can be also safely applied to tree-structured taxonomies (as FunCat), since obviously trees are DAGs. Finally HEMDAG scale nicely both in terms of the complexity of the taxonomy and in the cardinality of the examples.

Author(s)

Marco Notaro and Giorgio Valentini

[AnacletoLab](#)

DI, Dipartimento di Informatica

Universita' degli Studi di Milano

Maintainer: *Marco Notaro* <marco.notaro@unimi.it>

References

Marco Notaro, Max Schubach, Peter N. Robinson and Giorgio Valentini, *Prediction of Human Phenotype Ontology terms by means of Hierarchical Ensemble methods*, BMC Bioinformatics 2017

ancestors	<i>Build ancestors</i>
-----------	------------------------

Description

Compute the ancestors for each node of a graph

Usage

```
build.ancestors(g)
```

```
build.ancestors.per.level(g, levels)
```

```
build.ancestors.bottom.up(g, levels)
```

Arguments

<code>g</code>	a graph of class <code>graphNEL</code> . It represents the hierarchy of the classes
<code>levels</code>	a list of character vectors. Each component represents a graph level and the elements of any component correspond to nodes. The level 0 coincides with the root node.

Value

`build.ancestos` returns a named list of vectors. Each component corresponds to a node x of the graph and its vector is the set of its ancestors including also x .

`build.ancestors.per.level` returns a named list of vectors. Each component corresponds to a node x of the graph and its vector is the set of its ancestors including also x . The nodes are ordered from root (included) to leaves.

`build.ancestors.bottom.up` a named list of vectors. Each component corresponds to a node x of the graph and its vector is the set of its ancestors including also x . The nodes are ordered from leaves to root (included).

See Also

[graph.levels](#)

Examples

```
data(graph);
root <- root.node(g);
anc <- build.ancestors(g);
lev <- graph.levels(g, root=root);
anc.tod <- build.ancestors.per.level(g, lev);
anc.bup <- build.ancestors.bottom.up(g, lev);
```

AUPRC.single.class *AUPRC single class*

Description

High-level function to compute the Area under the Precision Recall Curve (AUPRC) just for a single class through **precrec** package

Usage

```
AUPRC.single.class(target, pred)
```

Arguments

<code>target</code>	vector of the true labels (0 negative, 1 positive examples)
<code>pred</code>	numeric vector of the values of the predicted labels (scores)

Value

a numeric value corresponding to the AUPRC for the considered class

See Also

[AUPRC.single.over.classes](#)

Examples

```
data(labels);
data(scores);
data(graph);
root <- root.node(g);
L <- L[, -which(colnames(L)==root)];
S <- S[, -which(colnames(S)==root)];
PRC <- AUPRC.single.class(L[,3],S[,3]);
```

AUPRC.single.over.classes

AUPRC over classes

Description

High-level function to compute the Area under the Precision Recall Curve (AUPRC) across a set of classes through **precrec** package

Usage

```
AUPRC.single.over.classes(target, pred)
```

Arguments

target	matrix with the target multilabels: rows correspond to examples and columns to classes. $target[i, j] = 1$ if example i belongs to class j , $target[i, j] = 0$ otherwise.
pred	a numeric matrix with predicted values (scores): rows correspond to examples and columns to classes.

Value

a list with two elements:

1. average: the average AUPRC across classes;
2. per.class: a named vector with AUPRC for each class. Names correspond to classes

See Also

[AUPRC.single.class](#)

Examples

```
data(labels);
data(scores);
data(graph);
root <- root.node(g);
L <- L[,-which(colnames(L)==root)];
S <- S[,-which(colnames(S)==root)];
PRC <- AUPRC.single.over.classes(L,S);
```

AUROC.single.class *AUROC single class*

Description

High-level function to compute the Area under the ROC Curve (AUPRC) just for a single class through **precrec** package

Usage

```
AUROC.single.class(target, pred)
```

Arguments

target	vector of the true labels (0 negative, 1 positive examples)
pred	numeric vector of the values of the predicted labels (scores)

Value

a numeric value corresponding to the AUROC for the considered class

See Also

[AUROC.single.over.classes](#)

Examples

```
data(labels);
data(scores);
data(graph);
root <- root.node(g);
L <- L[,-which(colnames(L)==root)];
S <- S[,-which(colnames(S)==root)];
AUC <- AUROC.single.class(L[,3],S[,3]);
```

`AUROC.single.over.classes`*AUROC over classes*

Description

High-level function to compute the Area under the ROC Curve (AUROC) for a set of classes through **precrec** package

Usage

```
AUROC.single.over.classes(target, pred)
```

Arguments

target	matrix with the target multilabels: rows correspond to examples and columns to classes. $target[i, j] = 1$ if example i belongs to class j , $target[i, j] = 0$ otherwise.
pred	a numeric matrix with predicted values (scores): rows correspond to examples and columns to classes.

Value

a list with two elements:

1. average: the average AUROC across classes;
2. per.class: a named vector with AUROC for each class. Names correspond to classes

See Also

[AUROC.single.class](#)

Examples

```
data(labels);
data(scores);
data(graph);
root <- root.node(g);
L <- L[, -which(colnames(L)==root)];
S <- S[, -which(colnames(S)==root)];
AUC <- AUROC.single.over.classes(L,S);
```

`check.annotation.matrix.integrity`*Annotation matrix checker*

Description

This function assess the integrity of an annotation table in which a transitive closure of annotations was performed

Usage

```
check.annotation.matrix.integrity(anc, ann.spec, hpo.ann)
```

Arguments

<code>anc</code>	list of the ancestors of the ontology.
<code>ann.spec</code>	the annotation matrix of the most specific annotations (0/1): rows are genes and columns are terms.
<code>hpo.ann</code>	the full annotations matrix (0/1), that is the matrix in which the transitive closure of the annotation was performed. Rows are examples and columns are classes.

Value

If the transitive closure of the annotations is well performed "OK" is returned, otherwise a message error is printed on the stdout

See Also

[build.ancestors](#), [transitive.closure.annotations](#), [full.annotation.matrix](#)

Examples

```
data(graph);  
data(labels);  
anc <- build.ancestors(g);  
tca <- transitive.closure.annotations(L, anc);  
check.annotation.matrix.integrity(anc, L, tca);
```

check.DAG.integrity *DAG checker*

Description

This function assess the integrity of a DAG

Usage

```
check.DAG.integrity(g, root = "00")
```

Arguments

g	a graph of class graphNEL. It represents the hierarchy of the classes.
root	name of the class that is on the top-level of the hierarchy (def:"00")

Value

If there are nodes not accessible from the root "OK" is printed, otherwise a message error and the list of the not accessible nodes is printed on the stdout

Examples

```
data(graph);
root <- root.node(g);
check.DAG.integrity(g, root=root);
```

children *Build children*

Description

Compute the children for each node of a graph

Usage

```
build.children(g)

get.children.top.down(g, levels)

get.children.bottom.up(g, levels)
```

Arguments

g	a graph of class graphNEL. It represents the hierarchy of the classes
levels	a list of character vectors. Each component represents a graph level and the elements of any component correspond to nodes. The level 0 coincides with the root node.

Value

build.children returns a named list of vectors. Each component corresponds to a node x of the graph and its vector is the set of its children

get.children.top.down returns a named list of character vectors. Each component corresponds to a node x of the graph (i.e. parent node) and its vector is the set of its children. The nodes are ordered from root (included) to leaves.

get.children.bottom.up returns a named list of character vectors. Each component corresponds to a node x of the graph (i.e. parent node) and its vector is the set of its children. The nodes are ordered from leaves (included) to root.

See Also

[graph.levels](#)

Examples

```
data(graph);
root <- root.node(g);
children <- build.children(g);
lev <- graph.levels(g, root=root);
children.tod <- get.children.top.down(g,lev);
children.bup <- get.children.bottom.up(g,lev);
```

compute.flipped.graph *Flip Graph*

Description

Compute a directed graph with edges in the opposite direction

Usage

```
compute.flipped.graph(g)
```

Arguments

`g` a graphNEL directed graph

Value

a graph (as an object of class graphNEL) with edges in the opposite direction w.r.t. `g`

Examples

```
data(graph);
g.flipped <- compute.flipped.graph(g);
```

constraints.matrix	<i>Constraints matrix</i>
--------------------	---------------------------

Description

This function returns a matrix with two columns and as many rows as there are edges. The entries of the first columns are the index of the node the edge comes from (i.e. children nodes), the entries of the second columns indicate the index of node the edge is to (i.e. parents nodes). Referring to a DAG this matrix defines a partial order.

Usage

```
constraints.matrix(g)
```

Arguments

`g` a graph of class graphNEL. Represent the hierarchy of the class

Value

a constraints matrix w.r.t the graph `g`

Examples

```
data(graph);  
m <- constraints.matrix(g);
```

descendants	<i>Build descendants</i>
-------------	--------------------------

Description

Compute the descendants for each node of a graph

Usage

```
build.descendants(g)
```

```
build.descendants.per.level(g, levels)
```

```
build.descendants.bottom.up(g, levels)
```

Arguments

`g` a graph of class graphNEL. It represents the hierarchy of the classes

`levels` a list of character vectors. Each component represents a graph level and the elements of any component correspond to nodes. The level 0 coincides with the root node.

Value

`build.descendants` returns a named list of vectors. Each component corresponds to a node x of the graph, and its vector is the set of its descendants including also x .

