

Package ‘MoEClust’

February 6, 2018

Type Package

Date 2018-02-06

Title Parsimonious Model-Based Clustering with Covariates

Version 1.1.0

Description Clustering via parsimonious Mixtures of Experts using the MoEClust models introduced by Murphy and Murphy (2017) <arXiv:1711.05632>. This package fits finite Gaussian mixture models with gating and expert network covariates using parsimonious covariance parameterisations from 'mclust' via the EM algorithm. Visualisation of the results of such models using generalised pairs plots is also facilitated.

Depends R (>= 2.10)

License GPL (>= 2)

Encoding UTF-8

URL <https://cran.r-project.org/package=MoEClust>

BugReports <https://github.com/Keefe-Murphy/MoEClust>

LazyData true

Imports lattice, matrixStats, mclust (>= 5.1), mvnfast, nnet, vcd

Suggests cluster, geometry, knitr, rmarkdown

RoxygenNote 6.0.1

VignetteBuilder knitr

NeedsCompilation no

Author Keefe Murphy [aut, cre],
Thomas Brendan Murphy [ctb]

Maintainer Keefe Murphy <keefe.murphy@ucd.ie>

Repository CRAN

Date/Publication 2018-02-06 14:41:03 UTC

R topics documented:

ais	2
as.Mclust	3
CO2data	5
drop_constants	6
drop_levels	7
expert_covar	8
MoEClust	9
MoE_aitken	11
MoE_clust	12
MoE_compare	17
MoE_control	20
MoE_crit	23
MoE_dens	25
MoE_estep	27
MoE_gpairs	29
MoE_mahala	33
MoE_plotCrit	34
MoE_plotGate	35
MoE_plotLogLik	36
MoE_qclass	37
noise_vol	37
plot.MoEClust	38
Index	41

ais	<i>Australian Institute of Sport data</i>
-----	---

Description

Data on 102 male and 100 female athletes collected at the Australian Institute of Sport, courtesy of Richard Telford and Ross Cunningham.

Usage

```
data(ais)
```

Format

A data frame with 202 observations on the following 13 variables:

- ,1 sex - categorical, levels = female, male
- ,2 sport - categorical, levels = B_Ball, Field, Gym, Netball, Row, Swim, T_400m, Tennis, T_Sprnt, W_Polo
- ,3 RCC - red cell count (numeric)
- ,4 WCC - white cell count (numeric)

- ,5 Hc - Hematocrit (numeric)
- ,6 Hg - Hemoglobin (numeric)
- ,7 Fe - plasma ferritin concentration (numeric)
- ,8 BMI - body mass index: $Wt/(Ht)^2$ (numeric)
- ,9 SSF - sum of skin folds (numeric)
- ,10 Bfat - body fat percentage (numeric)
- ,11 LBM - lean body mass (numeric)
- ,12 Ht - height, cm (numeric)
- ,13 Wt - weight, kg (numeric)

Details

The data have been made publicly available in connection with the book by Cook and Weisberg (1994).

References

Cook, R. D. and Weisberg, S. (1994), *An Introduction to Regression Graphics*. John Wiley & Sons, New York.

Examples

```
data(ais, package="MoEClust")
pairs(ais[,c(3:4, 10:13)], col=as.numeric(ais[,1]), main = "AIS data")
```

as.Mclust

Convert MoEClust objects to the Mclust class

Description

Converts an object of class "MoEClust" generated by [MoE_clust](#) and converts it to an object of class "Mclust" as generated by fitting [Mclust](#), to facilitate use of plotting and other functions for the "Mclust" class within the **mclust** package.

Usage

```
as.Mclust(x, resid = FALSE, signif = 0, ...)
```

Arguments

x	An object of class "MoEClust" generated by <code>MoE_clust</code> or an object of class "MoECompare" generated by <code>MoE_compare</code> .
resid	Logical indicating whether to treat the data as the raw data (when FALSE, the default) or the augmented data comprising the residuals from the expert network (when TRUE and the model actually contains expert covariates). In the latter case, the mean and (co)variance parameters are taken to be the mean and (co)variance of the residuals. Only relevant if expert network covariates were supplied to x, otherwise coerced to FALSE.
signif	Significance level for outlier removal. Must be a single number in the interval [0, 1). Corresponds to the percentage of data to be considered extreme and therefore removed (half of signif at each endpoint, on a column-wise basis). The default, 0, corresponds to no outlier removal. Only invoke this argument as an aid to visualisation via <code>plot.Mclust</code> .
...	Further arguments to be passed to other methods.

Details

Of course, the user is always encouraged to use the dedicated `plot` function for objects of the "MoEClust" class instead, but calling `plot` after converting via `as.Mclust` can be particularly useful for univariate mixtures.

In the presence of expert network covariates, the component-specific covariance matrices are modified for plotting purposes via the function `expert_covar`, in order to account for the extra variability of the means, usually resulting in bigger shapes & sizes for the MVN ellipses.

The `signif` argument is intended only to aid visualisation via `plot.Mclust`, as plots therein can be sensitive to outliers, particularly with regard to axis limits. This is especially true when `resid` is TRUE in the presence of expert network covariates.

Value

An object of class "Mclust". See `methods(class="Mclust")` for a list of functions which can be applied to this class.

Note

Of the functions which can be applied to the result of the conversion, `logLik.Mclust` shouldn't be trusted in the presence of either expert network covariates, or (for more models with more than 1 component) gating network covariates.

Mixing proportions are averaged over observations in components in the presence of gating network covariates during the coercion.

Plots may be misleading in the presence of expert covariates when the `what` argument is "density" within `plot.Mclust`.

Also note that plots may be misleading for models of univariate data with more than 1 component, in the presence of expert covariates when `resid` is TRUE and the `what` argument is either "classification" or "uncertainty" within `plot.Mclust`.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

References

C. Fraley and A. E. Raftery (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, 97:611-631.

See Also

[Mclust](#), [plot.Mclust](#), [MoE_clust](#), [plot.MoEClust](#), [expert_covar](#)

Examples

```
## Not run:
# Fit a mixture of experts model to the ais data
data(ais)
mod <- MoE_clust(ais[,3:7], G=3, gating= ~ sex, network.data=ais)

# Convert to the "Mclust" class and examine the classification
plot(as.Mclust(mod), what="classification")

# Examine the density using the augmented data in the expert network
plot(as.Mclust(mod, resid=TRUE), what="density")

# While we could have just used plot.MoEClust above,
# plot.Mclust is especially useful for univariate data
data(CO2data)
res <- MoE_clust(CO2data$CO2, G=2, expert = ~ GNP, network.data=CO2data)
plot(as.Mclust(res))
## End(Not run)
```

CO2data

GNP and CO2 Data Set

Description

This data set gives the gross national product (GNP) per capita in 1996 for various countries as well as their estimated carbon dioxide (CO2) emission per capita for the same year.

Usage

```
data(CO2data)
```

Format

This data frame consists of 28 countries and the following variables:

- GNP - The gross product per capita in 1996.
- CO2 - The estimated carbon dioxide emission per capita in 1996.
- country - An abbreviation pertaining to the country measures (e.g. "GRC" = Greece and "CH" = Switzerland).

References

Hurn, M., Justel, A. and Robert, C. P. (2003) Estimating Mixtures of Regressions, *Journal of Computational and Graphical Statistics*, 12(1): 55-79.

drop_constants	<i>Drop constant variables from a formula</i>
----------------	---

Description

Drops constant variables from the RHS of a formula taking the data set (`dat`), the formula (`formula`), and an optional subset vector (`sub`) as arguments.

Usage

```
drop_constants(dat, formula, sub = NULL)
```

Arguments

<code>dat</code>	A <code>data.frame</code> where rows correspond to observations and columns correspond to variables. Ideally column names should be present.
<code>formula</code>	An object of class " <code>formula</code> ": a symbolic description of the model to be fitted. Variables in the formula not present in the column names of <code>dat</code> will automatically be discarded.
<code>sub</code>	An optional vector specifying a subset of observations to be used in the fitting process.

Value

The updated formula with constant variables removed.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

Examples

```
data(ais)
hema <- as.matrix(ais[,3:7])
sex <- ais$sex
BMI <- ais$BMI
form1 <- as.formula(hema ~ sex + BMI)
sub <- ais$sex == "male"

# Try fitting a linear model
mod1 <- try(lm(form1, data=ais, subset=sub), silent=TRUE)
inherits(mod1, "try-error") # TRUE

# Remove redundant variables from formula & try again
form2 <- drop_constants(ais, form1, sub)
mod2 <- try(lm(form2, data=ais, subset=sub), silent=TRUE)
inherits(mod2, "try-error") # FALSE
```

drop_levels

Drop unused factor levels to predict from unseen data

Description

Drops unseen factor levels in new.data for which predictions are required from a [lm](#) model fit.

Usage

```
drop_levels(fit, newdata)
```

Arguments

`fit` A fitted [lm](#) model.
`newdata` A data.frame containing variables with which to predict.

Value

A data.frame like newdata with unseen factor levels replaced by NA.

Note

This function is untested for models other than [lm](#).

