

Package ‘boral’

October 10, 2016

Title Bayesian Ordination and Regression AnaLysis

Version 1.1.1

Date 2016-10-12

Author Francis K.C. Hui

Maintainer Francis Hui <fhui28@gmail.com>

Description Bayesian approaches for analyzing multivariate data in ecology. Estimation is performed using Markov Chain Monte Carlo (MCMC) methods via JAGS. Three types of models may be fitted: 1) With explanatory variables only, boral fits independent column GLMs to each column of the response matrix; 2) With latent variables only, boral fits a purely latent variable model for model-based unconstrained ordination; 3) With explanatory and latent variables, boral fits correlated column GLMs with latent variables to account for any residual correlation between the columns of the response matrix.

License GPL-2

Depends coda

Imports R2jags, mvtnorm, fishMod, MASS, stats, graphics, grDevices

Suggests mvabund (>= 3.8.4), corrplot

NeedsCompilation no

Repository CRAN

Date/Publication 2016-10-10 13:51:41

R topics documented:

boral-package	2
boral	3
calc.condlogLik	19
calc.logLik.lv0	22
calc.marglogLik	26
coefspplot	30
create.life	31
ds.residuals	37
fitted.boral	38
get.dic	39

get.enviro.cor	41
get.hpdiintervals	42
get.measures	45
get.more.measures	49
get.residual.cor	52
lvplot	55
make.jagsboralmmodel	57
make.jagsboralnullmodel	62
plot.boral	66
summary.boral	68

Index	70
--------------	-----------

boral-package	<i>Bayesian Ordination and Regression AnaLysis (boral)</i>
---------------	--

Description

boral is a package offering Bayesian model-based approaches for analyzing multivariate data in ecology. Estimation is performed using Bayesian/Markov Chain Monte Carlo (MCMC) methods via JAGS (Plummer, 2003). Three “types” of models may be fitted: 1) With covariates and no latent variables, boral fits independent response GLMs such that the columns of y are assumed to be independent; 2) With no covariates, boral fits a pure latent variable model (Skrondal and Rabe-Hesketh, 2004) to perform model-based unconstrained ordination (Hui et al., 2014); 3) With covariates and latent variables, boral fits correlated response GLMs, with latent variables accounting for any residual correlation between the columns of y (Warton et al., 2015).

Details

Package: boral
 Type: Package
 Version: 0.6
 Date: 2014-12-12
 License: GPL-2

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Hui et al. (2014). Model-based approaches to unconstrained ordination. *Methods in Ecology and Evolution*, 6, 399-411.

- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In Proceedings of the 3rd International Workshop on Distributed Statistical Computing. March (pp. 20-22).
- Skrondal, A., and Rabe-Hesketh, S. (2004). Generalized latent variable modeling: Multilevel, longitudinal, and structural equation models. CRC Press.
- Warton et al. (2015). So Many Variables: Joint Modeling in Community Ecology. Trends in Ecology and Evolution, to appear.
- Yi W. et al. (2013). mvabund: statistical methods for analysing multivariate abundance data. R package version 3.8.4.

Examples

```
## Please see examples in the help file for boral (?boral). Thanks!
```

boral	<i>Fitting boral (Bayesian Ordination and Regression AnaLysis) models</i>
-------	---

Description

Bayesian ordination and regression models for analyzing multivariate data in ecology. Three "types" of models may be fitted: 1) With covariates and no latent variables, boral fits independent response GLMs; 2) With no covariates, boral fits a pure latent variable model; 3) With covariates and latent variables, boral fits correlated response GLMs.

Usage

```
boral(y, ...)

## Default S3 method:
boral(y, X = NULL, traits = NULL, which.traits = NULL,
      family, trial.size = 1, num.lv = 0, row.eff = "none", row.ids = NULL,
      save.model = FALSE, calc.ics = TRUE, mcmc.control = list(n.burnin = 10000,
      n.iteration = 40000, n.thin = 30, seed = 123),
      prior.control = list(type = c("normal", "normal", "normal", "uniform"),
      hypparams = c(100, 20, 100, 50), ssvs.index = -1, ssvs.g = 1e-6),
      do.fit = TRUE, model.name = NULL, ...)

## S3 method for class 'boral'
print(x, ...)
```

Arguments

<code>y</code>	A response matrix of multivariate data e.g., counts, binomial or Bernoulli responses, continuous response, and so on. With multivariate abundance data ecology for instance, rows correspond to sites and columns correspond to species. Any categorical (multinomial) responses must be converted to integer values. For ordinal data, the minimum level of <code>y</code> must be 1 instead of 0.
----------------	--

<code>X</code>	A model matrix of covariates, which can be included as part of the boral model. Defaults to NULL, in which case no model matrix was used. No intercept column should be included in <code>X</code> .
<code>x</code>	An object for class "boral".
<code>traits</code>	A model matrix of species covariates, which can be included as part of the boral model. Defaults to NULL, in which case no matrix was used. No intercept column should be included in <code>traits</code> , as it is included automatically.
<code>which.traits</code>	<p>A list of length equal to (number of columns in <code>X</code> + 1), informing which columns of <code>traits</code> the column-specific intercepts and each of the column-specific regression coefficients should be regressed against. The first element in the list applies to the column-specific intercept, while the remaining elements apply to the regression coefficients. Each element of <code>which.traits</code> is a vector indicating which traits are to be used.</p> <p>For example, if <code>which.traits[[2]] = c(2, 3)</code>, then the regression coefficients corresponding to the first column in <code>X</code> are regressed against the second and third columns of <code>traits</code>. If <code>which.traits[[2]] = 0</code>, then the regression coefficients for each column are treated as independent. Please see help file below for more details.</p> <p>Defaults to NULL, in conjunction with <code>traits = NULL</code>.</p>
<code>family</code>	<p>Either a single element, or a vector of length equal to the number of columns in <code>y</code>. The former assumes all columns of <code>y</code> come from this distribution. The latter option allows for different distributions for each column of <code>y</code>. Elements can be one of "binomial" (with probit link), "poisson" (with log link), "negative.binomial" (with log link), "normal" (with identity link), "lnormal" for log-normal (with log link), "tweedie" (with log link), "exponential" (with log link), "gamma" (with log link), "beta" (with logit link), "ordinal" (cumulative probit regression).</p> <p>For the negative binomial distribution, the variance is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the column-specific dispersion parameter. For the normal distribution, the variance is parameterized as $Var(y) = \phi^2$, where ϕ is the column-specific standard deviation. For the tweedie distribution, the variance is parameterized as $Var(y) = \phi\mu^p$ where ϕ is the column-specific dispersion parameter and p is a power parameter common to all columns assumed to be tweedie, with $1 < p < 2$. For the gamma distribution, the variance is parameterized as $Var(y) = \mu/\phi$ where ϕ is the column-specific rate (henceforth referred to also as dispersion parameter). For the beta distribution, the parameterization is in terms of the mean μ and sample size ϕ (henceforth referred to also as dispersion parameter), so that the two shape parameters are given by $a = \mu\phi$ and $b = (1 - \mu)\phi$.</p> <p>All columns assumed to have ordinal responses are constrained to have the same cutoffs points, with a column-specific random intercept to account for differences between the columns (please see <i>Details</i> for formulation).</p>
<code>trial.size</code>	Either equal to a single element, or a vector of length equal to the number of columns in <code>y</code> . If a single element, then all columns assumed to be binomially distributed will have trial size set to this. If a vector, different trial sizes are allowed in each column of <code>y</code> . The argument is ignored for all columns not assumed to be binomially distributed. Defaults to 1, i.e. Bernoulli distribution.