`build.descendants.per.level` returns a named list of vectors. Each component corresponds to a node x of the graph and its vector is the set of its descendants including also x . The nodes are ordered from root (included) to leaves.

`build.descendants.bottom.up` returns a named list of vectors. Each component corresponds to a node x of the graph and its vector is the set of its descendants including also x . The nodes are ordered from leaves to root (included).

See Also

[graph.levels](#)

Examples

```
data(graph);
root <- root.node(g);
desc <- build.descendants(g);
lev <- graph.levels(g, root=root);
desc.tod <- build.descendants.per.level(g, lev);
desc.bup <- build.descendants.bottom.up(g, lev);
```

`distances.from.leaves` *Distances from leaves*

Description

This function returns the minimum distance of each node from one of the leaves of the graph

Usage

```
distances.from.leaves(g)
```

Arguments

`g` a graph of class `graphNEL`. It represents the hierarchy of the classes.

Value

a named vector. The names are the names of the nodes of the graph `g`, and their values represent the distance from the leaves. A value equal to 0 is assigned to the leaves, 1 to nodes with distance 1 from a leaf and so on

Examples

```
data(graph);
dist.leaves <- distances.from.leaves(g);
```

do.edges.from.HPO.obo *Parse an HPO OBO file*

Description

Read an HPO OBO file (**HPO**) and write the edges of the DAG on a plain text file. The format of the file is a sequence of rows and each row corresponds to an edge represented through a pair of vertices separated by blanks

Usage

```
do.edges.from.HPO.obo(file = "hp.obo", output.file = "edge.file")
```

Arguments

file	an HPO OBO file
output.file	name of the file of the edges to be written

Value

a text file representing the edges in the format: source destination (i.e. one row for each edge)

Examples

```
## Not run:  
hpobo <- "http://purl.obolibrary.org/obo/hp.obo";  
do.edges.from.HPO.obo(file=hpobo, output.file="hp.edge");  
## End(Not run)
```

Do.FLAT.scores.normalization

Flat scores normalization

Description

High level functions to normalize a flat scores matrix w.r.t. max normalization (MaxNorm) or quantile normalization (Qnorm)

Usage

```
Do.FLAT.scores.normalization(norm.type = "MaxNorm", flat.file = flat.file,  
dag.file = dag.file, flat.dir = flat.dir, dag.dir = dag.dir,  
flat.norm.dir = flat.norm.dir)
```

Arguments

<code>norm.type</code>	can assume two character values: <ul style="list-style-type: none"> • <code>MaxNorm</code>: each score is divided w.r.t. the max of each class; • <code>Qnorm</code>: a quantile normalization is applied. Library <code>preprocessCore</code> is used.
<code>flat.file</code>	name of the flat scores matrix (without rda extension)
<code>dag.file</code>	name of the graph that represents the hierarchy of the classes
<code>flat.dir</code>	relative path to folder where flat normalized scores matrix is stored
<code>dag.dir</code>	relative path to folder where graph is stored
<code>flat.norm.dir</code>	the directory where the normalized flat scores matrix must be stored

Details

To apply the quantile normalization the **preprocessCore** library is used.

Value

the matrix of the scores flat normalized w.r.t. `MaxNorm` or `Qnorm`

Examples

```

data(scores);
data(graph);
if (!dir.exists("data")){
  dir.create("data");
}
if (!dir.exists("results")){
  dir.create("results");
}
save(S,file="data/scores.rda");
save(g,file="data/graph.rda");
flat.dir <- dag.dir <- "data/";
flat.norm.dir <- "results/";
flat.file <- "scores";
dag.file <- "graph";
norm.types <- c("MaxNorm","Qnorm");
for(norm.type in norm.types){
  Do.FLAT.scores.normalization(norm.type=norm.type, flat.file=flat.file,
  dag.file=dag.file, flat.dir=flat.dir, dag.dir=dag.dir,
  flat.norm.dir=flat.norm.dir);
}

```

`Do.full.annotation.matrix`*Do full annotations matrix*

Description

High-level function to obtain a full annotation matrix, that is a matrix in which the transitive closure of annotations was performed, respect to a given weighted adjacency matrix

Usage

```
Do.full.annotation.matrix(anc.file.name = anc.file.name, anc.dir = anc.dir,  
  net.file = net.file, net.dir = net.dir, ann.file.name = ann.file.name,  
  ann.dir = ann.dir, output.name = output.name, output.dir = output.dir)
```

Arguments

<code>anc.file.name</code>	name of the file containing the list for each node the list of all its ancestor (without rda extension)
<code>anc.dir</code>	relative path to directory where the ancestor file is stored
<code>net.file</code>	name of the file containing the weighted adjacency matrix of the graph (without rda extension)
<code>net.dir</code>	relative path to directory where the weighted adjacency matrix is stored
<code>ann.file.name</code>	name of the file containing the matrix of the most specific annotations (without rda extension)
<code>ann.dir</code>	relative path to directory where the matrix of the most specific annotation is stored
<code>output.name</code>	name of the output file without rda extension (without rda extension)
<code>output.dir</code>	relative path to directory where the output file must be stored

Value

a full annotation matrix T , that is a matrix in which the transitive closure of annotations was performed. Rows correspond to genes of the input weighted adjacency matrix and columns to terms. $T[i, j] = 1$ means that gene i is annotated for the term j , $T[i, j] = 0$ means that gene i is not annotated for the term j .

See Also

[full.annotation.matrix](#)

Examples

```

data(graph);
data(labels);
data(wadj);
if (!dir.exists("data")){
  dir.create("data");
}
if (!dir.exists("results")){
  dir.create("results");
}
anc <- build.ancestors(g);
save(anc,file="data/ancestors.rda");
save(g,file="data/graph.rda");
save(L,file="data/labels.rda");
save(W,file="data/wadj.rda");
anc.dir <- net.dir <- ann.dir <- "data/";
output.dir <- "results/";
anc.file.name <- "ancestors";
net.file <- "wadj";
ann.file.name <- "labels";
output.name <- "full.ann.matrix";
Do.full.annotation.matrix(anc.file.name=anc.file.name, anc.dir=anc.dir, net.file=net.file,
net.dir=net.dir, ann.file.name=ann.file.name, ann.dir=ann.dir, output.name=output.name,
output.dir=output.dir);

```

Do.HTD

*HTD-DAG vanilla***Description**

High level function to compute hierarchical correction according to HTD-DAG algorithm

Usage

```

Do.HTD(norm = TRUE, norm.type = "NONE", flat.file = flat.file,
ann.file = ann.file, dag.file = dag.file, flat.dir = flat.dir,
ann.dir = ann.dir, dag.dir = dag.dir, flat.norm.dir = NULL,
n.round = 3, f.criterion = "F", hierScore.dir = hierScore.dir,
perf.dir = perf.dir)

```

Arguments

norm	boolean value: <ul style="list-style-type: none"> • TRUE (def.): the flat scores matrix has been already normalized in according to a normalization method; • FALSE: the flat scores matrix has not been normalized yet. See the parameter norm.type for which normalization can be applied.
norm.type	three values:

	<ol style="list-style-type: none"> 1. NONE (def.): set <code>norm.type</code> to NONE if and only if the parameter <code>norm</code> is set to TRUE; 2. MaxNorm: each score is divided w.r.t. the max of each class; 3. Qnorm: quantile normalization. preprocessCore package is used.
<code>flat.file</code>	name of the file containing the flat scores matrix to be normalized or already normalized (without rda extension)
<code>ann.file</code>	name of the file containing the the label matrix of the examples (without rda extension)
<code>dag.file</code>	name of the file containing the graph that represents the hierarchy of the classes (without rda extension)
<code>flat.dir</code>	relative path where flat scores matrix is stored
<code>ann.dir</code>	relative path where annotation matrix is stored
<code>dag.dir</code>	relative path where graph is stored
<code>flat.norm.dir</code>	relative path where flat normalized scores matrix must be stored. Use this parameter if and only if <code>norm</code> is set to FALSE, otherwise set <code>flat.norm.dir</code> to NULL (def.)
<code>n.round</code>	number of rounding digits to be applied to the hierarchical scores matrix (def. 3). It is used for choosing the best threshold on the basis of the best F-measure
<code>f.criterion</code>	character. Type of F-measure to be used to select the best F-measure. Two possibilities: <ol style="list-style-type: none"> 1. F (def.): corresponds to the harmonic mean between the average precision and recall 2. avF: corresponds to the per-example F-score averaged across all the examples
<code>hierScore.dir</code>	relative path where the hierarchical scores matrix must be stored
<code>perf.dir</code>	relative path where the term-centric and protein-centric measures must be stored

Value

Five rda files stored in the respective output directories:

1. `hierarchical scores matrix`: a matrix with examples on rows and classes on columns representing the computed hierarchical scores for each class and example considered. It stored in `hierScore.dir` directory.
2. `PCM (Protein Centric Measures) average and per-example`: compute F-score measure by `find.best.f` function. It stored in `perf.dir` directory.
3. `PRC (area under Precision-Recall Curve) average and per.class`: compute PRC by **precrec** package. It stored in `perf.dir` directory.
4. `AUC (Area Under ROC Curve) average and per-class`: compute AUC by **precrec** package. It stored in `perf.dir` directory.
5. `PxR (Precision at fixed Recall levels) average and per classes`: compute PxR by **PerfMeas** package. It stored in `perf.dir` directory.