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

Examples

```

data(ais)
hema <- as.matrix(ais[,3:7])
BMI <- ais$BMI
sport <- ais$sport
sub <- ais$sport != "Row"

# Fit a linear model
mod <- lm(hema ~ BMI + sport, data=ais, subset=sub)

# Make predictions
pred1 <- try(predict(mod, newdata=ais), silent=TRUE)
inherits(pred1, "try-error") #TRUE

# Remove unused levels and try again
pred2 <- try(predict(mod, newdata=drop_levels(mod, ais)), silent=TRUE)
inherits(pred2, "try-error") #FALSE
anyNA(pred2) #TRUE

```

expert_covar	<i>Account for extra variability in covariance matrices with expert covariates</i>
--------------	--

Description

In the presence of expert network covariates, this helper function modifies the component-specific covariance matrices of a "MoEClust" object, in order to account for the extra variability of the means, usually resulting in bigger shapes & sizes for the MVN ellipses. The function also works for univariate response data.

Usage

```
expert_covar(x)
```

Arguments

x An object of class "MoEClust" generated by [MoE_clust](#), or an object of class "MoECompare" generated by [MoE_compare](#).

Details

This function is used internally by [plot.MoEClust](#) and [as.Mclust](#), for visualisation purposes.

Value

The variance component only from the parameters list from the output of a call to [MoE_clust](#), modified accordingly.

Note

The `modelName` of the resulting variance object may not correspond to the model name of the "MoEClust" object, in particular scale, shape, &/or orientation may no longer be constrained across clusters.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

See Also

[MoE_clust](#), [MoE_gpairs](#), [plot.MoEClust](#), [as.Mclust](#)

Examples

```
data(ais)
res <- MoE_clust(ais[,3:7], G=2, gating= ~ BMI, expert= ~ sex,
               network.data=ais, modelNames="EVE")

# Extract the variance object
res$parameters$variance

# Modify the variance object
expert_covar(res)
```

MoEClust

MoEClust: Parsimonious Model-Based Clustering with Covariates

Description

Clustering via parsimonious Mixtures of Experts using the *MoEClust* models introduced by Murphy and Murphy (2017) <[arXiv:1711.05632](#)>. This package fits finite Gaussian mixture models with gating and expert network covariates using parsimonious covariance parameterisations from **mclust** via the EM algorithm. Visualisation of the results of such models using generalised pairs plots is also facilitated.

Usage

The most important function in the **MoEClust** package is: `MoE_clust`, for fitting the model via EM with gating and/or expert network covariates, supplied via formula interfaces.

Other functions also exist, e.g. `MoE_control`, `MoE_crit`, `MoE_dens`, `MoE_estep`, `MoE_compare`, and `MoE_aitken`, which are all used within `MoE_clust` but are nonetheless made available for standalone use.

A dedicated plotting function exists for visualising the results using generalised pairs plots, for examining the gating network &/or log-likelihood, and/or graphing model selection criteria values. The generalised pairs plots (`MoE_gpairs`) visualise all pairwise relationships between clustered response variables and associated gating &/or expert network continuous &/or categorical variables,

coloured according to the MAP classification, and also give the marginal distributions of each variable along the diagonal.

An `as.Mclust` method is provided to coerce the output of class "MoEClust" from `MoE_clust` to the "Mclust" class, to facilitate use of plotting and other functions for the "Mclust" class within the `mclust` package. As per `mclust`, **MoEClust** also facilitates modelling with an additional noise component (with or without the mixing proportion for the noise component depending on covariates).

The package also contains two data sets: `ais` and `CO2data`.

Details

- Type: Package
- Package: MoEClust
- Version: 1.1.0
- Date: 2018-02-06 (this version), 2017-11-28 (original release)
- Licence: GPL (>=2)

Author(s)

Keefe Murphy [aut, cre], Thomas Brendan Murphy [ctb]

Maintainer: Keefe Murphy - <<keefe.murphy@ucd.ie>>

References

K. Murphy and T. B. Murphy (2017). Parsimonious Model-Based Clustering with Covariates. *To appear*. <[arXiv:1711.05632](https://arxiv.org/abs/1711.05632)>.

Examples

```
## Not run:
data(ais)
res <- MoE_clust(ais[,3:7], G=2, gating=~BMI, expert=~sex,
                modelNames=c("EVE", "VVE", "VEE"), network.data=ais)
plot(res, what="gpairs")

data(CO2data)
GNP <- CO2data[,1]
CO2 <- CO2data[,2]
m1 <- MoE_clust(CO2, G=1:2)
m2 <- MoE_clust(CO2, G=2, gating= ~ GNP)
m3 <- MoE_clust(CO2, G=1:2, expert= ~ GNP)
m4 <- MoE_clust(CO2, G=2, gating= ~ GNP, expert= ~ GNP)
MoE_compare(m1, m2, m3, m4)
## End(Not run)
```

MoE_aitken	<i>Aitken Acceleration</i>
------------	----------------------------

Description

Calculates the Aitken acceleration estimate of the final converged maximised log-likelihood under the EM framework.

Usage

```
MoE_aitken(loglik)
```

Arguments

loglik	A vector of three consecutive log-likelihood values. These three values should be in ascending order, though this is not checked.
--------	---

Details

The final converged maximised log-likelihood can be used to determine convergence of the EM algorithm within `MoE_clust`, i.e. by checking whether the absolute difference in the current and previous estimates of the final converged maximised log-likelihood is less than some tolerance.

Value

A list with the following components:

ll	The most current estimate for the log-likelihood.
linf	The most current estimate of the final converged maximised log-likelihood.
a	The Aitken acceleration value where $0 \leq a \leq 1$.

Note

Within `MoE_clust`, as specified by the stopping argument of `MoE_control`, "aitken" is the default method used to assess convergence. The other option monitors the "relative" change in log-likelihood against some tolerance. See `MoE_control`.

When the "aitken" method is employed, the final estimate of the log-likelihood is the value of `linf` at convergence, rather than the value of `ll` at convergence under the "relative" option.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

References

Boehning, D., Dietz, E., Schaub, R., Schlattmann, P. and Lindsay, B. G. (1994). The distribution of the likelihood ratio for mixtures of densities from the one-parameter exponential family. *Annals of the Institute of Statistical Mathematics*, 46(2): 373-388.

See Also[MoE_control](#)**Examples**

```
(a1 <- MoE_aitken(-c(449.61534, 442.84221, 436.58999)))
a2 <- MoE_aitken(-c(442.84221, 436.58999, 436.58998))
abs(a2$linf - a1$linf) < 1e-05 #FALSE
a3 <- MoE_aitken(-c(436.58998, 436.58997, 436.58997))
abs(a3$linf - a2$linf) < 1e-05 #TRUE
(l1 <- a3$linf)
```

MoE_clust

*MoEClust: Parsimonious Model-Based Clustering with Covariates***Description**

Fits Mixture of Experts models with **mclust**-family covariance structures. In other words, performs model-based clustering via the EM algorithm where covariates are allowed to enter neither, either, or both the mixing proportions (gating network) and/or component densities (expert network).

Usage

```
MoE_clust(data, G = 1:9, modelNames = NULL, gating = NULL,
  expert = NULL, network.data = NULL, control = MoE_control(...), ...)
```

```
## S3 method for class 'MoEClust'
print(x, digits = 2, ...)
```

```
## S3 method for class 'MoEClust'
summary(object, ...)
```

Arguments

<code>data</code>	A numeric vector, matrix, or data frame of observations. Categorical variables are not allowed. If a matrix or data frame, rows correspond to observations and columns correspond to variables.
<code>G</code>	An integer vector specifying the numbers of mixture components (clusters) to fit. Defaults to $G=1:9$. Must be a strictly positive integer, unless a noise component is included in the estimation, in which case $G=0$ is allowed (see MoE_control).
<code>modelNames</code>	A vector of character strings indicating the models to be fitted in the EM phase of clustering. With n observations and d variables, the defaults are:
for univariate data	<code>c("E", "V")</code>
for multivariate data $n > d$	<code>mclust.options("emModelNames")</code>
for high-dimensional multivariate data $n \leq d$	<code>c("EII", "VII", "EEI", "EVI", "VEI", "VVI")</code>

For single-component models these options reduce to:

for univariate data	"E"
for multivariate data $n > d$	c("EII", "EEI", "EEE")
for high-dimensional multivariate data $n \leq d$	c("EII", "EEI")

For zero-component models with a noise component only the "E" and "EII" models will be fit for univariate and multivariate data, respectively. The help file for [mclustModelNames](#) further describes the available models (though the "X" in the single-component models will be coerced to "E" if supplied that way).

gating	A formula for determining the model matrix for the multinomial logistic regression in the gating network when covariates enter the mixing proportions. This will be ignored where $G=1$. Interactions etc. are permitted. The specification of the LHS of the formula is ignored.
expert	A formula for determining the model matrix for the (multivariate) WLS in the expert network when covariates are included in the component densities. Interactions etc. are permitted. The specification of the LHS of the formula is ignored.
network.data	An optional data frame in which to look for the covariates in the gating &/or expert network formulas, if any. If not found in network.data, any supplied gating &/or expert covariates are taken from the environment from which MoE_clust is called. Try to ensure the names of variables in network.data do not match any of those in data.
control	A list of control parameters for the EM and other aspects of the algorithm. The defaults are set by a call to MoE_control .
...	An alternative means of passing control parameters directly via the named arguments of MoE_control . Do not pass the output from a call to MoE_control here! This argument is only relevant for the MoE_clust function and will be ignored for the associated print and summary functions.
x, object, digits	Arguments required for the print and summary functions: x and object are objects of class "MoEClust" resulting from a call to MoE_clust , while digits gives the number of decimal places to round to for printing purposes (defaults to 2).