num.lv	Number of latent variables to fit. Can take any non-negative integer value. Defaults to 0.
row.eff	Single element indicating whether (multiple) row effects are included as fixed effects ("fixed"), random effects ("random") or not included ("none") in the boral model. If random effects, they are drawn from a normal distribution with mean zero and unknown variance, analogous to a random intercept in mixed models. Defaults to "none".
row.ids	A matrix with the number of rows equal to the number of rows in y , and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row i in y for random effect eqnj. This matrix is useful if one wants to specify more complicated row effect structures beyond a single, row effect unique to each row; please see details below as well as examples below. Whether these row effects are included as fixed or random effects is governed by row.eff. Defaults to NULL, so that if row.eff = "none" then the argument is ignored, otherwise if row.eff = "fixed" or "random", then row.ids = matrix(1:nrow(y), ncol = 1) i.e., a single, row effect unique to each row.
save.model	If save.model = TRUE, then the JAGS model file is saved as a text file (with name given by model.name) in the current working directory, as well as the MCMC samples from the call to JAGS. If saved, various functions available in the coda package can be applied to the MCMC samples. Note MCMC samples can take up a lot of memory. Defaults to FALSE.
calc.ics	If calc.ics = TRUE, then various information criteria values are also returned, which could be used to perform model selection (see get.measures). Defaults to TRUE.
mcmc.control	A list of parameters for controlling the MCMC sampling. Values not set will assume default values. These include: <ul style="list-style-type: none"> • <i>n.burnin</i>: Length of burnin i.e., the number of iterations to discard at the beginning of the MCMC sampler. Defaults to 10000. • <i>n.iteration</i>: Number of iterations including burnin. Defaults to 40000. • <i>n.thin</i>: Thinning rate. Must be a positive integer. Defaults to 30. • <i>seed</i>: Seed for JAGS sampler. A set.seed(seed) command is run immediately before starting the MCMC sampler. Defaults to the value 123.
prior.control	A list of parameters for controlling the prior distributions. Values not set will assume default values. These include: <ul style="list-style-type: none"> • <i>type</i>: Vector of four strings indicating the type of prior distributions to use. In order, these are: 1) priors for all column-specific intercepts, row effects, and cutoff points for ordinal data; 2) priors for the latent variable coefficients (ignored if num.lv = 0); 3) priors for all column-specific coefficients relating to the model matrix X (ignored if $X = \text{NULL}$). When traits are included in the model, this is also the prior for the trait regression coefficients (please see section <i>Including species traits</i> in the help file for boral for more information); 4) priors for any dispersion parameters. For elements 1-3, the prior distributions currently available include: I) "normal", which is normal prior with the variance controlled by the hypparams argument; II) "cauchy", which is a Cauchy prior with variance controlled

by the `hypparams` argument. Gelman, et al. (2008) considers using Cauchy priors with variance 2.5^2 as weakly informative priors for regression coefficients; III) "uniform", which is uniform prior with minimum values given by `-hypparams` and maximum values given by `+hypparams`.

For element 4, the prior distributions currently available include: I) "uniform", which is uniform prior with minimum zero and maximum controlled by `hypparams[4]`; II) "halfnormal", which is half-normal prior with variance controlled by `hypparams[4]`; III) "halfcauchy", which is a half-Cauchy prior with variance controlled by the `hypparams[4]` argument. Defaults to the vector `c("normal", "normal", "normal", "uniform")`.

- `hypparams`: Vector of four hyperparameters used in the set up of prior distributions. In order, these are: 1) affects the prior distribution for all column-specific intercepts, row effects, and cutoff points for ordinal data. If `row.eff = "random"`, this also controls the maximum of the uniform prior for the standard deviation of the random effects normal distribution. Also, if more than two of the columns are ordinal, then this also controls the maximum of the uniform prior for the standard deviation of the column-specific random intercepts for these columns; 2) affects the prior distribution for all latent variable coefficients (ignored if `num.lv = 0`); 3) affects the prior distribution for column-specific coefficients relating to the model matrix `X` (ignored if `X = NULL`). When traits are included in the model, it also affects the prior distribution for the trait regression coefficients, and controls the maximum of the uniform prior for the standard deviation of the normally distributed random effects; 4) affects the prior distribution used for any dispersion parameters.

Defaults to the vector `c(100, 20, 100, 50)`.

- `ssvs.index`: Indices to be used for stochastic search variable selection (SSVS, George and McCulloch, 1993). Either a single element or a vector with length equal to the number of columns in the implied model matrix `X`. Each element can take values of -1 (no SSVS is performed on this covariate), 0 (SSVS is performed on individual coefficients for this covariate), or any integer exceeding 1 (SSVS is performed on collectively all coefficients on this covariate/s.)

This argument is only read if `X.eff = TRUE`. Please see the [boral](#) help file for more information regarding the implementation of SSVS. Defaults to -1, in which case no model selection is performed on the fitted model at all.

- `ssvs.g`: Multiplicative, shrinkage factor for SSVS, which controls the strength of the "spike" in the SSVS mixture prior. In summary, if the coefficient is included in the model, the "slab" prior is a normal distribution with mean zero and variance given by `hypparams[3]`, while if the coefficient is not included in the model, the "spike" prior is normal distribution with mean zero and variance given by `hypparams[3]*ssvs.g`. Please see the [boral](#) help file for more information regarding the implementation of SSVS. Defaults to $1e-6$.

`do.fit` If `do.fit = FALSE`, then only the JAGS model file is written to the current working directory (as text file with name based on `model.name`). No MCMC sampling is performed, and *nothing else* is returned. Defaults to TRUE.

`model.name` Name of the text file that the JAGS model is written to. Defaults to NULL, in

which case the default of "jagsboralmode.txt" is used.

... Not used.

Details

The boral package is designed to fit three types models which may be useful in ecology (and probably outside of ecology as well =D).

Independent response models: boral allows explanatory variables to be entered into the model via a model matrix X . This model matrix can contain anything the user wants, provided factors have been parameterized as dummy variables. It should NOT include an intercept column.

Without latent variables, i.e. $\text{num.lv} = 0$, boral fits separate GLMs to each column of the $n \times p$ matrix y , where the columns are assumed to be independent.

$$g(\mu_{ij}) = \alpha_i + \beta_{0j} + \mathbf{x}_i^T \boldsymbol{\beta}_j; \quad i = 1, \dots, n; j = 1, \dots, p,$$

where the mean response for element (i,j) , denoted as mu_{ij} , is regressed against the covariates \mathbf{x}_i via a link function $g(\cdot)$. The quantities β_{0j} and $\boldsymbol{\beta}_j$ denote the column-specific intercepts and coefficients respectively, while α_i is an optional row effect that may be treated as a fixed or random effect. The latter assumes the row effects are drawn from a normal distribution with unknown variance ϕ^2 . One reason we might want to include row effects is to account differences in sampling intensity between sites: these can lead to differences in site total abundance, and so by including fixed effects they play the same role as an offset to account for these differences.

Note boral can also handle multiple, hierarchical row effects, which may be useful to account for sampling design. This is controlled using the `row.ids` argument. For example, if the first five rows of y correspond to replications from site 1, the next five rows correspond to replications from site 2, and so on, then one can set `row.ids = matrix(c(1,1,1,1,1,2,2,2,2,2,...), ncol = 1)` to take this in account. While this way of coding row effects via the `row.ids` argument takes some getting used to, it has been done this way partly to force the user to think more carefully about exactly the structure of the data i.e., with great power comes great responsibility...

Without row effects, the above independent response model is basically a Bayesian analog of the `manyglm` function in the `mvabund` package (Wang et al., 2013). Unlike `manyglm` though, row effects can be added easily as a type of "row-standardization". Also, a wider range of assumed distributions (families) are possible, as discussed below (please see the section later on distributions.)

Pure latent variable models: If no explanatory variables are included and $\text{num.lv} > 0$, boral fits a pure latent variable model to perform model-based unconstrained ordination (Hui et al., 2014),

$$g(\mu_{ij}) = \alpha_i + \beta_{0j} + \mathbf{z}_i^T \boldsymbol{\theta}_j,$$

where instead of measured covariates, we now have a vector of latent variables \mathbf{z}_i with $\boldsymbol{\theta}_j$ being the column-specific coefficients relating to these latent variables. The column-specific intercept, β_{0j} , accounts for differences between species prevalence, while the row effect, α_i , is included to account for differences in site total abundance (typically assuming a fixed effect, `row.eff = "fixed"`, although see Jamil and ter Braak, 2013, for a motivation for using random site effects), so that the ordination is then in terms of species composition. If α_i is omitted from the model i.e., `row.eff = FALSE`, then the ordination will be in terms of relative species abundance.

As mentioned previously, one reason for including fixed row effects is to account of any potential differences in sampling intensity between sites.

Unconstrained ordination is used for visualizing multivariate data in a low-dimensional space, without reference to covariates (Chapter 9, Legendre and Legendre, 2012). Typically, `num.lv = 1` to 3 latent variables is used, allowing the latent variables to be plotted (using `lvplot`, for instance). The resulting plot can be interpreted in the same manner as plots from Nonmetric Multi-dimensional Scaling (NMDS, Kruskal, 1964) and Correspondence Analysis (CA, Hill, 1974), for example. A biplot can also be constructed by setting `biplot = TRUE` when using `lvplot`, so that both the latent variables and their corresponding coefficients are plotted. For instance, with multivariate abundance data, biplots are used to visualize the relationships between sites in terms of species abundance or composition, as well as the indicator species for the sites.