Examples

```

data(graph);
data(scores);
data(labels);
if (!dir.exists("data")){
  dir.create("data");
}
if (!dir.exists("results")){
  dir.create("results");
}
save(g,file="data/graph.rda");
save(L,file="data/labels.rda");
save(S,file="data/scores.rda");
dag.dir <- flat.dir <- flat.norm.dir <- ann.dir <- "data/";
hierScore.dir <- perf.dir <- "results/";
dag.file <- "graph";
flat.file <- "scores";
ann.file <- "labels";
Do.HTD(norm=FALSE, norm.type= "MaxNorm", flat.file=flat.file, ann.file=ann.file,
dag.file=dag.file, flat.dir=flat.dir, ann.dir=ann.dir, dag.dir=dag.dir,
flat.norm.dir=flat.norm.dir, n.round=3, f.criterion ="F", hierScore.dir=hierScore.dir,
perf.dir=perf.dir);

```

Do.HTD.holdout

*HTD-DAG holdout***Description**

High level function to correct the scores with a hierarchy according to HTD-DAG algorithm applying a classical holdout procedure

Usage

```

Do.HTD.holdout(norm = TRUE, norm.type = "NONE", flat.file = flat.file,
  ann.file = ann.file, dag.file = dag.file, ind.test.set = ind.test.set,
  ind.dir = ind.dir, flat.dir = flat.dir, ann.dir = ann.dir,
  dag.dir = dag.dir, flat.norm.dir = NULL, n.round = 3,
  f.criterion = "F", hierScore.dir = "hierScore.dir/",
  perf.dir = "perf.dir/")

```

Arguments

norm	boolean value: <ul style="list-style-type: none"> • TRUE (def.): the flat scores matrix has been already normalized in according to a normalization method; • FALSE: the flat scores matrix has not been normalized yet. See the parameter norm for which normalization can be applied.
norm.type	three values:

	<ol style="list-style-type: none"> 1. NONE (def.): set <code>norm.type</code> to NONE if and only if the parameter <code>norm</code> is set to TRUE; 2. MaxNorm: each score is divided w.r.t. the max of each class; 3. Qnorm: quantile normalization. preprocessCore package is used.
<code>flat.file</code>	name of the file containing the flat scores matrix to be normalized or already normalized (without rda extension)
<code>ann.file</code>	name of the file containing the the label matrix of the examples (without rda extension)
<code>dag.file</code>	name of the file containing the graph that represents the hierarchy of the classes (without rda extension)
<code>ind.test.set</code>	name of the file containing a vector of integer numbers corresponding to the indices of the elements (rows) of scores matrix to be used in the test set
<code>ind.dir</code>	relative path to folder where <code>ind.test.set</code> is stored
<code>flat.dir</code>	relative path where flat scores matrix is stored
<code>ann.dir</code>	relative path where annotation matrix is stored
<code>dag.dir</code>	relative path where graph is stored
<code>flat.norm.dir</code>	relative path where flat normalized scores matrix must be stored. Use this parameter if and only if <code>norm</code> is set to FALSE, otherwise set <code>flat.norm.dir</code> to NULL (def.)
<code>n.round</code>	number of rounding digits to be applied to the hierarchical scores matrix (def. 3). It is used for choosing the best threshold on the basis of the best F-measure
<code>f.criterion</code>	character. Type of F-measure to be used to select the best F-measure. Two possibilities: <ol style="list-style-type: none"> 1. F (def.): corresponds to the harmonic mean between the average precision and recall 2. avF: corresponds to the per-example F-score averaged across all the examples
<code>hierScore.dir</code>	relative path where the hierarchical scores matrix must be stored
<code>perf.dir</code>	relative path where the term-centric and protein-centric measures must be stored

Value

Five rda files stored in the respective output directories:

1. `hierarchical scores matrix`: a matrix with examples on rows and classes on columns representing the computed hierarchical scores for each class and example considered. It stored in `hierScore.dir` directory.
2. PCM (Protein Centric Measures) average and per-example: compute F-score measure by `find.best.f` function. It stored in `perf.dir` directory.
3. PRC (area under Precision-Recall Curve) average and per.class: compute PRC by **precrec** package. It stored in `perf.dir` directory.
4. AUC (Area Under ROC Curve) average and per-class: compute AUC by **precrec** package. It stored in `perf.dir` directory.
5. PxR (Precision at fixed Recall levels) average and per classes: compute PxR by **PerfMeas** package. It stored in `perf.dir` directory.

Examples

```

data(graph);
data(scores);
data(labels);
data(test.index);
if (!dir.exists("data")){
  dir.create("data");
}
if (!dir.exists("results")){
  dir.create("results");
}
save(g,file="data/graph.rda");
save(L,file="data/labels.rda");
save(S,file="data/scores.rda");
save(test.index, file="data/test.index.rda");
ind.dir <- dag.dir <- flat.dir <- flat.norm.dir <- ann.dir <- "data/";
hierScore.dir <- perf.dir <- "results/";
ind.test.set <- "test.index";
dag.file <- "graph";
flat.file <- "scores";
ann.file <- "labels";
Do.HTD.holdout(norm=FALSE, norm.type= "MaxNorm", flat.file=flat.file, ann.file=ann.file,
dag.file=dag.file, ind.test.set=ind.test.set, ind.dir=ind.dir, flat.dir=flat.dir,
ann.dir=ann.dir, dag.dir=dag.dir, flat.norm.dir=flat.norm.dir, n.round=3, f.criterion ="F",
hierScore.dir=hierScore.dir, perf.dir=perf.dir);

```

do.subgraph

*Build subgraph***Description**

This function returns a subgraph with only the supplied nodes and any edges between them

Usage

```
do.subgraph(nd, g, edgemode = "directed")
```

Arguments

nd	a vector with the nodes for which the subgraph must be built
g	a graph of class graphNEL. It represents the hierarchy of the classes.
edgemode	can be "directed" or "undirected"

Value

a subgraph with only the supplied nodes

Examples

```
data(graph);
anc <- build.ancestors(g);
nd <- anc[["HP:0001371"]];
subg <- do.subgraph(nd, g, edgemode="directed");
```

do.submatrix	<i>Build submatrix</i>
--------------	------------------------

Description

Terms having less than n annotations are pruned. Terms having exactly n annotations are discarded as well.

Usage

```
do.submatrix(hpo.ann, n)
```

Arguments

hpo.ann	the annotations matrix (0/1). Rows are examples and columns are classes
n	integer number of annotations to be pruned

Value

Matrix of annotations having only those terms with more than n annotations

Examples

```
data(labels);
subm <- do.submatrix(L,5);
```

Do.tpr.threshold.free	<i>TPR-DAG vanilla</i>
-----------------------	------------------------

Description

High level function to compute hierarchical correction according to TPR-DAG vanilla algorithm

Usage

```
Do.tpr.threshold.free(norm = TRUE, norm.type = "NONE",
  flat.file = flat.file, ann.file = ann.file, dag.file = dag.file,
  flat.dir = flat.dir, ann.dir = ann.dir, dag.dir = dag.dir,
  flat.norm.dir = NULL, n.round = 3, f.criterion = "F",
  hierScore.dir = hierScore.dir, perf.dir = perf.dir)
```

Arguments

norm	boolean value: <ul style="list-style-type: none"> • TRUE (def.): the flat scores matrix has been already normalized in according to a normalization method; • FALSE: the flat scores matrix has not been normalized yet. See the parameter norm.type for which normalization can be applied.
norm.type	three values: <ol style="list-style-type: none"> 1. NONE (def.): set norm.type to NONE if and only if the parameter norm is set to TRUE; 2. MaxNorm: each score is divided w.r.t. the max of each class; 3. Qnorm: quantile normalization. preprocessCore package is used.
flat.file	name of the file containing the flat scores matrix to be normalized or already normalized (without rda extension)
ann.file	name of the file containing the the label matrix of the examples (without rda extension)
dag.file	name of the file containing the graph that represents the hierarchy of the classes (without rda extension)
flat.dir	relative path where flat scores matrix is stored
ann.dir	relative path where annotation matrix is stored
dag.dir	relative path where graph is stored
flat.norm.dir	relative path where flat normalized scores matrix must be stored. Use this parameter if and only if norm is set to FALSE, otherwise set flat.norm.dir to NULL (def.)
n.round	number of rounding digits to be applied to the hierarchical scores matrix (def. 3). It is used for choosing the best threshold on the basis of the best F-measure
f.criterion	character. Type of F-measure to be used to select the best F-measure. Two possibilities: <ol style="list-style-type: none"> 1. F (def.): corresponds to the harmonic mean between the average precision and recall 2. avF: corresponds to the per-example F-score averaged across all the examples
hierScore.dir	relative path where the hierarchical scores matrix must be stored
perf.dir	relative path where the term-centric and protein-centric measures must be stored

Value

Five rda files stored in the respective output directories:

1. hierarchical scores matrix: a matrix with examples on rows and classes on columns representing the computed hierarchical scores for each class and example considered. It stored in hierScore.dir directory.
2. PCM (Protein Centric Measures) average and per-example: compute F-score measure by find.best.f function. It stored in perf.dir directory.

3. PRC (area under Precision-Recall Curve) average and per.class: compute PRC by **precrec** package. It stored in perf.dir directory.
4. AUC (Area Under ROC Curve) average and per.class: compute AUC by **precrec** package. It stored in perf.dir directory.
5. PxR (Precision at fixed Recall levels) average and per classes: compute PxR by **PerfMeas** package. It stored in perf.dir directory.