Details

The function effectively allows 4 different types of Mixture of Experts model (as well as the different models in the mclust family, for each): i) the standard finite Gaussian mixture, ii) covariates only in the gating network, iii) covariates only in the expert network, iv) the full Mixture of Experts model with covariates entering both the mixing proportions and component densities. Note that having the same covariates in both networks is allowed.

While model selection in terms of choosing the optimal number of components and the **mclust** model type is performed within [MoE_clust](#), using one of the `criterion` options within [MoE_control](#), choosing between multiple fits with different combinations of covariates or different initialisation settings can be done by supplying objects of class "MoEClust" to [MoE_compare](#).

Where BIC, ICL, AIC, DF and `iters` contain NA entries, this corresponds to a model which was not run; for instance a VVV model is never run for single-component models as it is equivalent to EEE. As such, one can consider the value as not really missing, but equivalent to the EEE value. BIC, ICL, AIC and DF all inherit the class "MoECriterion", for which a dedicated print function exists.

Value

A list (of class "MoEClust") with the following named entries, mostly corresponding to the chosen optimal model (as determined by the criterion within `MoE_control`):

<code>call</code>	The matched call.
<code>data</code>	The input data, as a <code>data.frame</code> .
<code>modelName</code>	A character string denoting the mclust model type at which the optimal criterion occurs.
<code>n</code>	The number of observations in the data.
<code>d</code>	The dimension of the data.
<code>G</code>	The optimal number of mixture components.
<code>BIC</code>	A matrix of <i>all</i> BIC values with <code>length{G}</code> rows and <code>length(modelNames)</code> columns. May include missing entries: NA represents models which were not visited, <code>-Inf</code> represents models which were terminated due to error, for which a log-likelihood could not be estimated. Inherits the classes "MoECriterion" and "mclustBIC", for which a dedicated plotting function exists.
<code>ICL</code>	A matrix of <i>all</i> ICL values with <code>length{G}</code> rows and <code>length(modelNames)</code> columns. May include missing entries: NA represents models which were not visited, <code>-Inf</code> represents models which were terminated due to error, for which a log-likelihood could not be estimated. Inherits the classes "MoECriterion" and "mclustBIC", for which a dedicated plotting function exists.
<code>AIC</code>	A matrix of <i>all</i> AIC values with <code>length{G}</code> rows and <code>length(modelNames)</code> columns. May include missing entries: NA represents models which were not visited, <code>-Inf</code> represents models which were terminated due to error, for which a log-likelihood could not be estimated. Inherits the classes "MoECriterion" and "mclustBIC", for which a dedicated plotting function exists.
<code>bic</code>	The BIC value corresponding to the optimal model. May not necessarily be the optimal BIC.
<code>icl</code>	The ICL value corresponding to the optimal model. May not necessarily be the optimal ICL.
<code>aic</code>	The AIC value corresponding to the optimal model. May not necessarily be the optimal AIC.
<code>gating</code>	An object of class "MoE_gating" and either "multinom" or "glm" (for single-component models) giving the <code>multinom</code> regression coefficients of the gating network. If gating covariates were <i>NOT</i> supplied (or the best model has just one component), this corresponds to a RHS of <code>~1</code> , otherwise the supplied gating formula. As such, a fitted gating network is always returned even in the absence of supplied covariates. The number of parameters to penalise by for <code>MoE_crit</code> is given by <code>length(coef(gating))</code> , and the gating formula used is stored here

as an attribute. If there is a noise component, its coefficients are those for the *last* component. **Users are cautioned against making inferences about statistical significance from summaries of the coefficients in the gating network.**

expert	An object of class "MoE_expert" and "lm" giving the (multivariate) WLS regression coefficients of the expert network. If expert covariates were NOT supplied, this corresponds to a RHS of ~1, otherwise the supplied expert formula. As such, a fitted expert network is always returned even in the absence of supplied covariates. The number of parameters to penalise by for <code>MoE_crit</code> is given by $G * \text{length}(\text{coef}(\text{expert}[[1]]))$, and the expert formula used is stored here as an attribute. Users are cautioned against making inferences about statistical significance from summaries of the coefficients in the expert network.
loglik	The vector of increasing log-likelihood values for every EM iteration under the optimal model.
df	The number of estimated parameters in the optimal model (i.e. the number of 'used' degrees of freedom). Subtract this number from n to get the degrees of freedom. The number of parameters due to the gating network, expert network, and covariance matrices are also stored here as attributes of df.
hypvol	The hypervolume parameter for the noise component if required, otherwise set to NA (see <code>MoE_control</code>).
parameters	A list with the following components: <ul style="list-style-type: none"> pro The mixing proportions: either a vector of length G or, if gating covariates were supplied, a matrix with an entry for each observation (rows) and component (columns). mean The means of each component. If there is more than one component, this is a matrix whose <i>k</i>-th column is the mean of the <i>k</i>-th component of the mixture model. For models with expert network covariates, this is given by the posterior mean of the fitted values, otherwise the posterior mean of the response is reported. For models with expert network covariates, the <i>observation-specific</i> means can be accessed by calling <code>predict</code> on each object in the list given by <code>expert</code>. variance A list of variance parameters of each component of the model. The components of this list depend on the model type specification. See the help file for <code>mclustVariance</code> for details. Also see <code>expert_covar</code> for an alternative approach to summarising the variance parameters in the presence of expert covariates. Vinv The inverse of the hypervolume parameter for the noise component if required, otherwise set to NULL (see <code>MoE_control</code>).
z	The final responsibility matrix whose [i,k]-th entry is the probability that observation <i>i</i> belongs to the <i>k</i> -th component. If there is a noise component, its values are found in the <i>last</i> column.
classification	The vector of cluster labels for the chosen model corresponding to z, i.e. <code>max.col(z)</code> . Observations belonging to the noise component will belong to component 0.
uncertainty	The uncertainty associated with the classification.

net.covs	A data frame gathering the unique set of covariates used in the gating and expert networks, if any. Will contain zero columns in the absence of gating or expert network covariates. supplied gating covariates will be excluded if the optimal model has only one component. May have fewer columns than covariates supplied via the <code>network.data</code> argument also, as only the included covariates are gathered here.
resid.data	In the presence of expert network covariates, this is the augmented data (as a data frame) actually used in the clustering at convergence consisting of the $(n * G) * d$ matrix of (multivariate) WLS residuals. Will contain zero columns in the absence of expert network covariates.
DF	A matrix of giving numbers of estimated parameters (i.e. the number of 'used' degrees of freedom) for <i>all</i> visited models, with <code>length{G}</code> rows and <code>length(modelNames)</code> columns. Subtract these numbers from <code>n</code> to get the degrees of freedom. May include missing entries: NA represents models which were not visited, <code>-Inf</code> represents models which were terminated due to error, for which parameters could not be estimated. Inherits the classes "MoECriterion" and "mclustBIC", for which a dedicated plotting function exists.
iters	A matrix giving the total number of EM iterations with <code>length{G}</code> rows and <code>length(modelNames)</code> columns. May include missing entries: NA represents models which were not visited, <code>Inf</code> represents models which were terminated due to singularity/error and thus would never have converged.

Dedicated `plot`, `print` and `summary` functions exist for objects of class "MoEClust". The results can be coerced to the "Mclust" class to access other functions from the `mclust` package via `as.Mclust`.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

References

- K. Murphy and T. B. Murphy (2017). Parsimonious Model-Based Clustering with Covariates. *To appear*. <[arXiv:1711.05632](https://arxiv.org/abs/1711.05632)>.
- C. Fraley and A. E. Raftery (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, 97:611-631.

See Also

`MoE_compare`, `plot.MoEClust`, `MoE_control`, `as.Mclust`, `MoE_crit`, `MoE_estep`, `MoE_dens`, `mclustModelNames`, `mclustVariance`, `expert_covar`

Examples

```
## Not run:
data(ais)
hema <- ais[,3:7]
sex <- ais$sex
BMI <- ais$BMI
```



```

# Fit a standard finite mixture model
m1  <- MoE_clust(hema, G=2:3)

# Allow covariates to enter the mixing proportions
m2  <- MoE_clust(hema, G=2:3, gating= ~ sex + BMI)

# Allow covariates to enter the component densities
m3  <- MoE_clust(hema, G=2:3, expert= ~ sex)

# Allow covariates to enter both the gating & expert network
m4  <- MoE_clust(hema, G=2:3, gating= ~ BMI, expert= ~ sex)

# Extract the model with highest ICL
(comp <- MoE_compare(m1, m2, m3, m4, criterion="icl"))
(best <- comp$optimal)
(summ <- summary(best))

# Examine the gating and expert networks in greater detail
# (but refrain from inferring statistical significance!)
summary(best$gating)
summary(best$expert)

# Visualise the results, incl. the gating network and log-likelihood
plot(best, what="gpairs")
plot(best, what="gating")
plot(best, what="loglik")

# Visualise the results using the 'lattice' library
require("lattice")
z    <- factor(best$classification, labels=paste0("Cluster", seq_len(best$G)))
splom(~ hema | sex, groups=z)
splom(~ hema | z, groups=sex)
## End(Not run)

```

MoE_compare

Choose the best MoEClust model

Description

Takes one or more sets of MoEClust models fitted by `MoE_clust` and ranks them according to the BIC, ICL, or AIC. It's possible to respect the internal ranking within each set of models, or to discard models within each set which were already deemed sub-optimal.