Correlated response models: When one or more latent variables are included in conjunction with covariates i.e., X is given and `num.lv > 1`, `boral` fits separate GLMs to each column of y while allowing for residual correlation between columns via the latent variables. This is quite useful for multivariate abundance data in ecology, where a separate GLM is fitted to species (modeling its response against environmental covariates), while accounting for the fact species at a site are likely to be correlated for reason other than similarities in environmental responses, e.g. biotic interaction, phylogeny, and so on. Correlated response models take the following form,

$$g(\mu_{ij}) = \alpha_i + \beta_{0j} + \mathbf{x}_i^T \boldsymbol{\beta}_j + \mathbf{z}_i^T \boldsymbol{\theta}_j.$$

This model is thus a mash of the first two types of models. The linear predictor $\mathbf{z}_i^T \boldsymbol{\theta}_j$ induces a residual covariance between the columns of y (which is of rank `num.lv`). For multivariate abundance data, this leads to a parsimonious method of accounting for correlation between species not due to the shared environmental responses. After fitting the model, the residual correlation matrix then can be obtained via the `get.residual.cor` function. Note `num.lv > 1` is necessarily in order to flexibly model the residual correlations; see Pollock et al. (2014) for residual correlation matrices in the context of Joint Species Distribution Models, and Warton et al. (2015) for an overview of latent variable models in multivariate ecology.

Including species traits: When covariates X are included (i.e. both the independent and correlated response models), one has the option of also including traits to help explain differences in species environmental responses to these covariates. Specifically, when `traits` and `which.traits` are supplied, then the β_{0j} 's and $\boldsymbol{\beta}_j$'s are then regarded as random effects drawn from a normal distribution. For the species-specific intercepts, we have

$$\beta_{0j} \sim N(\kappa_{01} + \mathbf{traits}_j^T \boldsymbol{\kappa}_1, \sigma_1^2),$$

where $(\kappa_{01}, \boldsymbol{\kappa}_1)$ are the regression coefficients relating to the traits to the intercepts and σ_1 is the error standard deviation. These are now the "parameters" in the model, in the sense that priors are assigned to them and MCMC sampling is used to estimate them (see the next section on estimation).

In an analogous manner, each of the elements in $\boldsymbol{\beta}_j = (\beta_{j1}, \dots, \beta_{jd})$ are now drawn as random effects from a normal distribution. That is, for $k = 1, \dots, d$ where $d = \text{ncol}(X)$, we have,

$$\beta_{jk} \sim N(\kappa_{0k} + \mathbf{traits}_j^T \boldsymbol{\kappa}_k, \sigma_k^2),$$

Which traits are to be included (regressed) in the mean of the normal distributions is determined by the list `which.traits`. The first element in the list applies to `beta0j`, while the remaining elements

apply to the the β_j . Each element of which.traits is a vector indicating which traits are to be used. For example, if which.traits[[2]] = c(2, 3), then the β_{j1} 's are drawn from a normal distribution with mean depending only on the second and third columns of traits. If which.traits[[2]] = 0, then the regression coefficients are treated as independent, i.e. the values of β_{j1} are given their own priors and estimated separately from each other.

Including species traits in the model can be regarded as a method of simplifying the model – rather than each to estimates p sets of species-specific coefficients, we instead say that these coefficients are linearly related to the corresponding values of their traits (Warton et al., 2015). That is, we are using trait data to help explain similarities/differences in species responses to the environment. This idea has close relations to the fourth corner problem in ecology (Brown et al., 2014). Unlike the models of Brown et al. (2014) however, which treat the β_{0j} 's and β_{jk} 's are fixed effects and fully explained by the traits, boral adopts a random effects approach similar to Jamil et al. (2013) to "soak up" any additional between species differences in environmental responses not explained by traits.

Estimation: For boral models, estimation is performed using Bayesian Markov Chain Monte Carlo (MCMC) methods via JAGS (Plummer, 2003). Please note that only *one* MCMC chain in run – this point is discussed later in this help file. Regarding prior distributions, the default settings, based on the prior.control argument, are as follows:

- Normal priors with mean zero and variance given by hypparams[1] are assigned to all intercepts, cutoffs for ordinal responses, and row effects. If the row effects are assumed to random, then the standard deviation of the normal random effect is assigned a uniform prior with maximum hypparams[1]. If more than two columns are ordinal responses, then the standard deviation of the normal random species-specific intercepts is assigned a uniform prior with maximum hypparams[1].
- Normal priors with mean zero and variance given by hypparams[2] are assigned coefficients relating to latent variables, θ_j .
- Normal priors with mean zero and variance given by hypparams[3] are assigned to all coefficients relating to covariates in β_j . If traits are included, the same normal priors are assigned to the κ 's, and the standard deviation σ_k are assigned uniform priors with maximum equal to hypparams[4].
- For the negative binomial, normal, lognormal, and tweedie distributions, uniform priors with maximum equal to hypparams[4] are used on the dispersion parameters. Please note that for the normal and lognormal distributions, these uniform priors are assigned to the standard deviations ϕ (see Gelman, 2006).

With the default values of hypparams, all parameters are given uninformative prior distributions except for the priors of the latent variable coefficients θ_j . We recommend such a "weakly-informative" prior for the latent variable coefficients, as this tends to be produce more stable MCMC sampling particularly if the response matrix is large and sparse.

Using information criteria at your own risk: Using information criterion from calc.ics = TRUE for model selection should be done with extreme caution, for two reasons: 1) The implementation of some of these criteria is heuristic and experimental, 2) Deciding what model to fit should also be driven by the science. For example, it may be the case that a criterion suggests a model with 3 or 4 latent variables is more appropriate. However, if we are interested in visualizing the data for ordination purposes, then models with 1 or 2 latent variables are more appropriate. As another

example, whether or not we include row effects when ordinating multivariate abundance data depends on if we are interested in differences between sites in terms of relative species abundance (`row.eff = "none"`) or species composition (`row.eff = "fixed"`).

We also make the important point that if traits are included in the model, then the regression coefficients β_{0j}, β_j are now random effects. However, currently the calculation of all information criteria do not take this into account!

SSVS: As an alternative to using information criterion for model selection, stochastic search variable selection (SSVS, George and McCulloch, 1993) is also implemented for the column-specific coefficients β_j . Basically, SSVS works by placing a spike-and-slab priors on these coefficients, such that the spike is a narrow normal distribution concentrated around zero and the slab is a normal distribution with a large variance.

$$\rho(\beta) = I_{\beta=1} \times \mathcal{N}(0, \sigma^2) + (1 - I_{\beta=1}) \times \mathcal{N}(0, g * \sigma^2),$$

where σ^2 is determined by `prior.control$hypparams[3]` (see section on estimation above), g is determined by `ssvs.g`, and $I_{\beta=1} = P(\beta = 1)$ is an indicator function representing whether coefficient is included in the model. It is given a Bernoulli prior with probability of inclusion 0.5. After fitting, the posterior probability of β being included in the model is returned based on posterior mean of the indicator function $I_{\beta=1}$. Note this is NOT the same as a p -value seen in maximum likelihood estimation – a p -value provides an indication of how much evidence there is against the null hypothesis of $\beta = 0$, while the posterior probability provides a measure of how likely it is for $\beta \neq 0$ given the data.

In `boral`, SSVS can be applied at a grouped or individual coefficient level, and this is governed by `prior.control$ssvs.index`. For elements of `ssvs.index` equal to -1, SSVS is not applied on the corresponding covariate of the model matrix X . For elements equal to 0, SSVS is applied to each individual coefficient of the corresponding covariate in X . That is, the fitted model will return posterior probabilities for this covariate, one for each column of y . For elements taking positive integers 1,2,..., SSVS is applied to each group of coefficients of the corresponding covariate in X . That is, the fitted model will return a single posterior probability for this covariate, indicating whether this covariate should be included for all columns of y ; see O'Hara and Sillanpaa (2009) and Tenan et al. (2014) among many others for an discussion of Bayesian variable selection methods.

Note the last application of SSVS allows multiple covariates to be tested *simultaneously*. For example, suppose X consists of five columns – the first two columns are environmental covariates, while the last three correspond to quadratic terms of the two covariates as well as their interaction. If we want to "test" whether any quadratic terms are required, then we can set `prior.control$ssvs.index = c(-1, -1, 1, 1, 1)`, so a single posterior probability of inclusion is returned for the last three columns of X .

Finally, note using information criterion (and possibly residual analysis) should probably not be done at the same as when SSVS is used, and it is advised to separate out their applications e.g., choose the explanatory variables first using SSVS, and then use information criterion to select the number of latent variables???. Obtaining summaries such as posterior medians and HPD intervals of the coefficients from a `boral` model that is implementing SSVS is also (perhaps) problematic, because the posterior distribution is multi-modal.