Examples

```

data(graph);
data(scores);
data(labels);
if (!dir.exists("data")){
  dir.create("data");
}
if (!dir.exists("results")){
  dir.create("results");
}
save(g,file="data/graph.rda");
save(L,file="data/labels.rda");
save(S,file="data/scores.rda");
dag.dir <- flat.dir <- flat.norm.dir <- ann.dir <- "data/";
hierScore.dir <- perf.dir <- "results/";
dag.file <- "graph";
flat.file <- "scores";
ann.file <- "labels";
Do.tpr.threshold.free(norm=FALSE, norm.type= "MaxNorm", flat.file=flat.file,
ann.file=ann.file, dag.file=dag.file, flat.dir=flat.dir, ann.dir=ann.dir,
dag.dir=dag.dir, flat.norm.dir=flat.norm.dir, n.round=3, f.criterion = "F",
hierScore.dir=hierScore.dir, perf.dir=perf.dir);

```

Do.tpr.threshold.free.holdout

TPR-DAG vanilla holdout

Description

High level function to correct the scores with a hierarchy according to TPR-DAG vanilla algorithm applying a classical holdout procedure

Usage

```

Do.tpr.threshold.free.holdout(norm = TRUE, norm.type = "NONE",
  flat.file = flat.file, ann.file = ann.file, dag.file = dag.file,
  flat.dir = flat.dir, ann.dir = ann.dir, dag.dir = dag.dir,
  flat.norm.dir = NULL, ind.test.set = ind.test.set, ind.dir = ind.dir,
  n.round = 3, f.criterion = "F", hierScore.dir = hierScore.dir,
  perf.dir = perf.dir)

```

Arguments

norm	boolean value: <ul style="list-style-type: none"> • TRUE (def.): the flat scores matrix has been already normalized in according to a normalization method; • FALSE: the flat scores matrix has not been normalized yet. See the parameter norm for which normalization can be applied.
norm.type	three values: <ol style="list-style-type: none"> 1. NONE (def.): set norm.type to NONE if and only if the parameter norm is set to TRUE; 2. MaxNorm: each score is divided w.r.t. the max of each class; 3. Qnorm: quantile normalization. preprocessCore package is used.
flat.file	name of the file containing the flat scores matrix to be normalized or already normalized (without rda extension)
ann.file	name of the file containing the the label matrix of the examples (without rda extension)
dag.file	name of the file containing the graph that represents the hierarchy of the classes (without rda extension)
flat.dir	relative path where flat scores matrix is stored
ann.dir	relative path where annotation matrix is stored
dag.dir	relative path where graph is stored
flat.norm.dir	relative path where flat normalized scores matrix must be stored. Use this parameter if and only if norm is set to FALSE, otherwise set flat.norm.dir to NULL (def.)
ind.test.set	name of the file containing a vector of integer numbers corresponding to the indices of the elements (rows) of scores matrix to be used in the test set
ind.dir	relative path to folder where ind.test.set is stored
n.round	number of rounding digits to be applied to the hierarchical scores matrix (def. 3). It is used for choosing the best threshold on the basis of the best F-measure
f.criterion	character. Type of F-measure to be used to select the best F-measure. Two possibilities: <ol style="list-style-type: none"> 1. F (def.): corresponds to the harmonic mean between the average precision and recall 2. avF: corresponds to the per-example F-score averaged across all the examples
hierScore.dir	relative path where the hierarchical scores matrix must be stored
perf.dir	relative path where the term-centric and protein-centric measures must be stored

Value

Five rda files stored in the respective output directories:

1. hierarchical scores matrix: a matrix with examples on rows and classes on columns representing the computed hierarchical scores for each class and example considered. It stored in hierScore.dir directory.

2. PCM (Protein Centric Measures) average and per-example: compute F-score measure by `find.best.f` function. It stored in `perf.dir` directory.
3. PRC (area under Precision-Recall Curve) average and per.class: compute PRC by **precrec** package. It stored in `perf.dir` directory.
4. AUC (Area Under ROC Curve) average and per-class: compute AUC by **precrec** package. It stored in `perf.dir` directory.
5. PxR (Precision at fixed Recall levels) average and per classes: compute PxR by **PerfMeas** package. It stored in `perf.dir` directory.

Examples

```

data(graph);
data(scores);
data(labels);
data(test.index);
if (!dir.exists("data")){
  dir.create("data");
}
if (!dir.exists("results")){
  dir.create("results");
}
save(g,file="data/graph.rda");
save(L,file="data/labels.rda");
save(S,file="data/scores.rda");
save(test.index, file="data/test.index.rda");
ind.dir <- dag.dir <- flat.dir <- flat.norm.dir <- ann.dir <- "data/";
hierScore.dir <- perf.dir <- "results/";
ind.test.set <- "test.index";
dag.file <- "graph";
flat.file <- "scores";
ann.file <- "labels";
Do.tpr.threshold.free.holdout(norm=FALSE, norm.type= "MaxNorm", flat.file=flat.file,
ann.file=ann.file, dag.file=dag.file, ind.test.set=ind.test.set, ind.dir=ind.dir,
flat.dir=flat.dir, ann.dir=ann.dir, dag.dir=dag.dir, flat.norm.dir=flat.norm.dir,
n.round=3, f.criterion ="F", hierScore.dir=hierScore.dir, perf.dir=perf.dir);

```

do.unstratified.cv.data

Unstratified cross-validation

Description

This function splits a dataset in k-fold in an unstratified way (that is a fold may not have an equal amount of positive and negative examples). This function is used to perform k-fold cross-validation experiments in a hierarchical correction contest where splitting dataset in a stratified way is not needed.

Usage

```
do.unstratified.cv.data(S, kk = 5, seed = NULL)
```

Arguments

S	matrix of the flat scores. It must be a named matrix, where rows are example (e.g. genes) and columns are classes/terms (e.g. HPO terms)
kk	number of folds in which to split the dataset (def. k=5)
seed	seed for the random generator. If NULL (def.) no initialization is performed

Value

a list with $k = kk$ components (folds). Each component of the list is a character vector contains the names of the examples.

Examples

```
data(scores);
```

```
example.datasets
```

```
Small real example datasets
```

Description

Collection of real sub-datasets used in the examples of the **HEMDAG** package

Usage

```
data(graph)
data(labels)
data(scores)
data(wadj)
data(test.index)
```

Details

The DAG g contained in `graph` data is an object of class `graphNEL`. The graph g has 23 nodes and 30 edges. The nodes and the edges of g represent the terms and the between-term relationships of the leaf term "HP:0100490" (*Camptodactyly of finger*) of the Human Phenotype Ontology (HPO).

The matrix contained in the `labels` data is a 100 X 23 matrix L , whose rows correspondes to genes (*Entrez GeneID*) and columns to HPO classes. $L[i, j] = 1$ means that the gene i belong to class j , $L[i, j] = 0$ means that the gene i does not belong to class j . The classes of the matrix L correspond to the nodes of the graph g .

The matrix contained in the `scores` data is a named 100 X 23 flat scores matrix S , representing the likelihood that a given gene belongs to a given class: higher the value higher the likelihood. The classes of the matrix S correspond to the nodes of the graph g .

The matrix contained in the `wadj` data is a named 100 X 100 symmetric weighted adjacency matrix W , whose rows and columns correspond to genes. The genes names (*Entrez GeneID*) of the adjacency matrix W correspond to the genes names of the flat scores matrix S and to genes names of the target multilabel matrix L .

The vector of integer numbers contained in the `test.index` refers to the index of the examples of the scores matrix to be used in the test set. It is useful only in holdout experiments.

Note

Some examples of full data sets for the prediction of HPO terms are available at the following [link](#). Note that the processing of the full datasets should be done similarly to the processing of the small data examples provided directly in this package. Please read the README clicking the link above to know more details about the available full datasets.

find.best.f

Best hierarchical F-score

Description

Function to select the best hierarchical F-score by choosing an appropriate threshold in the scores

Usage

```
find.best.f(target, pred, n.round = 3, f.criterion = "F", verbose = TRUE,
            b.per.example = FALSE)
```

Arguments

<code>target</code>	matrix with the target multilabels: rows correspond to examples and columns to classes. $target[i, j] = 1$ if example i belongs to class j , $target[i, j] = 0$ otherwise
<code>pred</code>	a numeric matrix with predicted values (scores): rows correspond to examples and columns to classes
<code>n.round</code>	number of rounding digits to be applied to <code>pred</code> (default=3)
<code>f.criterion</code>	character. Type of F-measure to be used to select the best F-score. There are two possibilities: <ol style="list-style-type: none"> 1. F (def.) corresponds to the harmonic mean between the average precision and recall; 2. avF corresponds to the per-example F-score averaged across all the examples.
<code>verbose</code>	boolean. If TRUE (def.) the number of iterations are printed on stdout
<code>b.per.example</code>	boolean. <ul style="list-style-type: none"> • TRUE: results are returned for each example; • FALSE: only the average results are returned

Details

All the examples having no positive annotations are discarded. The predicted scores matrix (pred) is rounded according to parameter `n.round` and all the values of `pred` are divided by `max(pred)`. Then all the thresholds corresponding to all the different values included in `pred` are attempted, and the threshold leading to the maximum F-measure is selected.