Usage

```
MoE_compare(..., criterion = c("bic", "icl", "aic"), pick = 3L,
  optimal.only = FALSE)
```

```
## S3 method for class 'MoECompare'
print(x, index = seq_len(x$pick), noise = TRUE, ...)
```

Arguments

...	One or more objects of class "MoEClust" outputted by <code>MoE_clust</code> . All models must have been fit to the same data set. A single <i>named</i> list of such objects can also be supplied. This argument is only relevant for the <code>MoE_compare</code> function and will be ignored for the associated <code>print</code> function.
<code>criterion</code>	The criterion used to determine the ranking. Defaults to "bic".
<code>pick</code>	The (integer) number of models to be ranked and compared. Defaults to 3L. Will be constrained by the number of models within the "MoEClust" objects supplied via ... if <code>optimal.only</code> is FALSE, otherwise constrained simply by the number of "MoEClust" objects supplied. Setting <code>pick=Inf</code> is a valid way to select all models.
<code>optimal.only</code>	Logical indicating whether to only rank models already deemed optimal within each "MoEClust" object (TRUE), or to allow models which were deemed suboptimal enter the final ranking (FALSE, the default). See <code>details</code>
<code>x</code>	An object of class "MoECompare" resulting from a call to <code>MoE_compare</code> .
<code>index</code>	A logical or numeric vector giving the indices of the rows of the table of ranked models to print. This defaults to the full set of ranked models. It can be useful when the table of ranked models is large to examine a subset via this <code>index</code> argument, for display purposes.
<code>noise</code>	A logical which determines whether presence of a noise-component should be indicated by the method employed to estimate the hypervolume (defaults to TRUE) or, if FALSE, simply by TRUE. In the absence of a noise component, FALSE will be printed regardless. Only relevant if at least one of the models being compared has a noise component. If any of the compared models do have a noise component, this switch also controls whether the influence (or not) of gating covariates on the noise component's mixing proportion is indicated (either by TRUE or FALSE for models with a noise component, or else a blank entry for those without), for the models among those being compared which have gating covariates.

Details

The purpose of this function is to conduct model selection on "MoEClust" objects, fit to the same data set, with different combinations of gating/expert network covariates or different initialisation settings.

Model selection will have already been performed in terms of choosing the optimal number of components and `mclust` model type within each supplied set of results, but `MoE_compare` will respect the internal ranking of models when producing the final ranking if `optimal.only` is FALSE: otherwise only those models already deemed optimal within each "MoEClust" object will be ranked.

As such if two sets of results are supplied when `optimal.only` is FALSE, the 1st, 2nd and 3rd best models could all belong to the first set of results, meaning a model deemed suboptimal according to one set of covariates could be superior to one deemed optimal under another set of covariates.

Value

A list of class "MoECompare", for which a dedicated print function exists, containing the following elements, each of `length(pick)`:

optimal	The single optimal model (an object of class "MoEClust") among those supplied, according to the chosen criterion.
pick	The final number of ranked models. May be different (i.e. less than) the supplied pick value.
MoENames	The names of the supplied "MoEClust" objects.
modelName	The <code>mclustModelNames</code> .
G	The optimal numbers of components.
bic	The ranked BIC values.
icl	The ranked ICL values.
aic	The ranked AIC values.
gating	The gating formulas.
expert	The expert formulas.
equalPro	Logical indicating whether mixing proportions were constrained to be equal across components.
noise	Either a logical indicating the presence/absence of a noise component, or the type of noise component fitted (if any). Depends on the supplied value of noise. Only displayed if at least one of the compared models has a noise component.
noise.gate	Logical indicating whether gating covariates were allowed to influence the noise component's mixing proportion. Only printed for models with a noise component, when at least one of the compared models has gating covariates, and even then only when noise is supplied as TRUE.

Note

The criterion argument here need not comply with the criterion used for model selection within each "MoEClust" object, but be aware that a mismatch in terms of criterion *may* require the optimal model to be re-fit in order to be extracted, thereby slowing down `MoE_compare`.

A dedicated print function exists for objects of class "MoECompare".

`plot.MoEClust` and `as.Mclust` can both also be called on objects of class "MoECompare".

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

References

K. Murphy and T. B. Murphy (2017). Parsimonious Model-Based Clustering with Covariates. *To appear*. <[arXiv:1711.05632](https://arxiv.org/abs/1711.05632)>.

See Also

`MoE_clust`, `mclustModelNames`, `plot.MoEClust`, `as.Mclust`

Examples

```

data(CO2data)
GNP  <- CO2data[,1]
CO2  <- CO2data[,2]
m1   <- MoE_clust(CO2, G=1:2)
m2   <- MoE_clust(CO2, G=1:2, gating= ~ GNP)
m3   <- MoE_clust(CO2, G=1:2, expert= ~ GNP)
m4   <- MoE_clust(CO2, G=1:2, gating= ~ GNP, expert= ~ GNP)
m5   <- MoE_clust(CO2, G=1:2, equalPro=TRUE)
m6   <- MoE_clust(CO2, G=1:2, expert= ~ GNP, equalPro=TRUE)

# Rank only the optimal models and examine the best model
(comp <- MoE_compare(m1, m2, m3, m4, m5, m6, pick=6, optimal.only=TRUE))
(best <- comp$optimal)
(summ <- summary(best))

# Examine all models visited, including those already deemed suboptimal
# Only print models with expert covariates & more than one component
comp2 <- MoE_compare(m1, m2, m3, m4, m5, m6, pick=18)
print(comp2, comp2$expert != "None" & comp2$G > 1)

```

MoE_control

Set control values for use with MoEClust

Description

Supplies a list of arguments (with defaults) for use with [MoE_clust](#).

Usage

```

MoE_control(criterion = c("bic", "icl", "aic"), stopping = c("aitken",
  "relative"), exp.init = list(...), init.z = c("hc", "quantile", "kmeans",
  "mclust", "random"), eps = .Machine$double.eps, tol = c(1e-05,
  sqrt(.Machine$double.eps), 1e-08), itmax = c(.Machine$integer.max,
  .Machine$integer.max, 100), equalPro = FALSE, warn.it = 0,
  noise.init = NULL, noise.gate = TRUE, noise.meth = c("hypvol",
  "convexhull", "ellipsoidhull"), hc.meth = NULL, hcUse = "VARS",
  verbose = interactive(), ...)

```

Arguments

<p>criterion</p>	<p>When either G or modelNames is a vector, criterion determines whether the "bic" (Bayesian Information Criterion), "icl" (Integrated Complete Likelihood), "aic" (Akaike Information Criterion) is used to determine the 'best' model when gathering output. Note that all criteria will be returned in any case.</p>
<p>stopping</p>	<p>The criterion used to assess convergence of the EM algorithm. The default ("aitken") uses Aitken's acceleration method via MoE_aitken, otherwise the "relative" change in log-likelihood is monitored (which may be less strict).</p>

Both stopping rules are ultimately governed by `tol[1]`. When the "aitken" method is employed, the final estimate of the log-likelihood is the value of `linf` at convergence, rather than the value of `ll` at convergence under the "relative" option.

<code>exp.init</code>	<p>A list supplying select parameters to control the initialisation routine in the presence of <i>expert</i> network covariates (otherwise ignored):</p> <p><code>joint</code> A logical indicating whether the initial partition is obtained on the joint distribution of the response and (continuous only) expert network covariates (defaults to TRUE) or just the response variables (FALSE). Only relevant when <code>init.z</code> is not "random". This may render the "quantile" option to <code>init.z</code> for univariate data unusable.</p> <p><code>mahalanobis</code> A logical indicating whether to iteratively reallocate observations during the initialisation phase to the component corresponding to the expert network regression to which it's closest to the fitted values of in terms of Mahalanobis distance (defaults to TRUE). This will ensure that each component can be well modelled by a single expert prior to running the EM algorithm.</p> <p><code>max.init</code> The maximum number of iterations for the Mahalanobis distance-based reallocation procedure when <code>exp.init\$mahalanobis</code> is TRUE. Defaults to 100.</p>
<code>init.z</code>	<p>The method used to initialise the cluster labels. Defaults to a hierarchical clustering tree as per <code>hc</code> for multivariate data, or quantile-based clustering as per <code>MoE_qc.class</code> for univariate data (unless there are continuous expert network covariates, in which case the defaults is again <code>hc</code>). The "quantile" option is only available for univariate data without continuous expert network covariates. Other options include <code>kmeans</code>, random initialisation, and a full run of <code>Mclust</code>, although this last option is only permitted if there are gating &/or expert covariates within <code>MoE.clust</code>.</p>
<code>eps</code>	<p>A scalar tolerance associated with deciding when to terminate computations due to computational singularity in covariances. Smaller values of <code>eps</code> allow computations to proceed nearer to singularity. The default is the relative machine precision <code>.Machine\$double.eps</code>, which is approximately $2e-16$ on IEEE-compliant machines.</p>
<code>tol</code>	<p>A vector of length three giving relative convergence tolerances for 1) the log-likelihood of the EM algorithm, 2) parameter convergence in the inner loop for models with iterative M-step ("VEI", "EVE", "VEE", "VVE", "VEV"), and 3) optimisation in the multinomial logistic regression in the gating network, respectively. The default is <code>c(1e-05, sqrt(.Machine\$double.eps), 1e-08)</code>. If only one number is supplied, it is used as the tolerance for all three cases given.</p>
<code>itmax</code>	<p>A vector of length three giving integer limits on the number of iterations for 1) the EM algorithm, 2) the inner loop for models with iterative M-step ("VEI", "EVE", "VEE", "VVE", "VEV"), and 3) the multinomial logistic regression in the gating network, respectively.</p> <p>The default is <code>c(.Machine\$integer.max, .Machine\$integer.max, 100)</code> allowing termination to be completely governed by <code>tol</code> for the inner and outer</p>