Value

An object of class "boral" is returned, being a list containing the following components where applicable:

<code>call</code>	The matched call.
<code>lv.coefs.mean/median/sd/iqr</code>	Matrices containing the mean/median/standard deviation/interquartile range of the posterior distributions of the latent variable coefficients. This also includes the column-specific intercepts, and dispersion parameters if appropriate.
<code>lv.mean/median/sd/iqr</code>	A matrix containing the mean/median/standard deviation/interquartile range of the posterior distributions of the latent variables.
<code>X.coefs.mean/median/sd/iqr</code>	Matrices containing the mean/median/standard deviation/interquartile range of the posterior distributions of the column-specific coefficients relating to the model matrix X .
<code>traits.coefs.mean/median/sd/iqr</code>	Matrices containing the mean/median/standard deviation/interquartile range of the posterior distributions of the coefficients and standard deviation relating to the species traits (please see the section on including traits above).
<code>cutoffs.mean/median/sd/iqr</code>	Vectors containing the mean/median/standard deviation/interquartile range of the posterior distributions of the common cutoffs for ordinal responses (please see the not-so-brief tangent on distributions above).
<code>ordinal.sigma.mean/median/sd/iqr</code>	Scalars containing the mean/median/standard deviation/interquartile range of the posterior distributions of the standard deviation for the random intercept normal distribution corresponding to the ordinal response columns.
<code>powerparam.mean/median/sd/iqr</code>	Scalars for the mean/median/standard deviation/interquartile range of the posterior distributions of the common power parameter for tweedie responses (please see the not-so-brief tangent on distributions above).
<code>row.coefs.mean/median/sd/iq</code>	A list with each element containing the vectors of mean/median/standard deviation/interquartile range of the posterior distributions of the row effects. The length of the list is equal to the number of row effects included i.e., <code>ncol(row.ids)</code> .
<code>row.sigma.mean/median/sd/iqr</code>	A list with each element containing the mean/median/standard deviation/interquartile range of the posterior distributions of the standard deviation for the row random effects normal distribution. The length of the list is equal to the number of row effects included i.e., <code>ncol(row.ids)</code> .
<code>ssvs.indcoefs.mean/ssvs.indcoefs.sd</code>	Matrices containing the SSVS posterior probabilities and associated standard deviation of including individual coefficients in the model (please see the section on SSVS above).
<code>ssvs.gpcoefs.mean/ssvs.gpcoefs.sd</code>	Matrices containing the SSVS posterior probabilities and associated standard deviation of including grouped coefficients in the model (please see the section on SSVS above).
<code>hpdintervals</code>	A list containing components which correspond to the lower and upper bounds of highest posterior density (HPD) intervals for all the parameters indicated above. Please see get.hpdintervals for more details.

ics	If <code>calc.ics = TRUE</code> , then a list of different information criteria values for the model calculated using <code>get.measures</code> is run. Please see help file for <code>get.measures</code> regarding details on the criteria. Also, please note the ics returned are based on <code>get.measures</code> with <code>more.measures = FALSE</code> .
jags.model	If <code>save.model = TRUE</code> , the raw jags model fitted is returned. This can be quite large!
n, p, family, trial.size, num.lv, ...	Various attributes of the model fitted, including the dimension of y , the response and model matrix used, distributional assumptions and trial sizes, number of latent variables, the number of covariates and traits, whether information criteria values were calculated, hyperparameters used in the Bayesian estimation, indices for SSVS, the number of levels for ordinal responses, and <code>n.burnin</code> , <code>n.iteration</code> and <code>n.thin</code> .

Distributions

For multivariate abundance data in ecology (also known as community composition data, Legendre and Gallagher, 2001), species counts are often overdispersed. Using a negative binomial distribution (`family = "negative.binomial"`) to model the counts usually helps to account for this overdispersion. Please note the variance for the negative binomial distribution is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the dispersion parameter.

For non-negative continuous data such as biomass, the lognormal and tweedie distributions may be used (Foster and Bravington, 2013). Note however that a common power parameter is used for tweedie columns – there is almost always insufficient information to model column-specific power parameters. Normal responses are also implemented, just in case you encounter normal stuff in ecology (pun intended)!

The beta distribution can be used to model data between values between but *not* including 0 and 1. In principle, this would make it useful for percent cover data in ecology, if it not were for the fact that percent cover is commonly characterized by having lots of zeros (which are not permitted for beta regression). An *ad-hoc* fix to this would be to add a very small value to shift the data away from exact zeros and/or ones. This is however heuristic, and pulls the model towards producing conservative results (see Smithson and Verkuilen, 2006, for a detailed discussion on beta regression, and Korhonen et al., 2007, for an example of an application to forest canopy cover data). Note the parameterization of the beta distribution used here is directly in terms of the mean μ and the dispersion parameter ϕ (more commonly know as the "sample size"). In terms of the two shape parameters, this is equivalent to $shape1 = a = \mu\phi$ and $shape2 = b = (1 - \mu)\phi$.

For ordinal response columns, cumulative probit regression is used (Agresti, 2010). `boral` assumes all ordinal columns are measured using the same scale i.e., all columns have the same number of theoretical levels, even though some levels for some species may not be observed. The number of levels is then assumed to be given by the maximum value from all the ordinal columns of y . Because of this, all ordinal columns then assumed to have the *same* cutoffs, τ , while the column-specific intercept effect, β_{0j} , allow for deviations away from these common cutoffs. That is,

$$probit(P(y_{ij} \leq k)) = \tau_k + \beta_{0j} + \dots,$$

where $probit(\cdot)$ is the probit function, $P(y_{ij} \leq k)$ is the cumulative probability of element y_{ij} being less than or equal to level k , τ_k is the cutoff linking levels k and $k + 1$ (and which are

increasing in k), β_{0j} are the column effects, and \dots denotes what else is included in the model, e.g. latent variables and related coefficients. To ensure model identifiability, and also because they are interpreted as species-specific deviations from the common cutoffs, the β_{0j} 's are treated as random effects and drawn from a normal distribution with mean zero and unknown standard deviation.

The parameterization above is useful for modeling ordinal in ecology. When ordinal responses are recorded, usually the same scale is applied to all species e.g., level 1 = not there, level 2 = a bit there, level 3 = lots there, level 4 = everywhere! The quantity τ_k can thus be interpreted as this common scale, while β_{0j} allows for deviations away from these to account for differences in species prevalence. Admittedly, the current implementation of *boral* for ordinal data can be quite slow.

Finally, in the event different responses are collected for different columns, e.g., some columns of y are counts, while other columns are presence-absence, one can specify different distributions for each column. Aspects such as variable selection, residual analysis, and plotting of the latent variables are, in principle, not affected by having different distributions. Naturally though, one has to be more careful with interpretation of the row effects α_i and latent variables z_i , as different link functions will be applied to each column of y . A situation where different distributions may prove useful is when y is a species–traits matrix, where each row is a species and each column a trait such as specific leaf area. In this case, traits could be of different response types, and the goal perhaps is to perform unconstrained ordination to look for patterns between species on an underlying trait surface e.g., a defense index for a species (Moles et al., 2013; see also the discussion below on how to perform model-based unconstrained ordination).

Why is only one MCMC chain run?

Much like the `MCMCfactanal` function in the `MCMCpack` package (Martin et al., 2011) for conducting factor analysis, which is a special case of the pure latent variable model with Gaussian responses, *boral* deliberately runs only one MCMC chain. This runs contrary to the recommendation of most Bayesian analyses, where the advice is to run multiple MCMC chains and check convergence using (most commonly) the Gelman-Rubin statistic or “Rhat” (Gelman et al., 2013). The main reason for this is that, in the context of MCMC sampling, the latent variable model is invariant to a switch of the sign, i.e. $z_i^T \theta_j = (-z_i)^T (-\theta_j)$, and so is actually unidentifiable. This is similar to well-known problem of label switching that occurs during the course of MCMC sampling for mixture models (see for instance, Section 4.9, McLachlan and Peel, 2004), and is due to the fact that the sign of the latent variables (ordination coordinates) is inherently arbitrary.

As a result of this sign-switching problem, it means that different MCMC chains can produce latent variables and corresponding coefficients values that, while having similar magnitudes, will be different in sign. Consequently, combining MCMC chains and checking Rhats, computing posterior means and medians etc...becomes inappropriate (in principle, one way to resolve this problem would be to post-process the MCMC chains and deal with sign switching, but this is really hard!). Therefore, to alleviate this issue together, *boral* chooses to only run one MCMC chain.

What does this mean for the user?