Value

Two different outputs respect to the input parameter `b.per.example`:

- `b.per.example==FALSE`: a list with a single element average. A named vector with 7 elements relative to the best result in terms of the F-measure: Precision (P), Recall (R), Specificity (S), F-measure (F), av.F-measure (av.F), Accuracy (A) and the best selected Threshold (T). F is the F-measure computed as the harmonic mean between the average precision and recall; av.F is the F-measure computed as the average across examples and T is the best selected threshold;
- `b.per.example==TRUE`: a list with two elements:
 1. average: a named vector with with 7 elements relative to the best result in terms of the F-measure: Precision (P), Recall (R), Specificity (S), F-measure (F), av.F-measure (av.F), Accuracy (A) and the best selected Threshold (T).
 2. per.example: a named matrix with the Precision (P), Recall (R), Specificity (S), Accuracy (A), F-measure (F), av.F-measure (av.F) and the best selected Threshold (T) for each example. Row names correspond to examples, column names correspond respectively to Precision (P), Recall (R), Specificity (S), Accuracy (A), F-measure (F), av.F-measure (av.F) and the best selected Threshold (T).

Examples

```
data(labels);
data(scores);
root <- root.node(g);
L <- L[,-which(colnames(L)==root)];
S <- S[,-which(colnames(S)==root)];
FMM <- find.best.f(L,S,n.round=3, f.criterion="F", verbose=TRUE, b.per.example=TRUE);
```

find.leaves

Leaves

Description

Find the leaves of a directed graph

Usage

```
find.leaves(g)
```

Arguments

`g` a graph of class `graphNEL`. It represents the hierarchy of the classes.

Value

a vector with the names of the leaves of g

Examples

```
data(graph);
leaves <- find.leaves(g);
```

full.annotation.matrix

Full annotations matrix

Description

Construct a full annotations table using ancestors and the most specific annotations table w.r.t. a given weighted adjacency matrix (wadj). The rows of the full annotations matrix correspond to all the examples of the given weighted adjacency matrix and the columns to the class/terms. The transitive closure of the annotations is performed.

Usage

```
full.annotation.matrix(W, anc, ann.spec)
```

Arguments

W	symmetric adjacency weighted matrix of the graph
anc	list of the ancestors of the ontology.
ann.spec	the annotation matrix of the most specific annotations (0/1): rows are genes and columns are classes.

Details

The examples present in the annotation matrix (ann.spec) but not in the adjacency weighted matrix (W) are purged.

Value

a full annotation table T, that is a matrix in which the transitive closure of annotations was performed. Rows correspond to genes of the weighted adjacency matrix and columns to terms. $T[i, j] = 1$ means that gene i is annotated for the term j , $T[i, j] = 0$ means that gene i is not annotated for the term j .

See Also

[weighted.adjacency.matrix](#), [build.ancestors](#), [specific.annotation.matrix](#), [transitive.closure.annotations](#)

Examples

```
data(wadj);
data(graph);
data(labels);
anc <- build.ancestors(g);
full.ann <- full.annotation.matrix(W, anc, L);
```

graph.levels

Build Graph Levels

Description

This function groups a set of nodes in according to their maximum depth in the graph. It first inverts the weights of the graph and then applies the Bellman Ford algorithm to find the shortest path, achieving in this way the longest path

Usage

```
graph.levels(g, root = "00")
```

Arguments

g	an object of class graphNEL
root	name of the root node (def. root="00")

Value

a list of the nodes grouped w.r.t. the distance from the root: the first element of the list corresponds to the root node (level 0), the second to nodes at maximum distance 1 (level 1), the third to the node at maximum distance 3 (level 2) and so on.

Examples

```
data(graph);
root <- root.node(g);
lev <- graph.levels(g, root=root);
```

hierarchical.checkers *Hierarchical constraints checker*

Description

Check if the true path rule is violated or not. In other words this function checks if the score of a parent or an ancestor node is always larger or equal than that of its children or descendants nodes

Usage

```
check.hierarchy.single.sample(y.hier, g, root = "00")
```

```
check.hierarchy(S.hier, g, root = "00")
```

Arguments

y.hier	vector of scores relative to a single example. This must be a named numeric vector
g	a graph of class graphNEL. It represents the hierarchy of the classes
root	name of the class that it is the top-level (root) of the hierarchy (def:00)
S.hier	the matrix with the scores of the classes corrected in according to hierarchy

Value

return a list of 3 elements:

- Status:
 - OK if none hierarchical constraints have been broken;
 - NOTOK if there is at least one hierarchical constraint broken;
- Hierarchy_Constraints_Broken:
 - TRUE: example did not respect the hierarchical constraints;
 - FALSE: example broke the hierarchical constraints;
- Hierarchy_constraints_satisfied: how many terms satisfied the hierarchical constraint

See Also

[htd](#)

Examples

```
data(graph);
data(scores);
root <- root.node(g);
S.hier <- htd(S,g,root);
S.hier.single.example <- S.hier[sample(ncol(S.hier),1),];
check.hierarchy.single.sample(S.hier.single.example, g, root=root);
check.hierarchy(S.hier, g, root);
```

htd

HTD-DAG

Description

Implementation of a top-down procedure to correct the scores in the hierarchy according to the constraints that the score of a node cannot be greater than a score of its parents.

Usage

```
htd(S, g, root = "00")
```

Arguments

S	a named flat scores matrix with examples on rows and classes on columns
g	a graph of class graphNEL. It represents the hierarchy of the classes
root	name of the class that it is the top-level (root) of the hierarchy (def: 00)

Details

The HTD-DAG algorithm modifies the flat scores according to the hierarchy of a DAG through a unique run across the nodes of the graph. For a given example $x \in X$, the flat predictions $f(x) = \hat{y}$ are hierarchically corrected to \bar{y} , by per-level visiting the nodes of the DAG from top to bottom according to the following simple rule:

$$\bar{y}_i := \begin{cases} \hat{y}_i & \text{if } i \in \text{root}(G) \\ \min_{j \in \text{par}(i)} \bar{y}_j & \text{if } \min_{j \in \text{par}(i)} \bar{y}_j < \hat{y}_i \\ \hat{y}_i & \text{otherwise} \end{cases}$$

The node levels correspond to their maximum path length from the root.

Value

a matrix with the scores of the classes corrected according to the HTD-DAG algorithm.

See Also

[graph.levels](#), [hierarchical.checkers](#)

Examples

```
data(graph);
data(scores);
root <- root.node(g);
S.htd <- htd(S,g,root);
```

Multilabel.F.measure *Multilabel F-measure*

Description

Method for computing Precision, Recall, Specificity, Accuracy and F-measure for multiclass multilabel classification

Usage

```
F.measure.multilabel(target, predicted, b.per.example = FALSE)
```

```
## S4 method for signature 'matrix,matrix'
```

```
F.measure.multilabel(target, predicted,
```

```
  b.per.example = FALSE)
```

Arguments

- | | |
|---------------|---|
| target | matrix with the target multilabels: rows correspond to examples and columns to classes. $target[i, j] = 1$ if example i belongs to class j , $target[i, j] = 0$ otherwise |
| predicted | a numeric matrix with predicted values (scores): rows correspond to examples and columns to classes |
| b.per.example | boolean. <ul style="list-style-type: none"> • TRUE: results are returned for each example; • FALSE: only the average results are returned |

Value

Two different outputs respect to the input parameter `b.per.example`:

- `b.per.example==FALSE`: a list with a single element average. A named vector with average precision (P), recall (R), specificity (S), F-measure (F), average F-measure (avF) and Accuracy (A) across examples. F is the F-measure computed as the harmonic mean between the average precision and recall; av.F is the F-measure computed as the average across examples.
- `b.per.example==TRUE`: a list with two elements:
 1. average: a named vector with average precision (P), recall (R), specificity (S), F-measure (F), average F-measure (avF) and Accuracy (A) across examples;
 2. per.example: a named matrix with the Precision (P), Recall (R), Specificity (S), Accuracy (A), F-measure (F) and av.F-measure (av.F) for each example. Row names correspond to examples, column names correspond respectively to Precision (P), Recall (R), Specificity (S), Accuracy (A), F-measure (F) and av.F-measure (av.F)

Examples

```
data(labels);
data(scores);
data(graph);
root <- root.node(g);
L <- L[, -which(colnames(L)==root)];
S <- S[, -which(colnames(S)==root)];
S[S>0.7] <- 1;
S[S<0.7] <- 0;
FMM <- F.measure.multilabel(L,S);
```

`normalize.max`*Max normalization*

Description

Function to normalize the scores of a flat scores matrix per class

Usage

```
normalize.max(S)
```

Arguments

`S` matrix with the raw non normalized scores. Rows are examples and columns are classes

Details

The scores of each class are normalized by dividing the score values for the maximum score of that class. If the max score of a class is zero, no normalization is needed, otherwise NaN value will be printed as results of 0 out of 0 division.

Value

A score matrix with the same dimensions of `S`, but with scores max/normalized separately for each class

Examples

```
data(scores);
maxnorm <- normalize.max(S);
```

parents

Build parents

Description

Compute the parents for each node of a graph

Usage

```
get.parents(g, root = "00")
```

```
get.parents.top.down(g, levels, root = "00")
```

```
get.parents.bottom.up(g, levels, root = "00")
```

```
get.parents.topological.sorting(g, root = "00")
```

Arguments

<code>g</code>	a graph of class graphNEL. It represents the hierarchy of the classes
<code>root</code>	name of the root node (def. <code>root="00"</code>)
<code>levels</code>	a list of character vectors. Each component represents a graph level and the elements of any component correspond to nodes. The level 0 coincides with the root node.