	loops of the EM. If only one number is supplied, it is used as the iteration limit for the outer loop only.
<code>equalPro</code>	Logical variable indicating whether or not the mixing proportions are to be constrained to be equal in the model. Default: <code>equalPro = FALSE</code> . Only relevant when gating covariates are <i>not</i> supplied within <code>MoE_clust</code> , otherwise ignored. In the presence of a noise component (see below), only the mixing proportions for the non-noise components are constrained to be equal, after accounting for the noise component.
<code>warn.it</code>	A single number giving the iteration count at which a warning will be printed if the EM algorithm has failed to converge. Defaults to 0, i.e. no warning (which is true for any <code>warn.it</code> value less than 3), otherwise the message is printed regardless of the value of <code>verbose</code> . If non-zero, <code>warn.it</code> should be moderately large, but obviously less than <code>itmax[1]</code> . A warning will always be printed if one of more models fail to converge in <code>itmax[1]</code> iterations.
<code>noise.init</code>	A logical or numeric vector indicating an initial guess as to which observations are noise in the data. If numeric, the entries should correspond to row indices of the data. If supplied, a noise term will be added to the model in the estimation.
<code>noise.gate</code>	A logical indicating whether gating network covariates influence the mixing proportion for the noise component, if any. Defaults to <code>TRUE</code> , but leads to greater parsimony if <code>FALSE</code> . Only relevant in the presence of a noise component; only effects estimation in the presence of gating covariates.
<code>noise.meth</code>	The method use to estimate the volume when observations are labelled as noise via <code>noise.init</code> . Defaults to <code>hypvol</code> . For univariate data, this argument is ignored and the range of the data is used instead. The options <code>"convexhull"</code> and <code>"ellipsoidhull"</code> require loading the <code>geometry</code> and <code>cluster</code> libraries, respectively.
<code>hc.meth</code>	A character string indicating the model to be used when hierarchical clustering (see <code>hc</code> is employed for initialisation according to <code>init.z</code> . Defaults to <code>"EII"</code> for high-dimensional data, or <code>"VVV"</code> otherwise.
<code>hcUse</code>	A string specifying the type of input variables to be used when <code>init.z</code> is set to <code>hc</code> . Unlike <code>mclust</code> , this defaults to <code>"VARS"</code> here.
<code>verbose</code>	Logical indicating whether to print messages pertaining to progress to the screen during fitting. By default is <code>TRUE</code> if the session is interactive, and <code>FALSE</code> otherwise. If <code>FALSE</code> , warnings and error messages will still be printed to the screen, but everything else will be suppressed.
<code>...</code>	Catches unused arguments.

Details

`MoE_control` is provided for assigning values and defaults within `MoE_clust`.

While the `criterion` argument controls the choice of the optimal number of components and `mclust` model type, `MoE_compare` is provided for choosing between fits with different combinations of covariates or different initialisation settings.

Value

A named list in which the names are the names of the arguments and the values are the values supplied to the arguments.

Note

When initialising using the "hc" option for the `init.z` argument, the "EII" model is used for high-dimensional data, otherwise "VVV" is used.

Note also that, unlike `mclust`, the `hcUse` argument to `mclust.options` is set to "VARS" by default.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

See Also

[MoE_clust](#), [MoE_aitken](#), [hc](#), [mclust.options](#), [MoE_qclass](#), [hypvol](#), [convhulln](#), [ellipsoidhull](#), [MoE_compare](#)

Examples

```
ctrl <- MoE_control(criterion="icl", itmax=100, warn.it=15, init.z="random")

data(CO2data)
GNP <- CO2data$GNP
res <- MoE_clust(CO2data$CO2, G=2, expert = ~ GNP, control=ctrl)

# Alternatively, specify control arguments directly
res2 <- MoE_clust(CO2data$CO2, G=2, expert = ~ GNP, stopping="relative")

# Supplying ctrl without naming it as control throws an error,
# when any of {modelName, gating, expert} are not supplied
## Not run:
res3 <- MoE_clust(CO2data$CO2, G=2, expert = ~ GNP, ctrl)
## End(Not run)
```

MoE_crit

MoEClust BIC, ICL, and AIC Model-Selection Criteria

Description

Computes the BIC (Bayesian Information Criterion), ICL (Integrated Complete Likelihood), and AIC (Akaike Information Criterion) for parsimonious mixture of experts models given the log-likelihood, the dimension of the data, the number of mixture components in the model, the numbers of parameters in the gating and expert networks respectively, and, for the ICL, the numbers of observations in each component.

Usage

```
MoE_crit(modelName, loglik, n, d, G, gating.pen = G - 1, expert.pen = G * d,
         z = NULL, df = NULL, delta = 0.5)
```

Arguments

modelName	A character string indicating the model. The help file for mclustModelNames describes the available models.
loglik	The log-likelihood for a data set with respect to the Gaussian mixture model specified in the modelName argument.
n, d, G	The number of observations in the data, dimension of the data, and number of components in the Gaussian mixture model, respectively, used to compute loglik. d & G are not necessary if df is supplied.
gating.pen	The number of parameters of the <i>gating</i> network of the MoEClust model. Defaults to $G - 1$, which corresponds to no gating covariates. If covariates are included, this should be the number of regression coefficients in the fitted object. If there are no covariates and mixing proportions are further assumed to be present in equal proportion, gating.pen should be \emptyset . Not necessary if df is supplied.
expert.pen	The number of parameters of the <i>expert</i> network of the MoEClust model. Defaults to $G * d$, which corresponds to no expert covariates. If covariates are included, this should be the number of regression coefficients in the fitted object. Not necessary if df is supplied.
z	The n times G responsibility matrix whose [i,k]-th entry is the probability that observation <i>i</i> belongs to the <i>k</i> -th component.. If supplied the ICL is also computed and returned, otherwise only the BIC and AIC.
df	An alternative way to specify the number of estimated parameters (or 'used' degrees of freedom) exactly. If supplied, the arguments d, G, gating.pen and expert.pen, which are used to calculate the number of parameters, will be ignored. The number of parameters used in the estimation of the noise component, if any, should also be included.
delta	Dirichlet hyperparameter for the prior on the mixing proportions. Defaults to 0.5. Only relevant for the ICL computation.

Details

The function is vectorized with respect to the arguments modelName and loglik.

If model is an object of class "MoEClust" with G components, the number of parameters for the gating.pen and expert.pen are `length(coef(model$gating))` and `G * length(coef(model$expert[[1]]))`, respectively.

Value

A simplified array containing the BIC, AIC, number of estimated parameters (df) and, if z is supplied, also the ICL, for each of the given input arguments.

Note

In order to speed up repeated calls to the function inside `MoE_clust`, no checks take place.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

References

Biernacki, C., Celeux, G., Govaert, G. (2000). Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(7): 719-725.

See Also

`MoE_clust`, `nVarParams`, `mclustModelNames`

Examples

```
MoE_crit(modelName=c("VVI", "VVE", "VVV"), n=120, d=8,
         G=3, loglik=c(-4036.99, -3987.12, -3992.45))

data(CO2data)
GNP <- CO2data$GNP
model <- MoE_clust(CO2data$CO2, G=1:2, expert= ~ GNP)
G <- model$G
name <- model$modelName
ll <- max(model$loglik)
n <- length(CO2data$CO2)
z <- model$z

# Compare BIC from MoE_crit to the BIC of the model
(bic2 <- MoE_crit(modelName=name, loglik=ll, n=n, d=1, G=G, z=z,
                 expert.pen=G * length(coef(model$expert[[1]])))["bic",])
identical(bic2, unname(model$bic)) #TRUE

# Make the same comparison with the known number of estimated parameters
(bic3 <- MoE_crit(modelName=name, loglik=ll, n=n, G=G, df=model$df, z=z)["bic",])
identical(bic3, bic2) #TRUE
```

Description

Computes densities (or log-densities) of observations in MoEClust mixture models.

Usage

```
MoE_dens(modelName, data, mus, sigs, log.tau = 0, Vinv = NULL,
         logarithm = TRUE)
```

Arguments

modelName	A character string indicating the model. The help file for mclustModelNames describes the available models.
data	If there are no expert network covariates, data should be a numeric matrix or data frame, wherein rows correspond to observations (n) and columns correspond to variables (d). If there are expert network covariates, this should be a list of length G containing matrices/data.frames of (multivariate) WLS residuals for each component.
mus	The mean for each of G components. If there is more than one component, this is a matrix whose k-th column is the mean of the k-th component of the mixture model. For the univariate models, this is a G-vector of means. In the presence of expert network covariates, all values should be equal to zero.
sigs	A list of length G of variance parameters of the model. The components of this list depend on the specification of modelName.
log.tau	If covariates enter the gating network, an n times G matrix of mixing proportions, otherwise a G-vector of mixing proportions for the components of the mixture. Must be on the log-scale in both cases. The default of 0 effectively means densities (or log-densities) aren't scaled by the mixing proportions.
Vinv	An estimate of the reciprocal hypervolume of the data region. The default is determined by applying the function hypvol to the data. Used only if an initial guess as to which observations are noise is supplied. Mixing proportion(s) must be included for the noise component also.
logarithm	A logical value indicating whether or not the logarithm of the component densities should be returned. This defaults to TRUE, otherwise component densities are returned, obtained from the component log-densities by exponentiation. The log -densities can be passed to MoE_estep .