- For checking convergence, we recommend you look at trace plots of the MCMC chains. Using the `coda` package, which is automatically loaded when the *boral* package is loaded, try something like `traceplot(fit$jags.model, ask = TRUE)`. You could also try `geweke.diag` for Geweke’s convergence diagnostic, although no promises this necessarily does what is meant it!

- If you have a lot of data, e.g. lots of sites compared to species, sign-switching tends to be less of a problem and pops up less often.
- **IMPORTANTLY**, if the goal of your analysis is to inference while account for residual correlations between the columns of y , and not for model-based ordination, then the sign-switching problem is not a problem at all! This is because while the signs of the latent variables and associated coefficients may switch, the correlation and their signs are unaffected. In other words, looking the point estimates and credible intervals of regression coefficients β_j , and functions like `get.residual.cor` are unaffected by sign-switching.

Warnings

- *No* intercept column is required in X . Column-specific intercepts are estimated automatically and given by the first column of `lv.coefs`. Similarly, *no* intercept column is required in `traits`, as it is included automatically.
- If `num.lv > 5`, a warning is printed asking whether you really want to fit an `boral` with more than five latent variables. A warning is also printed if `num.lv == 1`, as this is not going to be successful in modeling between the correlation between columns.
- For models including both explanatory covariates and latent variables, one requires `num.lv > 1` to allow flexible modeling of the residual correlation matrix.
- MCMC can take a long time to run, especially with if the response matrix y is large! The calculation of information criteria (`calc.ics = TRUE`) can also take a while. Apologies for this advance =(
- MCMC with lots of ordinal columns take an especially long time to run! Moreover, estimates for the cutoffs in cumulative probit regression may be poor for levels with little data. Major apologies for this advance =(
- As discussed in the details, the use of information criterion should be done so with caution. What model to select should be first and foremost driven by the question of interest. Also, the use of information criterion in the presence of model selection using `SSVS` is problematic.
- Summaries of the coefficients such as posterior medians and HPD intervals may also be problematic when `SSVS` is being used, since the posterior distribution will be multi-modal.
- If `save.model = TRUE`, the raw `jags` model is also returned. This can be quite very memory-consuming, since it indirectly saves all the MCMC samples.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Agresti, A. (2010). Analysis of Ordinal Categorical Data. Wiley.
- Brown, et al. (2014). The fourth-corner solution - using predictive models to understand how species traits interact with the environment. *Methods in Ecology and Evolution* 5, 344-352.
- Foster, S. D. and Bravington, M. V. (2013). A Poisson-Gamma model for analysis of ecological non-negative continuous data. *Journal of Environmental and Ecological Statistics*, 20, 533-552.

- Gelman A. (2006) Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis* 1, 515-533.
- Gelman, et al. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2, 1360-1383.
- Gelman et al. (2013) *Bayesian Data Analysis*. CRC Press.
- George, E. I. and McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 85, 398-409.
- Hui et al. (2014). Model-based approaches to unconstrained ordination. *Methods in Ecology and Evolution*, 6, 399-411.
- Hill, M. O. (1974). Correspondence analysis: a neglected multivariate method. *Applied statistics*, 23, 340-354.
- Jamil, T., and ter Braak, C.J.F. (2013). Generalized linear mixed models can detect unimodal species-environment relationships. *PeerJ* 1: e95.
- Jamil, T. et al. (2013). Selecting traits that explain species-environment relationships: a generalized linear mixed model approach. *Journal of Vegetation Science* 24, 988-1000
- Korhonen, L., et al. (2007). Local models for forest canopy cover with beta regression. *Silva Fennica*, 41, 671-685.
- Kruskal, J. B. (1964). Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, 29, 115-129.
- Legendre, P. and Gallagher, E. D. (2001). Ecologically meaningful transformations for ordination of species data. *Oecologia*, 129, 271-280. *Numerical ecology*, Volume 20. Elsevier.
- Legendre, P. and Legendre, L. (2012). *Numerical ecology*, Volume 20. Elsevier.
- Martin et al. (2011). MCMCpack: Markov Chain Monte Carlo in R. *Journal of Statistical Software*, 42, 1-21. URL: <http://www.jstatsoft.org/v42/i09/>.
- McLachlan, G., and Peel, D. (2004). *Finite Mixture Models*. Wiley.
- Moles et al. (2013). Correlations between physical and chemical defences in plants: Trade-offs, syndromes, or just many different ways to skin a herbivorous cat? *New Phytologist*, 198, 252-263.
- O'Hara, B., and Sillianpaa, M.J. (2009). A Review of Bayesian Variable Selection Methods: What, How and Which. *Bayesian Analysis*, 4, 85-118.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*. March (pp. 20-22).
- Pollock, L. J. et al. (2014). Understanding co-occurrence by modelling species simultaneously with a Joint Species Distribution Model (JSDM). *Methods in Ecology and Evolution*, 5, 397-406.
- Skrondal, A., and Rabe-Hesketh, S. (2004). *Generalized latent variable modeling: Multilevel, longitudinal, and structural equation models*. CRC Press.
- Smithson, M., and Verkuilen, J. (2006). A better lemon squeezer? Maximum-likelihood regression with beta-distributed dependent variables. *Psychological methods*, 11, 54-71.
- Tenan, S., et al. (2014). Bayesian model selection: The steepest mountain to climb. *Ecological Modelling*, 283, 62-69.

- Warton et al. (2015). So Many Variables: Joint Modeling in Community Ecology. Trends in Ecology and Evolution, to appear.
- Warton et al. (2012). Distance-based multivariate analyses confound location and dispersion effects. Methods in Ecology and Evolution, 3, 89-101.
- Wang et al. (2013). mvabund: statistical methods for analysing multivariate abundance data. R package version 3.8.4.

See Also

[lvplot](#) for a scatter plot of the latent variables (and their coefficients if applicable) when `num.lv = 1` or 2, [coefplot](#) for horizontal line or "caterpillar plot" of the regression coefficients corresponding to X (if applicable), [summary.boral](#) for a summary of the fitted boral model, [get.measures](#) and [get.more.measures](#) for information criteria from the fitted boral model, [get.residual.cor](#) for calculating the residual correlation matrix.

Examples

```
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
X <- scale(spider$x)
n <- nrow(y); p <- ncol(y);

## NOTE: The values below MUST NOT be used in a real application;
## they are only used here to make the examples run quick!!!
example.mcmc.control <- list(n.burnin = 10, n.iteration = 100,
  n.thin = 1)

## Example 1 - model with two latent variables, site effects,
## and no environmental covariates
spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
  row.eff = "fixed", calc.ics = FALSE,
  mcmc.control = example.mcmc.control)

summary(spider.fit.nb)

par(mfrow = c(2,2))
plot(spider.fit.nb) ## Plots used in residual analysis,
dev.off()
## Used to check if assumptions such as a mean-variance relationship
## are adequately satisfied.

lvplot(spider.fit.nb) ## Biplot of the latent variables,
## which can be interpreted in the same manner as an ordination plot.

## Not run:
## Example 2a - model with no latent variables, no site effects,
## and environmental covariates
spider.fit.nb <- boral(y, X = X, family = "negative.binomial",
  num.lv = 0, mcmc.control = example.mcmc.control)
```



```

summary(spider.fit.nb)
## The results can be compared with the default example from
## the manyglm() function in mvabund. Hopefully they are similar =>

## Caterpillar plots for the coefficients
par(mfrow=c(2,3), mar = c(5,6,1,1))
sapply(colnames(spider.fit.nb$X), coefplot, x = spider.fit.nb)

## Example 2b - suppose now, for some reason, the 28 rows were
## sampled such into four replications of seven sites
## Let us account for this as a fixed effect
spider.fit.nb <- boral(y, X = X, family = "negative.binomial",
num.lv = 0, row.eff = "fixed", row.ids = matrix(rep(1:7,each=4),ncol=1),
mcmc.control = example.mcmc.control)

spider.fit.nb$row.coefs

## Example 2c - suppose now, for some reason, the 28 rows reflected
## a nested design with seven regions, each with four sub-regions
## We can account for this nesting as a random effect
spider.fit.nb <- boral(y, X = X, family = "negative.binomial",
num.lv = 0, row.eff = "random",
row.ids = cbind(1:n, rep(1:7,each=4)),
mcmc.control = example.mcmc.control)

spider.fit.nb$row.coefs

## Example 3a - Extend example 2 to demonstrate grouped covariate selection
## on the last three covariates.
set.prior <- list(type = c("normal","normal","normal","uniform"),
hypparams = c(100, 20, 100, 50), ssvs.index = c(-1,-1,-1,1,2,3))
spider.fit.nb2 <- boral(y, X = X, family = "negative.binomial",
num.lv = 0, calc.ics = FALSE, mcmc.control = example.mcmc.control,
prior.control = set.prior)

summary(spider.fit.nb2)

## Example 3b - Extend example 2 to demonstrate individual covariate selection
## on the last three covariates.
set.prior <- list(type = c("normal","normal","normal","uniform"),
hypparams = c(100, 20, 100, 50), ssvs.index = c(-1,-1,-1,0,0,0))
spider.fit.nb3 <- boral(y, X = X, family = "negative.binomial",
num.lv = 0, calc.ics = FALSE, mcmc.control = example.mcmc.control,
prior.control = set.prior)
summary(spider.fit.nb3)

## Example 4 - model fitted to presence-absence data, no site effects, and
## two latent variables