Value

`get.parents` returns a named list of character vectors. Each component corresponds to a node x of the graph (i.e. child node) and its vector is the set of its parents (the root node is not included)

`get.parents.top.down` returns a named list of character vectors. Each component corresponds to a node x of the graph (i.e. child node) and its vector is the set of its parents. The nodes order follows the levels of the graph from root (excluded) to leaves.

`get.parents.bottom.up` returns a named list of character vectors. Each component corresponds to a node x of the graph (i.e. child node) and its vector is the set of its parents. The nodes are ordered from leaves to root (excluded).

`get.parents.topological.sorting` a named list of character vectors. Each component corresponds to a node x of the graph (i.e. child node) and its vector is the set of its parents. The nodes are ordered according to a topological sorting, i.e. parents node come before children node.

See Also

[graph.levels](#)

Examples

```

data(graph);
root <- root.node(g)
parents <- get.parents(g, root=root);
lev <- graph.levels(g, root=root);
parents.tod <- get.parents.top.down(g, lev, root=root);
parents.bup <- get.parents.bottom.up(g, lev, root=root);
parents.tsort <- get.parents.topological.sorting(g, root=root);

```

read.graph	<i>Read a directed graph from a file</i>
------------	--

Description

A directed graph is read from a file and a graphNEL object is built

Usage

```
read.graph(file = "graph.txt")
```

Arguments

file	name of the file to be read. The format of the file is a sequence of rows and each row corresponds to an edge represented through a pair of vertices separated by blanks
------	--

Value

an object of class graphNEL

Examples

```

ed <- system.file("extdata/graph.edges.txt", package= "HEMDAG");
g <- read.graph(file=ed);

```

read.undirected.graph	<i>Read an undirected graph from a file</i>
-----------------------	---

Description

The graph is read from a file and a graphNEL object is built. The format of the input file is a sequence of rows. Each row corresponds to an edge represented through a pair of vertices separated by blanks, and the weight of the edge.

Usage

```
read.undirected.graph(file = "graph.txt")
```

Arguments

file name of the file to be read

Value

a graph of class graphNEL

Examples

```
edges <- system.file("extdata/edges.txt" ,package="HEMDAG");  
g <- read.undirected.graph(file=edges);
```

<i>root.node</i>	<i>Root node</i>
------------------	------------------

Description

Find the root node of a directed graph

Usage

```
root.node(g)
```

Arguments

g a graph of class graphNEL. It represents the hierarchy of the classes.

Value

name of the root node

Examples

```
data(graph);  
root <- root.node(g);
```

specific.annotation.list

Specific annotations list

Description

Construct a list of the most specific annotations starting from the table of the most specific annotations

Usage

```
specific.annotation.list(ann)
```

Arguments

ann annotation matrix (0/1). Rows are examples and columns are most specific terms. It must be a named matrix.

Value

a named list, where the names of each component correspond to an examples (genes) and the elements of each component are the most specific classes associated to that genes

See Also

[specific.annotation.matrix](#)

Examples

```
data(labels);  
spec.list <- specific.annotation.list(L);
```

specific.annotation.matrix

HPO specific annotations matrix

Description

Construct the labels matrix of the most specific HPO terms

Usage

```
specific.annotation.matrix(file = "gene2pheno.txt", genename = "TRUE")
```

Arguments

- file** text file representing the most specific associations gene-HPO term (def: "gene2pheno.txt"). The file must be written as sequence of rows. Each row represents a gene and all its associations with abnormal phenotype tab separated, e.g.: *gene_1 <tab> phen1 <tab> ... phen_N*
- genename** boolean value:
- TRUE (def.): the names of genes are *gene symbol* (i.e. characters);
 - FALSE: the names of gene are *entrez gene ID* (i.e. integer numbers);

Value

the annotation matrix of the most specific annotations (0/1): rows are genes and columns are HPO terms. Let's denote M the labels matrix. If $M[i, j] = 1$, means that the gene i is annotated with the class j , otherwise $M[i, j] = 0$.

Examples

```
gene2pheno <- system.file("extdata/gene2pheno.txt", package="HEMDAG");
spec.ann <- specific.annotation.matrix(gene2pheno, genename=TRUE);
```

```
stratified.cross.validation
      Stratified cross validation
```

Description

Generate data for the stratified cross-validation

Usage

```
do.stratified.cv.data.single.class(examples, positives, kk = 5, seed = NULL)
do.stratified.cv.data.over.classes(labels, examples, kk = 5, seed = NULL)
```

Arguments

- examples** indices or names of the examples. Can be either a vector of integers or a vector of names.
- positives** vector of integers or vector of names. The indices (or names) refer to the indices (or names) of 'positive' examples
- kk** number of folds (def=5)
- seed** seed of the random generator (def=NULL). If is set to NULL no initialization is performed
- labels** labels matrix. Rows are genes and columns are classes. Let's denote M the labels matrix. If $M[i, j] = 1$, means that the gene i is annotated with the class j , otherwise $M[i, j] = 0$.

Value

`do.stratified.cv.data.single.class` returns a list with 2 two component:

- `fold.non.positives`: a list with k components. Each component is a vector with the indices (or names) of the non-positive elements. Indices (or names) refer to row numbers (or names) of a data matrix.
- `fold.positives`: a list with k components. Each component is a vector with the indices (or names) of the positive elements. Indices (or names) refer to row numbers (or names) of a data matrix.

`do.stratified.cv.data.over.classes` returns a list with n components, where n is the number of classes of the labels matrix. Each component n is in turn a list with k elements, where k is the number of folds. Each fold contains an equal amount of examples positives and negatives.

Examples

```
data(labels);
examples.index <- 1:nrow(L);
examples.name <- rownames(L);
positives <- which(L[,3]==1);
x <- do.stratified.cv.data.single.class(examples.index, positives, kk=5, seed=23);
y <- do.stratified.cv.data.single.class(examples.name, positives, kk=5, seed=23);
z <- do.stratified.cv.data.over.classes(L, examples.index, kk=5, seed=23);
k <- do.stratified.cv.data.over.classes(L, examples.name, kk=5, seed=23);
```

 TPR-DAG

TPR-DAG variants

Description

Different variants of the TPR-DAG algorithm are implemented. In their more general form the TPR-DAG algorithms adopt a two step learnig strategy:

1. in the first step they compute a *per-level bottom-up* visit from the leaves to the root to propagate positive predictions across the hierarchy;
2. in the second step they compute a *per-level top-down* visit from the root to the leaves in order to assure the hierarchical consistency of the predictions

Usage

```
tpr.threshold(S, g, root = "00", t = 0.5)
```

```
tpr.threshold.free(S, g, root = "00")
```

```
tpr.weighted.threshold.free(S, g, root = "00", w = 0.5)
```

```
tpr.weighted.threshold(S, g, root = "00", t = 0.5, w = 0.5)
```

Arguments

S	a named flat scores matrix with examples on rows and classes on columns
g	a graph of class graphNEL. It represents the hierarchy of the classes
root	name of the class that it is on the top-level of the hierarchy (def. root="00")
t	threshold for the choice of positive children (def. t=0.5)
w	weight to balance between the contribution of the node i and that of its positive children

Details

The *vanilla* TPR-DAG adopts a per-level bottom-up traversal of the DAG to correct the flat predictions \hat{y}_i :

$$\bar{y}_i := \frac{1}{1 + |\phi_i|} (\hat{y}_i + \sum_{j \in \phi_i} \bar{y}_j)$$

where ϕ_i are the positive children of i . Different strategies to select the positive children ϕ_i can be applied:

1. **Threshold-Free** strategy: the positive nodes are those children that can increment the score of the node i , that is those nodes that achieve a score higher than that of their parents:

$$\phi_i := \{j \in \text{child}(i) | \bar{y}_j > \hat{y}_i\}$$

2. **Threshold** strategy: the positive children are selected on the basis of a threshold that can be selected in two different ways:

- (a) for each node a constant threshold \bar{t} is a priori selected:

$$\phi_i := \{j \in \text{child}(i) | \bar{y}_j > \bar{t}\}$$

For instance if the predictions represent probabilities it could be meaningful to a priori select $\bar{t} = 0.5$.

- (b) the threshold is selected to maximize some performance metric \mathcal{M} estimated on the training data, as for instance the F-score or the AUPRC. In other words the threshold is selected to maximize some measure of accuracy of the predictions $\mathcal{M}(j, t)$ on the training data for the class j with respect to the threshold t . The corresponding set of positives $\forall i \in V$ is:

$$\phi_i := \{j \in \text{child}(i) | \bar{y}_j > t_j^*, t_j^* = \arg \max_t \mathcal{M}(j, t)\}$$

For instance t_j^* can be selected from a set of $t \in (0, 1)$ through internal cross-validation techniques.

Furthermore it is possible to add a weight $w \in [0, 1]$ to balance between the contribution of the node i and that of its positive children ϕ , through their convex combination, obtaining in this way the weighted TPR-DAG version:

$$\bar{y}_i := w \hat{y}_i + \frac{(1-w)}{|\phi_i|} \sum_{j \in \phi_i} \bar{y}_j$$

Value

a named matrix with the scores of the classes corrected according to the TPR-DAG algorithm.