Value

A numeric matrix whose [i, k]-th entry is the density or log-density of observation *i* in component *k*, scaled by the mixing proportions.

Note

This function is intended for joint use with [MoE_estep](#), using the **log**-densities.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

See Also

[MoE_estep](#), [MoE_clust](#), [mclustModelNames](#)

Examples

```

data(ais)
hema <- ais[,3:7]
model <- MoE_clust(hema, G=3, gating= ~ BMI + sex, model="EEE", network.data=ais)
Dens <- MoE_dens(modelName=model$modelName, data=hema,
                mus=model$parameters$mean, sigs=model$parameters$variance$sigma,
                log.tau=log(model$parameters$pro))

# Construct the z matrix and compute the log-likelihood
Estep <- MoE_estep(Dens=Dens)
(ll <- Estep$loglik)

# The z matrix will be close but not exactly the same as that from the model
# as the EM algorithm finishes on an M-step, but the classification should be
identical(max.col(Estep$z), as.integer(unname(model$classification))) #TRUE
round(sum(Estep$z - model$z), options()$digits) == 0 #TRUE

```

MoE_estep	<i>Compute the Responsibility Matrix and Log-likelihood for MoEClust Mixture Models</i>
-----------	---

Description

Softmax function to compute the responsibility matrix z and the log-likelihood for MoEClust models, with the aid of [MoE_dens](#).

Usage

```

MoE_estep(modelName, data, mus, sigs, log.tau = 0, Vinv = NULL,
          Dens = NULL)

```

Arguments

modelName	A character string indicating the model. The help file for mclustModelNames describes the available models.
data	If there are no expert network covariates, data should be a numeric matrix or data frame, wherein rows correspond to observations (n) and columns correspond to variables (d). If there are expert network covariates, this should be a list of length G containing matrices/data.frames of (multivariate) WLS residuals for each component.
mus	The mean for each of G components. If there is more than one component, this is a matrix whose k -th column is the mean of the k -th component of the mixture model. For the univariate models, this is a G -vector of means. In the presence of expert network covariates, all values should be equal to zero.
sigs	A list of length G of variance parameters of the model. The components of this list depend on the specification of modelName.

log.tau	If covariates enter the gating network, an n times G matrix of mixing proportions, otherwise a G-vector of mixing proportions for the components of the mixture. Must be on the log-scale in both cases. The default of 0 effectively means densities (or log-densities) aren't scaled by the mixing proportions.
Vinv	An estimate of the reciprocal hypervolume of the data region. The default is determined by applying the function <code>hypvol</code> to the data. Used only if an initial guess as to which observations are noise is supplied. Mixing proportion(s) must be included for the noise component also.
Dens	(Optional) A numeric matrix whose [i,k]-th entry is the log -density of observation <i>i</i> in component <i>k</i> , scaled by the mixing proportions, to which the softmax function is to be applied, typically obtained by <code>MoE_dens</code> but this is not necessary. If this is supplied, all other arguments are ignored, otherwise <code>MoE_dens</code> is called according to the other supplied arguments.

Value

A list containing two elements:

z	A matrix with n rows and G columns containing the probability of cluster membership for each of n observations and G clusters
loglik	The log-likelihood, computed efficiently via <code>rowLogSumExps</code>

Note

This softmax function is intended for joint use with `MoE_dens`, using the **log**-densities.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

See Also

`MoE_dens`, `MoE_clust`, `rowLogSumExps`, `mclustModelNames`

Examples

```
data(ais)
hema <- ais[,3:7]
model <- MoE_clust(hema, G=3, gating= ~ BMI + sex, model="EEE", network.data=ais)
Dens <- MoE_dens(modelName=model$modelName, data=hema,
                mus=model$parameters$mean, sigs=model$parameters$variance$sigma,
                log.tau=log(model$parameters$pro))

# Construct the z matrix and compute the log-likelihood
Estep <- MoE_estep(Dens=Dens)
(ll <- Estep$loglik)

# The z matrix will be close but not exactly the same as that from the model
# as the EM algorithm finishes on an M-step, but the classification should be
identical(max.col(Estep$z), as.integer(unnamed(model$classification))) #TRUE
```

```

round(sum(Estep$z - model$z), options()$digits) == 0           #TRUE

# Call MoE_estep directly
Estep2 <- MoE_estep(modelName=model$modelName, data=hema,
                    mus=model$parameters$mean, sigs=model$parameters$variance$sigma,
                    log.tau=log(model$parameters$pro))
identical(Estep2$loglik, ll)                                   #TRUE

```

MoE_gpairs

*Generalised Pairs Plots for MoEClust Mixture Models***Description**

Produces a matrix of plots showing pairwise relationships between continuous response variables and continuous/categorical associated covariates, as well as the clustering achieved, according to fitted MoEClust mixture models.

Usage

```

MoE_gpairs(res, response.type = c("points", "uncertainty", "density"),
           subset = list(...), residuals = FALSE, scatter.type = c("lm", "points"),
           conditional = c("stripplot", "boxplot"), addEllipses = c("outer", "yes",
           "no", "inner", "both"), border.col = c("purple", "black", "brown", "brown",
           "navy"), bg.col = c("cornsilk", "white", "palegoldenrod", "palegoldenrod",
           "cornsilk"), diagonal = c("on", "off"), outer.margins = list(bottom =
           grid::unit(2, "lines"), left = grid::unit(2, "lines"), top = grid::unit(2,
           "lines"), right = grid::unit(2, "lines")), outer.labels = NULL,
           outer.rot = c(0, 90), gap = 0.05, buffer = 0.02,
           scatter.pars = list(...), density.pars = list(...),
           stripplot.pars = list(...), barcode.pars = list(...),
           mosaic.pars = list(...), axis.pars = list(...), diag.pars = list(...),
           ...)

```

Arguments

- | | |
|---------------|---|
| res | An object of class "MoEClust" generated by MoE_clust , or an object of class "MoECompare" generated by MoE_compare . |
| response.type | The type of plot desired for the scatter plots comparing continuous response variables. Defaults to "points", but points can also be sized according to their associated clustering "uncertainty", or the bivariate "density" contours can be displayed (see density.pars). Caution is advised, until further updates, when interpreting "density" plots in the presence of covariates. |
| subset | A list giving named arguments for producing only a subset of panels:
show.map Logical indicating whether to show panels involving the MAP classification (defaults to TRUE, unless there is only one component, in which case the MAP classification is never plotted.). |

	<p><code>data.ind</code> For subsetting response variables: a vector of column indices corresponding to the variables in the columns of <code>res\$data</code> which should be shown. Defaults to all.</p> <p><code>cov.ind</code> For subsetting covariates: a vector of column indices corresponding to the covariates in the columns <code>res\$net.covs</code> which should be shown. Defaults to all.</p>
<code>residuals</code>	<p>Logical indicating whether to treat the data as the raw data (when FALSE, the default) or the augmented data comprising the residuals from the expert network (when TRUE). In the latter case, the mean and (co)variance parameters are taken to be the mean and (co)variance of the residuals. Only relevant if expert network covariates were supplied under <code>res</code>, otherwise coerced to FALSE.</p>
<code>scatter.type</code>	<p>A vector of length 2 (or 1) giving the plot type for the upper and lower triangular portions of the plot, respectively, pertaining to the associated covariates. Defaults to "lm" for covariate vs. response panels and "points" otherwise. Only relevant for models with continuous covariates in the gating &/or expert network. "ci" and "lm" type plots are only produced for plots pairing covariates with response, and never response vs. response or covariate vs. covariate. Note that lines &/or confidence intervals will only be drawn for continuous covariates included in the expert network; to include covariates included only in the gating network also, use the options "lm2" or "ci2".</p>
<code>conditional</code>	<p>A vector of length 2 (or 1) giving the plot type for the upper and lower triangular portions of the plot, respectively, for plots involving a mix of categorical and continuous variables. Defaults to "stripplot" in the upper triangle and "boxplot" in the lower triangle (see panel.stripplot and panel.bwplot). "barcode" and "violin" plots can also be produced. Only relevant for models with categorical covariates in the gating &/or expert network. Comparisons of two categorical variables (which can only ever be covariates) are always displayed via mosaic plots (see strucplot).</p>
<code>addEllipses</code>	<p>Controls whether to add MVN ellipses with axes corresponding to the within-cluster covariances for the response data ("yes" or "no"). The options "inner" and "outer" (the default) will colour the axes or the perimeter of those ellipses, respectively, according to the cluster they represent (according to <code>scatter.pars\$lc.col</code>). The option "both" will obviously colour both the axes and the perimeter. Ellipses are only ever drawn for multivariate data, and only when <code>response.type</code> is "points" or "uncertainty".</p> <p>Ellipses are centered on the posterior mean of the fitted values when there are expert network covariates, otherwise on the posterior mean of the response variables. In the presence of expert network covariates, the component-specific covariance matrices are also modified for plotting purposes via the function expert_covar, in order to account for the extra variability of the means, usually resulting in bigger shapes & sizes for the MVN ellipses.</p>
<code>border.col</code>	<p>A vector of length 5 (or 1) containing <i>border</i> colours for plots against the MAP classification, response vs. response, covariate vs. response, response vs. covariate, and covariate vs. covariate panels, respectively.</p> <p>Defaults to <code>c("purple", "black", "brown", "brown", "navy")</code>.</p>
<code>bg.col</code>	<p>A vector of length 5 (or 1) containing <i>background</i> colours for plots against the</p>