```

```

data(tikus)
y <- tikus$abun
y[y > 0] <- 1
y <- y[1:20,] ## Consider only years 1981 and 1983
y <- y[,apply(y > 0,2,sum) > 2] ## Consider only spp with more than 2 presences

tikus.fit <- boral(y, family = "binomial", num.lv = 2,
  calc.ics = FALSE, mcmc.control = example.mcmc.control)

lvplot(tikus.fit, biplot = FALSE)
## A strong location between the two sampling years

## Example 5 - model fitted to count data, no site effects, and
## two latent variables, plus traits included to explain environmental responses
data(antTraits)
y <- antTraits$abun
X <- as.matrix(scale(antTraits$env))
## Include only traits 1, 2, and 5
traits <- as.matrix(antTraits$traits[,c(1,2,5)])
which.traits <- vector("list",ncol(X)+1)
for(i in 1:length(which.traits)) which.traits[[i]] <- 1:ncol(traits)
## Just for fun, the regression coefficients for the second column of X,
## corresponding to the third element in the list which.traits,
## will be estimated separately and not regressed against traits.
which.traits[[3]] <- 0

fit.traits <- boral(y, X = X, traits = traits, which.traits = which.traits,
  family = "negative.binomial", calc.ics = FALSE,
  mcmc.control = example.mcmc.control)

summary(fit.traits)

## Example 6 - simulate Bernoulli data, based on a model with two latent variables,
## no site variables, with two traits and one environmental covariates
## This example is a proof of concept that traits can used to
## explain environmental responses
library(mvtnorm)

n <- 100; s <- 50
X <- as.matrix(scale(1:n))
colnames(X) <- c("elevation")

traits <- cbind(rbinom(s,1,0.5), rnorm(s))
## one categorical and one continuous variable
colnames(traits) <- c("thorns-dummy","SLA")

simfit <- list(true.lv = rmvnorm(n, mean = rep(0,2)),
  lv.coefs = cbind(rnorm(s), rmvnorm(s, mean = rep(0,2))),
  traits.coefs = matrix(c(0.1,1,-0.5,1,0.5,0,-1,1), 2, byrow = TRUE))
rownames(simfit$traits.coefs) <- c("beta0","elevation")
colnames(simfit$traits.coefs) <- c("kappa0","thorns-dummy","SLA","sigma")

```

```

simy = create.life(true.lv = simfit$true.lv, lv.coefs = simfit$lv.coefs, X = X,
traits = traits, traits.coefs = simfit$traits.coefs, family = "binomial")

which.traits <- vector("list",ncol(X)+1)
for(i in 1:length(which.traits)) which.traits[[i]] <- 1:ncol(traits)
fit.traits <- boral(y = simy, X = X, traits = traits, which.traits = which.traits,
family = "binomial", num.lv = 2, save.model = TRUE, calc.ics = FALSE,
mcmc.control = example.mcmc.control)

## End(Not run)

```

calc.condlogLik *Conditional log-likelihood for an boral model*

Description

Calculates the conditional log-likelihood for a set of parameter estimates from an boral model, where everything is treated as "fixed effects" including latent variables, row effects, and so on.

Usage

```

calc.condlogLik(y, X = NULL, family, trial.size = 1, lv.coefs,
X.coefs = NULL, row.coefs = NULL, row.ids = NULL,
lv = NULL, cutoffs = NULL, powerparam = NULL)

```

Arguments

y	The response matrix the boral model was fitted to.
X	The model matrix used in the boral model. Defaults to NULL, in which case it is assumed no model matrix was used.
family	Either a single element, or a vector of length equal to the number of columns in y. The former assumes all columns of y come from this distribution. The latter option allows for different distributions for each column of y. Elements can be one of "binomial" (with probit link), "poisson" (with log link), "negative.binomial" (with log link), "normal" (with identity link), "lnormal" for log-normal (with log link), "tweedie" (with log link), "exponential" (with log link), "gamma" (with log link), "beta" (with logit link), "ordinal" (cumulative probit regression).

For the negative binomial distribution, the variance is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the column-specific dispersion parameter. For the normal distribution, the variance is parameterized as $Var(y) = \phi^2$, where ϕ is the column-specific standard deviation. For the tweedie distribution, the variance is parameterized as $Var(y) = \phi\mu^p$ where ϕ is the column-specific dispersion

parameter and p is a power parameter common to all columns assumed to be tweedie, with $1 < p < 2$. For the gamma distribution, the variance is parameterized as $Var(y) = \mu/\phi$ where ϕ is the column-specific rate (henceforth referred to also as dispersion parameter). For the beta distribution, the parameterization is in terms of the mean μ and sample size ϕ (henceforth referred to also as dispersion parameter), so that the two shape parameters are given by $a = \mu\phi$ and $b = (1 - \mu)\phi$.

All columns assumed to have ordinal responses are constrained to have the same cutoffs points, with a column-specific intercept to account for differences between the columns (please see *Details* for formulation).

trial.size	Either equal to a single element, or a vector of length equal to the number of columns in y . If a single element, then all columns assumed to be binomially distributed will have trial size set to this. If a vector, different trial sizes are allowed in each column of y . The argument is ignored for all columns not assumed to be binomially distributed. Defaults to 1, i.e. Bernoulli distribution.
lv.coefs	The column-specific intercept, coefficient estimates relating to the latent variables, and dispersion parameters from the boral model.
X.coefs	The coefficients estimates relating to the model matrix X from the boral model. Defaults to NULL, in which it is assumed there are no covariates in the model.
row.coefs	Row effect estimates for the boral model. The conditional likelihood is defined conditional on these estimates i.e., they are also treated as "fixed effects". Defaults to NULL, in which case it is assumed there are no row effects in the model.
row.ids	A matrix with the number of rows equal to the number of rows in y , and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row i in y for random effect eqnj; please see the help file for the main boral function for details. Defaults to NULL, so that if row.coefs = NULL then the argument is ignored, otherwise if row.coefs is supplied then row.ids = matrix(1:nrow(y), ncol = 1) i.e., a single, row effect unique to each row. An internal check is done to see row.coefs and row.ids are consistent in terms of arguments supplied.
lv	Latent variables "estimates" from the boral model, which the conditional likelihood is based on. Defaults to NULL, in which case it is assumed no latent variables were included in the model.
cutoffs	Common cutoff estimates from the boral model when any of the columns of y are ordinal responses. Defaults to NULL.
powerparam	Common power parameter from the boral model when any of the columns of y are tweedie responses. Defaults to NULL.

Details

For an $n \times p$ response matrix y , suppose we fit an boral model with one or more latent variables. If we denote the latent variables by $z_i; i = 1, \dots, n$, then the conditional log-likelihood is given by,

$$\log(f) = \sum_{i=1}^n \sum_{j=1}^p \log(f(y_{ij}|z_i, \theta_j, \beta_{0j}, \dots)),$$

where $f(y_{ij}|\cdot)$ is the assumed distribution for column j , z_i are the latent variables and θ_j are the coefficients relating to them, β_{0j} are column-specific intercepts, and \dots denotes anything else included in the model, such as row effects, regression coefficients related X and traits, etc...

The key difference between this and the marginal likelihood (see [calc.marglogLik](#)) is that the conditional likelihood treats everything as "fixed effects" i.e., conditions on them. These include the latent variables z_i and other parameters that were included in the model as random effects e.g., row effects if `row.eff = "random"`, regression coefficients related to X if traits were included in the model (see the section titled "Including species traits" in the main `boral` fitting function), and so on.

The conditional DIC, WAIC, EAIC, and EBIC returned from [get.measures](#) are based on the conditional likelihood calculated from this function. Additionally, [get.measures](#) returns the conditional likelihood evaluated at all MCMC samples of a fitted `boral` model.

Value

A list with the following components:

<code>logLik</code>	Value of the conditional log-likelihood.
<code>logLik.comp</code>	A matrix of the log-likelihood values for each element in <code>y</code> , such that <code>sum(logLik.comp) = logLik</code> .