Examples

```
data(graph);
data(scores);
data(labels);
root <- root.node(g);
TPR.TF <- tpr.threshold.free(S,g,root);
TPR.T <- tpr.threshold(S,g,root,t=0.5);
TPR.W <- tpr.weighted.threshold.free(S,g,root,w=0.5);
TPR.WT <- tpr.weighted.threshold(S,g,root,w=0.5, t=0.5);
```

TPR.DAG.CV

TPR-DAG cross-validation

Description

High level function to correct the scores with a hierarchy according to one of the TPR-DAG variants performing cross-validation experiments

Usage

```
Do.tpr.threshold.cv(threshold = seq(from = 0.1, to = 0.9, by = 0.1), kk = 5,
  seed = NULL, norm = TRUE, norm.type = "NONE", flat.file = flat.file,
  ann.file = ann.file, dag.file = dag.file, flat.dir = flat.dir,
  ann.dir = ann.dir, dag.dir = dag.dir, flat.norm.dir = NULL,
  n.round = 3, f.criterion = "F", hierScore.dir = hierScore.dir,
  perf.dir = perf.dir)
```

```
Do.tpr.weighted.threshold.free.cv(weight = seq(from = 0.1, to = 1, by = 0.1),
  kk = 5, seed = NULL, norm = TRUE, norm.type = "NONE",
  flat.file = flat.file, ann.file = ann.file, dag.file = dag.file,
  flat.dir = flat.dir, ann.dir = ann.dir, dag.dir = dag.dir,
  flat.norm.dir = NULL, n.round = 3, f.criterion = "F",
  hierScore.dir = "hierScore.dir/", perf.dir = "perf.dir/")
```

```
Do.tpr.weighted.threshold.cv(threshold = seq(from = 0.1, to = 0.9, by = 0.1),
  weight = seq(from = 0.1, to = 1, by = 0.1), kk = 5, seed = NULL,
  norm = TRUE, norm.type = "NONE", flat.file = flat.file,
  ann.file = ann.file, dag.file = dag.file, flat.dir = flat.dir,
  ann.dir = ann.dir, dag.dir = dag.dir, flat.norm.dir = NULL,
  n.round = 3, f.criterion = "F", hierScore.dir = hierScore.dir,
  perf.dir = perf.dir)
```

Arguments

threshold	range of threshold values to be tested in order to find the best threshold (def: from:0.1, to:0.9, by:0.1 step). The denser the range is, the higher the probability to find the best threshold is, but obviously the execution time will be increasing.
kk	number of folds of the cross validation (def: kk=5);
seed	initialization seed for the random generator to create folds (def:0). If NULL (def.) no initialization is performed
norm	boolean value: <ul style="list-style-type: none"> • TRUE (def.): the flat scores matrix has been already normalized in according to a normalization method; • FALSE: the flat scores matrix has not been normalized yet. See the parameter norm.type for which normalization can be applied.
norm.type	three values: <ol style="list-style-type: none"> 1. NONE (def.): set norm.type to NONE if and only if the parameter norm is set to TRUE; 2. MaxNorm: each score is divided w.r.t. the max of each class; 3. Qnorm: quantile normalization. preprocessCore package is used.
flat.file	name of the file containing the flat scores matrix to be normalized or already normalized (without rda extension)
ann.file	name of the file containing the the label matrix of the examples (without rda extension)
dag.file	name of the file containing the graph that represents the hierarchy of the classes (without rda extension)
flat.dir	relative path where flat scores matrix is stored
ann.dir	relative path where annotation matrix is stored
dag.dir	relative path where graph is stored
flat.norm.dir	relative path where flat normalized scores matrix must be stored. Use this parameter if and only if norm is set to FALSE, otherwise set flat.norm.dir to NULL (def.)
n.round	number of rounding digits to be applied to the hierarchical scores matrix (def. 3). It is used for choosing the best threshold on the basis of the best F-measure
f.criterion	character. Type of F-measure to be used to select the best F-measure. Two possibilities: <ol style="list-style-type: none"> 1. F (def.): corresponds to the harmonic mean between the average precision and recall 2. avF: corresponds to the per-example F-score averaged across all the examples
hierScore.dir	relative path where the hierarchical scores matrix must be stored
perf.dir	relative path where the term-centric and protein-centric measures must be stored
weight	range of weight values to be tested in order to find the best weight (def: from:0.1, to:0.9, by:0.1 step). The denser the range is, the higher the probability to find the best weight is, but obviously the execution time will be increasing.

Details

These high level functions perform a classical cross-validation procedure to find the best threshold maximizing on F-score measure.

Value

Five rda files stored in the respective output directories:

1. `hierarchical_scores_matrix`: a matrix with examples on rows and classes on columns representing the computed hierarchical scores for each class and example considered. It stored in `hierScore.dir` directory.
2. PCM (Protein Centric Measures) average and per-example: compute F-score measure by `find.best.f` function. It stored in `perf.dir` directory.
3. PRC (area under Precision-Recall Curve) average and per.class: compute PRC by **precrec** package. It stored in `perf.dir` directory.
4. AUC (Area Under ROC Curve) average and per-class: compute AUC by **precrec** package. It stored in `perf.dir` directory.
5. PxR (Precision at fixed Recall levels) average and per classes: compute PxR by **PerfMeas** package. It stored in `perf.dir` directory.

Examples

```
data(graph);
data(scores);
data(labels);
if (!dir.exists("data")){
  dir.create("data");
}
if (!dir.exists("results")){
  dir.create("results");
}
save(g,file="data/graph.rda");
save(L,file="data/labels.rda");
save(S,file="data/scores.rda");
dag.dir <- flat.dir <- flat.norm.dir <- ann.dir <- "data/";
hierScore.dir <- perf.dir <- "results/";
dag.file <- "graph";
flat.file <- "scores";
ann.file <- "labels";
threshold <- seq(from=0.1, to=0.5, by=0.1);
Do.tpr.threshold.cv(threshold=threshold, kk=5, seed=23, norm=FALSE, norm.type= "MaxNorm",
flat.file=flat.file, ann.file=ann.file, dag.file=dag.file, flat.dir=flat.dir, ann.dir=ann.dir,
dag.dir=dag.dir, flat.norm.dir=flat.norm.dir, n.round=3, f.criterion ="F",
hierScore.dir=hierScore.dir, perf.dir=perf.dir);
```

TPR.DAG.HOLDOUT

TPR-DAG holdout

Description

High level function to correct the scores with a hierarchy according to one of the TPR-DAG variants performing a classical holdout procedure.

Usage

```
Do.tpr.threshold.holdout(threshold = seq(from = 0.1, to = 0.9, by = 0.1),
  kk = 5, seed = NULL, norm = TRUE, norm.type = "NONE",
  flat.file = flat.file, ann.file = ann.file, dag.file = dag.file,
  ind.test.set = ind.test.set, ind.dir = ind.dir, flat.dir = flat.dir,
  ann.dir = ann.dir, dag.dir = dag.dir, flat.norm.dir = NULL,
  n.round = 3, f.criterion = "F", hierScore.dir = hierScore.dir,
  perf.dir = perf.dir)
```

```
Do.tpr.weighted.threshold.free.holdout(weight = seq(from = 0.1, to = 1, by =
  0.1), kk = 5, seed = NULL, norm = TRUE, norm.type = "NONE",
  flat.file = flat.file, ann.file = ann.file, dag.file = dag.file,
  ind.test.set = ind.test.set, ind.dir = ind.dir, flat.dir = flat.dir,
  ann.dir = ann.dir, dag.dir = dag.dir, flat.norm.dir = NULL,
  n.round = 3, f.criterion = "F", hierScore.dir = hierScore.dir,
  perf.dir = perf.dir)
```

```
Do.tpr.weighted.threshold.holdout(threshold = seq(from = 0.1, to = 0.9, by =
  0.1), weight = seq(from = 0.1, to = 1, by = 0.1), kk = 5, seed = NULL,
  norm = TRUE, norm.type = "NONE", flat.file = flat.file,
  ann.file = ann.file, dag.file = dag.file, ind.test.set = ind.test.set,
  ind.dir = ind.dir, flat.dir = flat.dir, ann.dir = ann.dir,
  dag.dir = dag.dir, flat.norm.dir = NULL, n.round = 3,
  f.criterion = "F", hierScore.dir = hierScore.dir, perf.dir = perf.dir)
```

Arguments

- | | |
|-----------|--|
| threshold | range of threshold values to be tested in order to find the best threshold (def: from:0.1, to:0.9, by:0.1 step). The denser the range is, the higher the probability to find the best threshold is, but obviously the execution time will be increasing. |
| kk | number of folds of the cross validation (def: kk=5); |
| seed | initialization seed for the random generator to create folds (def:0). If NULL (def.) no initialization is performed |
| norm | boolean value: <ul style="list-style-type: none"> TRUE (def.): the flat scores matrix has been already normalized in according to a normalization method; |