	MAP classification, response vs. response, covariate vs. response, response vs. covariate, and covariate vs. covariate panels, respectively. Defaults to <code>c("cornsilk", "white", "palegoldenrod", "palegoldenrod", "cornsilk")</code> .
<code>diagonal</code>	By default, the diagonal from the top left to the bottom right is used for displaying the marginal distributions of variables. Specifying "off" will place the diagonal running from the top right down to the bottom left.
<code>outer.margins</code>	A list of length 4 with units as components named bottom, left, top, and right, giving the outer margins; the defaults uses two lines of text. A vector of length 4 with units (ordered properly) will work, as will a vector of length 4 with numeric variables (interpreted as lines).
<code>outer.labels</code>	The default is NULL, for alternating labels around the perimeter. If "all", all labels are printed, and if "none", no labels are printed.
<code>outer.rot</code>	A 2-vector (x, y) rotating the top/bottom outer labels x degrees and the left/right outer labels y degrees. Only works for categorical labels of boxplot and mosaic panels. Defaults to <code>c(0, 90)</code> .
<code>gap</code>	The gap between the tiles; defaulting to 0.05 of the width of a tile.
<code>buffer</code>	The fraction by which to expand the range of quantitative variables to provide plots that will not truncate plotting symbols. Defaults to 2 percent of range currently.
<code>scatter.pars</code>	A list supplying select parameters for the continuous vs. continuous scatter plots. NULL is equivalent to <code>list(scatter.pch=res\$classification, scatter.size=unit(0.25, "char"), scatter.lci.col</code> gives the colour of the fitted lines &/or confidence intervals when <code>scatter.type</code> is one of "ci" or "lm" and the colour of the ellipses when <code>addEllipses</code> is one of "outer", "inner", or "both". Note that <code>scatter.pars\$scatter.size</code> will be modified on an observation by observation level when <code>response.type</code> is "uncertainty".
<code>density.pars</code>	A list supplying select parameters for visualising the bivariate density contours, only when <code>response.type</code> is "density". NULL is equivalent to <code>list(grid.size=c(100, 100), dcol="grey30", nlevels=11, show.labels=T)</code> where <code>density.pars\$grid.size</code> is a vector of length two giving the number of points in the x & y direction of the grid over which the density is evaluated, respectively.
<code>stripplot.pars</code>	A list supplying select parameters for continuous vs. categorical panels when one of the entries of <code>conditional</code> is "stripplot". NULL is equivalent to <code>list(strip.pch=res\$classification, strip.size=unit(0.5, "char"), strip</code>
<code>barcode.pars</code>	A list supplying select parameters for continuous vs. categorical panels when one of the entries of <code>conditional</code> is "boxplot". NULL is equivalent to <code>list(nint=0, ptsize=unit(0.25, "char"), ptpch=1, bcspace=NULL, use.p</code> See the help file for <code>barcode::barcode</code> .
<code>mosaic.pars</code>	A list supplying select parameters for categorical vs. categorical panels. NULL. Currently <code>shade</code> , <code>gp_labels</code> , <code>gp</code> and <code>gp_args</code> are passed through to <code>strucplot</code> for producing mosaic tiles.

<code>axis.pars</code>	A list supplying select parameters for controlling axes. NULL is equivalent to <code>list(n.ticks=5, axis.fontsize=9)</code> . The argument <code>n.ticks</code> will be overwritten for categorical variables with fewer than 5 levels.
<code>diag.pars</code>	A list supplying select parameters for panels along the diagonal. NULL is equivalent to <code>list(diag.fontsize=9, show.hist=TRUE, hist.color=hist.color, show.c)</code> where <code>hist.color</code> is a vector of length 4, giving the colours for the response variables, gating covariates, expert covariates, and covariates entering both networks, respectively. By default, response variables are "black" and covariates of any kind are "grey". The MAP classification is always coloured by cluster membership. <code>show.counts</code> is only relevant for categorical variables.
<code>...</code>	Catches unused arguments. Alternatively, named arguments can be passed directly here to any/all of <code>scatter.pars</code> , <code>barcode.pars</code> , <code>mosaic.pars</code> , <code>axis.pars</code> and <code>diag.pars</code> .

Value

A generalised pairs plot showing all pairwise relationships between clustered response variables and associated gating &/or expert network continuous &/or categorical variables, coloured according to the MAP classification, with the marginal distributions of each variable along the diagonal.

Note

For MoEClust models with more than one associated covariate (entering either network), fitted lines produced in continuous covariate vs. continuous response scatter plots via `scatter.type="lm"` or `scatter.type="ci"` will **NOT** correspond to the coefficients in the expert network (`res$expert`).

`plot.MoEClust` is a wrapper to `MoE_gpairs` which accepts the default arguments, and also produces other types of plots. Caution is advised producing generalised pairs plots when the dimension of the data is large.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

References

K. Murphy and T. B. Murphy (2017). Parsimonious Model-Based Clustering with Covariates. *To appear*. <[arXiv:1711.05632](https://arxiv.org/abs/1711.05632)>.

Emerson, J.W., Green, W.A., Schloerke, B., Crowley, J., Cook, D., Hofmann, H. and Wickham, H. (2013). The Generalized Pairs Plot. *Journal of Computational and Graphical Statistics*, 22(1):79-91.

See Also

`MoE_clust`, `plot.MoEClust`, `panel.stripplot`, `panel.bwplot`, `panel.violin`, `strucplot`

Examples

```
## Not run:
data(ais)
res  <- MoE_clust(ais[,3:7], G=2, gating= ~ BMI, expert= ~ sex,
                 network.data=ais, modelNames="EVE")
MoE_gpairs(res)

# Produce the same plot, but with a violin plot in the lower triangle.
# Add confidence intervals to the scatter plots. Remove the "Fe" variable.
# Size points in the response vs. response panels by their clustering uncertainty.
MoE_gpairs(res, conditional=c("stripplot", "violin"), data.ind=1:4,
           scatter.type=c("ci", "points"), response.type="uncertainty")

# Instead show the bivariate density contours of the response variables.
# Use different colours for histograms of covariates in the gating/expert/both networks.
# Also use different colours for response vs. covariate & covariate vs. response panels.
MoE_gpairs(res, response.type="density", show.labels=FALSE,
           hist.color=c("black", "cyan", "hotpink", "chartreuse"),
           bg.col=c("whitesmoke", "white", "mintcream", "mintcream", "floralwhite"))
## End(Not run)
```

MoE_mahala

*Mahalanobis Distance Outlier Detection for Multivariate Response***Description**

Computes the Mahalanobis distance between the fitted values and residuals of linear regression models with multivariate or univariate responses.

Usage

```
MoE_mahala(fit, resid, squared = FALSE)
```

Arguments

<code>fit</code>	A fitted <code>lm</code> model, inheriting either the <code>"mlm"</code> or <code>"lm"</code> class.
<code>resid</code>	The residuals. Can be residuals for observations included in the model, or residuals arising from predictions on unseen data.
<code>squared</code>	A logical. By default (<code>FALSE</code>), the generalized interpoint distance is computed. Set this flag to <code>TRUE</code> for the squared value.

Value

A vector giving the Mahalanobis distance (or square Mahalanobis distance) between fitted values and residuals for each observation.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

Examples

```
data(ais)
hema <- as.matrix(ais[,3:7])
mod <- lm(hema ~ sex + BMI, data=ais)
res <- hema - predict(mod)
MoE_mahala(mod, res)
```

MoE_plotCrit

Model Selection Criteria Plot for MoEClust Mixture Models

Description

Plots the BIC, ICL, or AIC values of a fitted MoEClust object.

Usage

```
MoE_plotCrit(res, criterion = c("bic", "icl", "aic"), ...)
```

Arguments

res	An object of class "MoEClust" generated by MoE_clust , or an object of class "MoECompare" generated by MoE_compare .
criterion	The criterion to be plotted. Defaults to "bic".
...	Catches other arguments, or additional arguments to be passed to plot.mclustBIC (or equivalent functions for the other criterion arguments).

Value

A plot of the values of the chosen criterion.

Note

[plot.MoEClust](#) is a wrapper to [MoE_plotCrit](#) which accepts the default arguments, and also produces other types of plots.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

See Also

[MoE_clust](#), [plot.MoEClust](#), [plot.mclustBIC](#)

Examples

```
## Not run:
data(ais)
res <- MoE_clust(ais[,3:7], expert= ~ sex, network.data=ais)
MoE_plotCrit(res)
## End(Not run)
```

MoE_plotGate

Plot MoEClust Gating Network

Description

Plots the gating network for fitted MoEClust models, i.e. the observation index against the mixing proportions for that observation, coloured by cluster.

Usage

```
MoE_plotGate(res, type = "l", xlab = "Observation",
             ylab = expression(tau[g]), ylim = c(0, 1), col = NULL, ...)
```

Arguments

`res` An object of class "MoEClust" generated by [MoE_clust](#), or an object of class "MoECompare" generated by [MoE_compare](#).