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

See Also

[get.measures](#) for some information criteria based on the conditional log-likelihood; [calc.logLik.lv0](#) to calculate the conditional/marginal log-likelihood for an `boral` model with no latent variables; [calc.marglogLik](#) for calculation of the marginal log-likelihood;

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

## Example 1 - model with 2 latent variables, site effects,
## and no environmental covariates
spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
row.eff = "fixed", save.model = TRUE, calc.ics = FALSE)

## Extract all MCMC samples
fit.mcmc <- mcmc(spider.fit.nb$jags.model$BUGSoutput$sims.matrix)

## Find the posterior medians
coef.mat <- matrix(apply(fit.mcmc[,grep("all.params", colnames(fit.mcmc))],
2, median), nrow=p)
```

```

site.coef <- list(ID1 = apply(fit.mcmc[,grep("row.params", colnames(fit.mcmc))],
  2,median))
lvs.mat <- matrix(apply(fit.mcmc[,grep("lvs", colnames(fit.mcmc))],2,median),nrow=n)

## Calculate the conditional log-likelihood at the posterior median
calc.condlogLik(y, family = "negative.binomial",
  lv.coefs = coef.mat, row.coefs = site.coef, lv = lvs.mat)

## Example 2 - model with no latent variables and environmental covariates
X <- scale(spider$x)
spider.fit.nb2 <- boral(y, X = X, family = "negative.binomial", num.lv = 0,
  save.model = TRUE, calc.ics = FALSE)

## Extract all MCMC samples
fit.mcmc <- mcmc(spider.fit.nb2$jags.model$BUGSoutput$sims.matrix)

## Find the posterior medians
coef.mat <- matrix(apply(fit.mcmc[,grep("all.params", colnames(fit.mcmc))],
  2,median),nrow=p)
X.coef.mat <- matrix(apply(fit.mcmc[,grep("X.params", colnames(fit.mcmc))],
  2,median),nrow=p)

## Calculate the log-likelihood at the posterior median
calc.condlogLik(y, X = X, family = "negative.binomial",
  lv.coefs = coef.mat, X.coefs = X.coef.mat)

## End(Not run)

```

calc.logLik.lv0

Log-likelihood for a boral model with no latent variables

Description

Calculates the log-likelihood for a set of parameter estimates from an boral model with no latent variables. If the row effects are assumed to be random, they are integrated over using Monte Carlo integration.

Usage

```

calc.logLik.lv0(y, X = NULL, family, trial.size = 1, lv.coefs,
  X.coefs = NULL, row.eff = "none", row.params = NULL,
  row.ids = NULL, cutoffs = NULL, powerparam = NULL)

```

Arguments

y The response matrix the boral model was fitted to.

X The model matrix used in the boral model. Defaults to NULL, in which case it is assumed no model matrix was used.

family	<p>Either a single element, or a vector of length equal to the number of columns in y. The former assumes all columns of y come from this distribution. The latter option allows for different distributions for each column of y. Elements can be one of "binomial" (with probit link), "poisson" (with log link), "negative.binomial" (with log link), "normal" (with identity link), "lnormal" for log-normal (with log link), "tweedie" (with log link), "exponential" (with log link), "gamma" (with log link), "beta" (with logit link), "ordinal" (cumulative probit regression).</p> <p>For the negative binomial distribution, the variance is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the column-specific dispersion parameter. For the normal distribution, the variance is parameterized as $Var(y) = \phi^2$, where ϕ is the column-specific standard deviation. For the tweedie distribution, the variance is parameterized as $Var(y) = \phi\mu^p$ where ϕ is the column-specific dispersion parameter and p is a power parameter common to all columns assumed to be tweedie, with $1 < p < 2$. For the gamma distribution, the variance is parameterized as $Var(y) = \mu/\phi$ where ϕ is the column-specific rate (henceforth referred to also as dispersion parameter). For the beta distribution, the parameterization is in terms of the mean μ and sample size ϕ (henceforth referred to also as dispersion parameter), so that the two shape parameters are given by $a = \mu\phi$ and $b = (1 - \mu)\phi$.</p> <p>All columns assumed to have ordinal responses are constrained to have the same cutoffs points, with a column-specific intercept to account for differences between the columns (please see <i>Details</i> for formulation).</p>
trial.size	<p>Either equal to a single element, or a vector of length equal to the number of columns in y. If a single element, then all columns assumed to be binomially distributed will have trial size set to this. If a vector, different trial sizes are allowed in each column of y. The argument is ignored for all columns not assumed to be binomially distributed. Defaults to 1, i.e. Bernoulli distribution.</p>
lv.coefs	<p>The column-specific intercept, coefficient estimates relating to the latent variables, and dispersion parameters from the boral model.</p>
X.coefs	<p>The coefficients estimates relating to the model matrix X from the boral model. Defaults to NULL, in which it is assumed there are no covariates in the model.</p>
row.eff	<p>Single element indicating whether row effects are included as fixed effects ("fixed"), random effects ("random") or not included ("none") in the boral model. If random effects, they are drawn from a normal distribution with mean zero and standard deviation given by row.params. Defaults to "none".</p>
row.params	<p>Parameters corresponding to the row effect from the boral model. If row.eff = "fixed", then these are the fixed effects and should have length equal to the number of columns in y. If row.eff = "random", then this is the standard deviation for the random effects normal distribution. If row.eff = "none", then this argument is ignored.</p>
row.ids	<p>A matrix with the number of rows equal to the number of rows in y, and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row i in y for random effect eqnj; please see the help file for the main boral function for details. Defaults to NULL, so that if row.params = NULL then the argument is ignored, otherwise if row.params is supplied then row.ids = matrix(1:nrow(y), ncol = 1)</p>

	i.e., a single, row effect unique to each row. An internal check is done to see row.params and row.ids are consistent in terms of arguments supplied.
cutoffs	Common cutoff estimates from the boral model when any of the columns of y are ordinal responses. Defaults to NULL.
powerparam	Common power parameter from the boral model when any of the columns of y are tweedie responses. Defaults to NULL.

Details

For an $n \times p$ response matrix y , the log-likelihood for a model with no latent variables included is given by,

$$\log(f) = \sum_{i=1}^n \sum_{j=1}^p \log(f(y_{ij} | \beta_{0j}, \alpha_i, \dots)),$$

where $f(y_{ij} | \cdot)$ is the assumed distribution for column j , β_{0j} is the column-specific intercepts, α_i is the row effect, and \dots generically denotes anything else included in the model, e.g. row effects, dispersion parameters etc...

Please note the function is written conditional on all regression coefficients. Therefore, if traits are included in the model, in which case the regression coefficients β_{0j}, β_j become random effects instead (please see section titled “Including species traits” in the main boral function), then the calculation of the log-likelihood does NOT take this into account, i.e. does not marginalize over them!

Likewise if more than two columns are ordinal responses, then the regression coefficients β_{0j} corresponding to these columns become random effects, and the calculation of the log-likelihood also does NOT take this into account, i.e. does not marginalize over them!

In the special case where α_i is a random row effect, then the log-likelihood is calculated by integrating over this,

$$\log(f) = \sum_{i=1}^n \log\left(\int \prod_{j=1}^p (f(y_{ij} | \beta_{0j}, \alpha_i, \dots)) f(\alpha_i) d\alpha_i\right),$$

where $f(\alpha_i)$ is the random effects distribution with mean zero and standard deviation given by the row.params. The integration is performed using standard Monte Carlo integration. Beyond this however, for more complicated random row effects structures the function currently does not permit the calculation of the log-likelihood in such cases, as the integration over these random effects becomes too complicated to perform in such case. Sorry!

Value

A list with the following components:

logLik	Value of the log-likelihood
logLik.comp	A vector of the log-likelihood values for each row of y, such that $\text{sum}(\text{logLik.comp}) = \text{logLik}$.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

See Also

[calc.marglogLik](#) for calculation of the log-likelihood marginalizing over one or more latent variables, and [calc.condlogLik](#) for calculation of the conditional log-likelihood for models where everything is treated as "fixed effects", including latent variables, row effects, and so on.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

## Example 1 - NULL model with site effects only
spider.fit.nb <- boral(y, family = "negative.binomial",
  row.eff = "fixed", save.model = TRUE, calc.ics = FALSE)

## Extract all MCMC samples
fit.mcmc <- mcmc(spider.fit.nb$jags.model$BUGSoutput$sims.matrix)

## Find the posterior medians
coef.mat <- matrix(apply(fit.mcmc[,grep("all.params", colnames(fit.mcmc))],
  2, median), nrow=p)
site.coef <- list(ID1 = apply(fit.mcmc[,grep("row.params", colnames(fit.mcmc))],
  2, median))

## Calculate the log-likelihood at the posterior median
calc.logLik.lv0(y, family = "negative.binomial",
  lv.coefs = coef.mat, row.eff = "fixed", row.params = site.coef)

## Example 2 - Model with environmental covariates and random row effects
X <- scale(spider$x)
spider.fit.nb2 <- boral(y, X = X, family = "negative.binomial", row.eff = "random",
  save.model = TRUE, calc.ics = FALSE)

## Extract all MCMC samples
fit.mcmc <- mcmc(spider.fit.nb2$jags.model$BUGSoutput$sims.matrix)

## Find the posterior medians
coef.mat <- matrix(apply(fit.mcmc[,grep("all.params", colnames(fit.mcmc))],
  2, median), nrow=p)
X.coef.mat <- matrix(apply(fit.mcmc[,grep("X.params", colnames(fit.mcmc))],
  2, median), nrow=p)
site.sigma <- list(ID1 =
  median(fit.mcmc[,grep("row.ranef.sigma", colnames(fit.mcmc))]))
```

```
## Calculate the log-likelihood at the posterior median
calc.logLik.lv0(y, X = spider$x, family = "negative.binomial", row.eff = "random",
lv.coefs = coef.mat, X.coefs = X.coef.mat, row.params = site.sigma)

## End(Not run)
```