	<ul style="list-style-type: none"> • FALSE: the flat scores matrix has not been normalized yet. See the parameter <code>norm.type</code> for which normalization can be applied.
<code>norm.type</code>	<p>three values:</p> <ol style="list-style-type: none"> 1. NONE (def.): set <code>norm.type</code> to NONE if and only if the parameter <code>norm</code> is set to TRUE; 2. MaxNorm: each score is divided w.r.t. the max of each class; 3. Qnorm: quantile normalization. preprocessCore package is used.
<code>flat.file</code>	name of the file containing the flat scores matrix to be normalized or already normalized (without rda extension)
<code>ann.file</code>	name of the file containing the the label matrix of the examples (without rda extension)
<code>dag.file</code>	name of the file containing the graph that represents the hierarchy of the classes (without rda extension)
<code>ind.test.set</code>	name of the file containing a vector of integer numbers corresponding to the indices of the elements (rows) of scores matrix to be used in the test set
<code>ind.dir</code>	relative path to folder where <code>ind.test.set</code> is stored
<code>flat.dir</code>	relative path where flat scores matrix is stored
<code>ann.dir</code>	relative path where annotation matrix is stored
<code>dag.dir</code>	relative path where graph is stored
<code>flat.norm.dir</code>	relative path where flat normalized scores matrix must be stored. Use this parameter if and only if <code>norm</code> is set to FALSE, otherwise set <code>flat.norm.dir</code> to NULL (def.)
<code>n.round</code>	number of rounding digits to be applied to the hierarchical scores matrix (def. 3). It is used for choosing the best threshold on the basis of the best F-measure
<code>f.criterion</code>	<p>character. Type of F-measure to be used to select the best F-measure. Two possibilities:</p> <ol style="list-style-type: none"> 1. F (def.): corresponds to the harmonic mean between the average precision and recall 2. avF: corresponds to the per-example F-score averaged across all the examples
<code>hierScore.dir</code>	relative path where the hierarchical scores matrix must be stored
<code>perf.dir</code>	relative path where the term-centric and protein-centric measures must be stored
<code>weight</code>	range of weight values to be tested in order to find the best weight (def: from:0.1, to:0.9, by:0.1 step). The denser the range is, the higher the probability to find the best weight is, but obviously the execution time will be increasing.

Details

These high level functions perform a classical holdout procedure to find the best threshold maximizing on F.score measure.

Value

Five rda files stored in the respective output directories:

1. hierarchical scores matrix: a matrix with examples on rows and classes on columns representing the computed hierarchical scores for each class and example considered. It stored in hierScore.dir directory.
2. PCM (Protein Centric Measures) average and per-example: compute F-score measure by find.best.f function. It stored in perf.dir directory.
3. PRC (area under Precision-Recall Curve) average and per.class: compute PRC by **precrec** package. It stored in perf.dir directory.
4. AUC (Area Under ROC Curve) average and per-class: compute AUC by **precrec** package. It stored in perf.dir directory.
5. PxR (Precision at fixed Recall levels) average and per classes: compute PxR by **PerfMeas** package. It stored in perf.dir directory.

Examples

```

data(graph);
data(scores);
data(labels);
data(test.index);
if (!dir.exists("data")){
  dir.create("data");
}
if (!dir.exists("results")){
  dir.create("results");
}
save(g,file="data/graph.rda");
save(L,file="data/labels.rda");
save(S,file="data/scores.rda");
save(test.index, file="data/test.index.rda");
ind.dir <- dag.dir <- flat.dir <- flat.norm.dir <- ann.dir <- "data/";
hierScore.dir <- perf.dir <- "results/";
ind.test.set <- "test.index";
dag.file <- "graph";
flat.file <- "scores";
ann.file <- "labels";
threshold <- seq(from=0.1, to=0.3, by=0.1);
Do.tpr.threshold.holdout(threshold=threshold, kk=5, seed=23, norm=FALSE, norm.type="MaxNorm",
  flat.file=flat.file, ann.file=ann.file, dag.file=dag.file, ind.test.set=ind.test.set,
  ind.dir=ind.dir, flat.dir=flat.dir, ann.dir=ann.dir, dag.dir=dag.dir, flat.norm.dir=flat.norm.dir,
  n.round=3, f.criterion="F", hierScore.dir=hierScore.dir, perf.dir=perf.dir);

```

transitive.closure.annotations

Transitive closure of annotations

Description

Performs the transitive closure of the annotations using ancestors and the most specific annotation table. The annotations are propagated from bottom to top, enriching the most specific annotations table. The rows of the matrix correspond to the genes of the most specific annotation table and the columns to the HPO terms/classes

Usage

```
transitive.closure.annotations(ann.spec, anc)
```

Arguments

ann.spec	the annotation matrix of the most specific annotations (0/1): rows are genes and columns are HPO terms.
anc	list of the ancestors of the ontology.

Value

an annotation table T: rows correspond to genes and columns to HPO terms. $T[i, j] = 1$ means that gene i is annotated for the term j , $T[i, j] = 0$ means that gene i is not annotated for the term j .

See Also

[specific.annotation.matrix](#), [build.ancestors](#)

Examples

```
data(graph);
data(labels);
anc <- build.ancestors(g);
tca <- transitive.closure.annotations(L, anc);
```

weighted.adjacency.matrix
Weighted Adjacency Matrix

Description

Construct a Weighted Adjacency Matrix (wadj matrix) of a graph

Usage

```
weighted.adjacency.matrix(file = "edges.txt", compressed = TRUE,
  nodename = TRUE)
```

Arguments

file	name of the plain text file to be read (def. edges). The format of the file is a sequence of rows. Each row corresponds to an edge represented through a pair of vertices separated by blanks and the weight of the edge. For instance: nodeX nodeY score
compressed	boolean value: <ul style="list-style-type: none"> • TRUE (def.): the input file must be in a .gz compressed format; • FALSE: the input file must be in a plain text format;
nodename	boolean value: <ul style="list-style-type: none"> • TRUE (def.): the names of nodes are gene symbol (i.e. characters); • FALSE: the names of the nodes are entrez gene ID (i.e. integer numbers);

Details

The input parameter nodename sorts the row names of the wadj matrix in increasing order if they are integer number or in alphabetic order if they are characters.

Value

a named symmetric weighted adjacency matrix of the graph

Examples

```
edges <- system.file("extdata/edges.txt", package="HEMDAG");
W <- weighted.adjacency.matrix(file=edges, compressed=FALSE, nodename=TRUE);
```

write.graph	<i>Write a directed graph on file</i>
-------------	---------------------------------------

Description

An object of class graphNEL is read and the graph is written on a plain text file as sequence of rows

Usage

```
write.graph(g, file = "graph.txt")
```

Arguments

g	a graph of class graphNEL
file	name of the file to be written

Value

a plain text file representing the graph. Each row corresponds to an edge represented through a pair of vertices separated by blanks

Examples

```
## Not run:  
data(graph);  
write.graph(g, file="graph.edges.txt");  
## End(Not run)
```

Index

*Topic **package**

- HEMDAG-package, 3
- ancestors, 3
- AUPRC.single.class, 4, 5
- AUPRC.single.over.classes, 5, 5
- AUROC.single.class, 6, 7
- AUROC.single.over.classes, 6, 7
- build.ancestors, 8, 29, 48
- build.ancestors (ancestors), 3
- build.children (children), 9
- build.descendants (descendants), 11
- check.annotation.matrix.integrity, 8
- check.DAG.integrity, 9
- check.hierarchy
 - (hierarchical.checkers), 31
- children, 9
- compute.flipped.graph, 10
- constraints.matrix, 11
- descendants, 11
- distances.from.leaves, 12
- do.edges.from.HPO.obo, 13
- Do.FLAT.scores.normalization, 13
- Do.full.annotation.matrix, 15
- Do.HTD, 16
- Do.HTD.holdout, 18
- do.stratified.cv.data.over.classes
 - (stratified.cross.validation), 39
- do.stratified.cv.data.single.class
 - (stratified.cross.validation), 39
- do.subgraph, 20
- do.submatrix, 21
- Do.tpr.threshold.cv (TPR.DAG.CV), 42
- Do.tpr.threshold.free, 21
- Do.tpr.threshold.free.holdout, 23
- Do.tpr.threshold.holdout
 - (TPR.DAG.HOLDOUT), 45
- Do.tpr.weighted.threshold.cv
 - (TPR.DAG.CV), 42
- Do.tpr.weighted.threshold.free.cv
 - (TPR.DAG.CV), 42
- Do.tpr.weighted.threshold.free.holdout
 - (TPR.DAG.HOLDOUT), 45
- Do.tpr.weighted.threshold.holdout
 - (TPR.DAG.HOLDOUT), 45
- do.unstratified.cv.data, 25
- example.datasets, 26
- F.measure.multilabel
 - (Multilabel.F.measure), 33
- F.measure.multilabel,matrix,matrix-method
 - (Multilabel.F.measure), 33
- find.best.f, 27
- find.leaves, 28
- full.annotation.matrix, 8, 15, 29
- g (example.datasets), 26
- get.children.bottom.up (children), 9
- get.children.top.down (children), 9
- get.parents (parents), 35
- graph.levels, 4, 10, 12, 30, 32, 35
- HEMDAG (HEMDAG-package), 3
- HEMDAG-package, 3
- hierarchical.checkers, 31, 32
- htd, 31, 32
- L (example.datasets), 26
- Multilabel.F.measure, 33
- normalize.max, 34
- parents, 35
- read.graph, 36

read.undirected.graph, 36
root.node, 37

S (example.datasets), 26
specific.annotation.list, 38
specific.annotation.matrix, 29, 38, 38,
48
stratified.cross.validation, 39

test.index (example.datasets), 26
TPR-DAG, 40
TPR.DAG.CV, 42
TPR.DAG.HOLDOUT, 45
tpr.threshold (TPR-DAG), 40
tpr.weighted.threshold (TPR-DAG), 40
transitive.closure.annotations, 8, 29,
47

W (example.datasets), 26
weighted.adjacency.matrix, 29, 48
write.graph, 49