`type`, `xlab`, `ylab`, `ylim`, `col` These graphical parameters retain their definitions from [matplot](#).

`...` Catches unused arguments, or additional arguments to be passed to [matplot](#).

Value

A plot of the gating network of the fitted model.

Note

[plot.MoEClust](#) is a wrapper to [MoE_plotGate](#) which accepts the default arguments, and also produces other types of plots.

By default, the noise component (if any) will be coloured "grey65".

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

See Also

[MoE_clust](#), [plot.MoEClust](#), [matplot](#)

Examples

```
data(ais)
res <- MoE_clust(ais[,3:7], gating= ~ sex, G=3, modelNames="EEE", network.data=ais,
               noise.init=sample(c(TRUE, FALSE), replace=TRUE, size=nrow(ais)))
MoE_plotGate(res)
```

MoE_plotLogLik

Plot the Log-Likelihood of a MoEClust Mixture Model

Description

Plots the log-likelihood at every iteration of the EM algorithm used to fit a MoEClust mixture model.

Usage

```
MoE_plotLogLik(res, type = "l", xlab = "Iteration",
              ylab = "Log-Likelihood", xaxt = "n", ...)
```

Arguments

`res` An object of class "MoEClust" generated by [MoE_clust](#), or an object of class "MoECompare" generated by [MoE_compare](#).

`type`, `xlab`, `ylab`, `xaxt` These graphical parameters retain their usual definitions from [plot](#).

`...` Catches unused arguments, or additional arguments to be passed to [plot](#).

Value

A plot of the log-likelihood versus the EM iterations.

Note

[plot.MoEClust](#) is a wrapper to [MoE_plotLogLik](#) which accepts the default arguments, and also produces other types of plots.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

See Also

[MoE_clust](#), [plot.MoEClust](#),

Examples

```
data(ais)
res <- MoE_clust(ais[,3:7], gating= ~ BMI, expert= ~ sex,
                G=2, modelNames="EVE", network.data=ais)
MoE_plotLogLik(res)
```

MoE_qclass	<i>Quantile-Based Clustering for Univariate Data</i>
------------	--

Description

Returns a quantile-based clustering for univariate data.

Usage

```
MoE_qclass(x, G)
```

Arguments

x A vector of numeric data.
G The desired number of clusters.

Value

The vector of cluster labels.

Examples

```
data(CO2data)
MoE_qclass(CO2data$CO2, G=2)
```

noise_vol	<i>Approximate Hypervolume Estimate</i>
-----------	---

Description

Computes simple approximations to the hypervolume of univariate and multivariate data sets.

Usage

```
noise_vol(data, method = c("hypvol", "convexhull", "ellipsoidhull"),
          reciprocal = FALSE)
```

Arguments

data	A numeric vector, matrix, or data frame of observations. Categorical variables are not allowed. If a matrix or data frame, rows correspond to observations and columns correspond to variables.
method	The method used to estimate the hypervolume. The "convexhull" and "ellipsoidhull" options require loading the geometry and cluster libraries, respectively.
reciprocal	A logical variable indicating whether or not the reciprocal hypervolume is desired rather than the hypervolume itself. The default is to return the hypervolume.

Value

A hypervolume estimate (or its inverse), to be used as the hypervolume parameter for the noise component when observations are designated as noise in `MoE_clust`.

Author(s)

Keefe Murphy - <<keefe.murphy@ucd.ie>>

See Also

[hypvol](#), [convhulln](#), [ellipsoidhull](#)

Examples

```
data(ais)
noise_vol(ais[,3:7], reciprocal=TRUE)
```

plot.MoEClust

Plot MoEClust Results

Description

Plot results for fitted `MoE_clust` mixture models with gating &/or expert network covariates: generalised pairs plots, model selection criteria, the log-likelihood vs. the EM iterations, and the gating network are all currently visualisable.

Usage

```
## S3 method for class 'MoEClust'
plot(x, what = c("gpairs", "gating", "criterion",
  "loglik"), ...)
```

Arguments

x	An object of class "MoEClust" generated by MoE_clust , or an object of class "MoECompare" generated by MoE_compare .
what	The type of graph requested: <p>gpairs A generalised pairs plot. To further customise this plot, arguments to MoE_gpairs can be supplied.</p> <p>gating The gating network. To further customise this plot, arguments to MoE_plotGate and matplot can be supplied.</p> <p>criterion The model selection criteria. To further customise this plot, arguments to MoE_plotCrit and plot.mclustBIC can be supplied.</p> <p>loglik The log-likelihood vs. the iterations of the EM algorithm. To further customise this plot, arguments to MoE_plotLogLik and plot can be supplied.</p> <p>By default, all of the above graphs are produced.</p>
...	Optional arguments to be passed to MoE_gpairs , MoE_plotGate , MoE_plotCrit , MoE_plotGate , matplot , plot.mclustBIC and plot .

Details

For more flexibility in plotting, use [MoE_gpairs](#), [MoE_plotGate](#), [MoE_plotCrit](#) and [MoE_plotLogLik](#).

Value

The visualisation according to "what" of the results of a fitted MoEClust model.

Note

Caution is advised producing generalised pairs plots when the dimension of the data is large.

Other types of plots are available by first calling [as.Mclust](#) on the fitted object, and then calling [plot.Mclust](#) on the results.

Author(s)

Keefe Murphy - <keefe.murphy@ucd.ie>

References

K. Murphy and T. B. Murphy (2017). Parsimonious Model-Based Clustering with Covariates. *To appear*. <[arXiv:1711.05632](https://arxiv.org/abs/1711.05632)>.

See Also

[MoE_clust](#), [MoE_gpairs](#), [MoE_plotGate](#), [MoE_plotCrit](#), [as.Mclust](#), [plot.Mclust](#)

Examples

```
## Not run:
data(ais)
res <- MoE_clust(ais[,3:7], gating= ~ BMI, expert= ~ sex,
                G=2, modelNames="EVE", network.data=ais)

# Plot the gating network
plot(res, what="gating")

# Plot the log-likelihood
plot(res, what="loglik")

# Produce a generalised pairs plot
plot(res, what="gpairs")

# Modify the gpairs plot by passing arguments to MoE_gpairs()
plot(res, what="gpairs", response.type="density",
      scatter.type="ci", jitter=FALSE, show.counts=FALSE)
## End(Not run)
```


Index

- *Topic **clustering**
 - MoE_clust, 12
 - MoE_compare, 17
 - MoE_crit, 23
 - MoE_dens, 25
 - MoE_estep, 27
 - MoE_qclass, 37
 - *Topic **control**
 - MoE_aitken, 11
 - MoE_control, 20
 - noise_vol, 37
 - *Topic **datasets**
 - ais, 2
 - CO2data, 5
 - *Topic **main**
 - MoE_clust, 12
 - MoE_compare, 17
 - plot.MoEClust, 38
 - *Topic **package**
 - MoEClust, 9
 - *Topic **plotting**
 - MoE_gpairs, 29
 - MoE_plotCrit, 34
 - MoE_plotGate, 35
 - MoE_plotLogLik, 36
 - plot.MoEClust, 38
 - *Topic **utility**
 - as.Mclust, 3
 - drop_constants, 6
 - drop_levels, 7
 - expert_covar, 8
 - MoE_mahala, 33
- ais, 2
- as.Mclust, 3, 4, 8, 9, 16, 19, 39
- CO2data, 5
- convhulln, 23, 38
- drop_constants, 6
- drop_levels, 7
- ellipsoidhull, 23, 38
- expert_covar, 4, 5, 8, 15, 16, 30
- formula, 6
- hc, 21–23
- hypvol, 22, 23, 26, 28, 38
- lm, 7, 33
- logLik.Mclust, 4
- matplot, 35, 39
- Mclust, 3, 5, 21
- mclust.options, 23
- mclustModelNames, 13, 16, 19, 24–28
- mclustVariance, 15, 16
- MoE_aitken, 9, 11, 20, 23
- MoE_clust, 3–5, 8–11, 12, 13, 17–23, 25, 26, 28, 29, 32, 34–36, 38, 39
- MoE_compare, 4, 8, 9, 13, 16, 17, 18, 19, 22, 23, 29, 34–36, 39
- MoE_control, 9, 11–16, 20, 22
- MoE_crit, 9, 14–16, 23
- MoE_dens, 9, 16, 25, 27, 28
- MoE_estep, 9, 16, 26, 27
- MoE_gpairs, 9, 29, 32, 39
- MoE_mahala, 33
- MoE_plotCrit, 34, 34, 39
- MoE_plotGate, 35, 35, 39
- MoE_plotLogLik, 36, 36, 39
- MoE_qclass, 21, 23, 37
- MoEClust, 9
- MoEClust-package (MoEClust), 9
- multinom, 14
- noise_vol, 37
- nVarParams, 25
- panel.bwplot, 30, 32

panel.stripplot, [30](#), [32](#)
panel.violin, [32](#)
plot, [4](#), [16](#), [36](#), [39](#)
plot.Mclust, [4](#), [5](#), [39](#)
plot.mclustBIC, [34](#), [39](#)
plot.MoEClust, [5](#), [8](#), [9](#), [16](#), [19](#), [32](#), [34–36](#), [38](#)
print.MoEClust (MoE_clust), [12](#)
print.MoECompare (MoE_compare), [17](#)

rowLogSumExps, [28](#)

strucplot, [30–32](#)
summary.MoEClust (MoE_clust), [12](#)