calc.marglogLik

Marginal log-likelihood for an boral model

Description

Calculates the marginal log-likelihood for a set of parameter estimates from an boral model, whereby the latent variables and random effects (if applicable) are integrated out. The integration is performed using Monte Carlo integration.

Usage

```
calc.marglogLik(y, X = NULL, family, trial.size = 1, lv.coefs,
X.coefs = NULL, row.eff = "none", row.params = NULL, row.ids = NULL,
num.lv, lv.mc = NULL, cutoffs = NULL, powerparam = NULL)
```

Arguments

y	The response matrix that the boral model was fitted to.
X	The model matrix used in the boral model. Defaults to NULL, in which case it is assumed no model matrix was used.
family	Either a single element, or a vector of length equal to the number of columns in y. The former assumes all columns of y come from this distribution. The latter option allows for different distributions for each column of y. Elements can be one of "binomial" (with probit link), "poisson" (with log link), "negative.binomial" (with log link), "normal" (with identity link), "lnormal" for log-normal (with log link), "tweedie" (with log link), "exponential" (with log link), "gamma" (with log link), "beta" (with logit link), "ordinal" (cumulative probit regression).

For the negative binomial distribution, the variance is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the column-specific dispersion parameter. For the normal distribution, the variance is parameterized as $Var(y) = \phi^2$, where ϕ is the column-specific standard deviation. For the tweedie distribution, the variance is parameterized as $Var(y) = \phi\mu^p$ where ϕ is the column-specific dispersion parameter and p is a power parameter common to all columns assumed to be tweedie, with $1 < p < 2$. For the gamma distribution, the variance is parameterized as $Var(y) = \mu/\phi$ where ϕ is the column-specific rate (henceforth referred to also as dispersion parameter). For the beta distribution, the parameterization is in terms of the mean μ and sample size ϕ (henceforth referred to also as dispersion parameter), so that the two shape parameters are given by $a = \mu\phi$ and $b = (1 - \mu)\phi$.

All columns assumed to have ordinal responses are constrained to have the same cutoffs points, with a column-specific intercept to account for differences between the columns (please see *Details* for formulation).

trial.size	Either equal to a single element, or a vector of length equal to the number of columns in y . If a single element, then all columns assumed to be binomially distributed will have trial size set to this. If a vector, different trial sizes are allowed in each column of y . The argument is ignored for all columns not assumed to be binomially distributed. Defaults to 1, i.e. Bernoulli distribution.
lv.coefs	The column-specific intercept, coefficient estimates relating to the latent variables, and dispersion parameters from the boral model.
X.coefs	The coefficients estimates relating to the model matrix X from the boral model. Defaults to NULL, in which it is assumed there are no covariates in the model.
row.eff	Single element indicating whether row effects are included as fixed effects ("fixed"), random effects ("random") or not included ("none") in the boral model. If random effects, they are drawn from a normal distribution with mean zero and standard deviation given by row.params. Defaults to "none".
row.params	Parameters corresponding to the row effect from the boral model. If row.eff = "fixed", then these are the fixed effects and should have length equal to the number of columns in y . If row.eff = "random", then this is standard deviation for the random effects normal distribution. If row.eff = "none", then this argument is ignored.
row.ids	A matrix with the number of rows equal to the number of rows in y , and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row i in y for random effect eqnj; please see the help file for the main boral function for details. Defaults to NULL, so that if row.params = NULL then the argument is ignored, otherwise if row.params is supplied then row.ids = matrix(1:nrow(y), ncol = 1) i.e., a single, row effect unique to each row. An internal check is done to see row.params and row.ids are consistent in terms of arguments supplied.
num.lv	The number of latent variables used in the boral model. For boral models with no latent variables, please use <code>calc.logLik.lv0</code> to calculate the log-likelihood.
lv.mc	A matrix used for performing the Monte Carlo integration. Defaults to NULL, in which case a matrix is generated within the function.
cutoffs	Common cutoff estimates from the boral model when any of the columns of y are ordinal responses. Defaults to NULL.
powerparam	Common power parameter from the boral model when any of the columns of y are tweedie responses. Defaults to NULL.

Details

For an $n \times p$ response matrix y , suppose we fit an boral model with one or more latent variables. If we denote the latent variables by $\mathbf{z}_i; i = 1, \dots, n$, then the marginal log-likelihood is given by

$$\log(f) = \sum_{i=1}^n \log\left(\int \prod_{j=1}^p f(y_{ij} | \mathbf{z}_i, \beta_{0j}, \boldsymbol{\theta}_j, \dots) f(\mathbf{z}_i) d\mathbf{z}_i\right),$$

where $f(y_{ij}|\cdot)$ is the assumed distribution for column j , β_{0j} are the column-specific intercepts, θ_j are the column-specific latent variable coefficients, and \dots generically denotes anything else included in the model, e.g. row effects, dispersion parameters etc... The quantity $f(z_i)$ denotes the distribution of the latent variable, which is assumed to be standard multivariate Gaussian. Standard Monte Carlo integration is used for calculating the marginal likelihood. If `lv.mc = NULL`, the function automatically generates a matrix as

`lv.mc <- cbind(1, rmvnorm(2000, rep(0, num.lv)))`. If there is a need to apply this function numerous times, we recommend a matrix be inserted into `lv.mc` to speed up computation.

The key difference between this and the conditional likelihood (see `calc.condlogLik`) is that the marginal likelihood treats the latent variables as "random effects" and integrates over them, whereas the conditional likelihood treats the latent variables as "fixed effects".

Please note the function is written conditional on all regression coefficients. Therefore, if traits are included in the model, in which case the regression coefficients β_{0j}, β_j become random effects instead (please see section titled "Including species traits" in the main `boral` function), then the calculation of the log-likelihood does NOT take this into account, i.e. does not marginalize over them!

Likewise if more than two columns are ordinal responses, then the regression coefficients β_{0j} corresponding to these columns become random effects, and the calculation of the log-likelihood also does NOT take this into account, i.e. does not marginalize over them!

In the special case where α_i is a random row effect, then the log-likelihood is calculated by integrating over this as well,

$$\log(f) = \sum_{i=1}^n \log\left(\int \prod_{j=1}^p (f(y_{ij}|z_i, \beta_{0j}, \theta_j, \alpha_i, \dots)) f(z_i) f(\alpha_i) dz_i d\alpha_i\right),$$

where $f(\alpha_i)$ is the random effects distribution with mean zero and standard deviation given by the `row.params`. The integration is performed using standard Monte Carlo integration. Beyond this however, for more complicated random row effects structures the function currently does not permit the calculation of the log-likelihood in such cases, as the integration over these random effects becomes too complicated to perform in such case. Sorry!

Value

A list with the following components:

<code>logLik</code>	Value of the marginal log-likelihood.
<code>logLik.comp</code>	A vector of the log-likelihood values for each row of <code>y</code> , such that <code>sum(logLik.comp) = logLik</code> .

Note

The AIC and BIC at posterior median returned from `get.measures` are all based on the marginal log-likelihood calculated from this function. Additionally, `get.more.measures` returns even more information criteria based on the marginal log-likelihood. As mentioned in the details though, these information criteria do not take into account that traits are included in the model!

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

See Also

[get.measures](#) and [get.more.measures](#) for information criteria based on the marginal log-likelihood; [calc.condlogLik](#) for calculation of the conditional log-likelihood; [calc.logLik.lv0](#) to calculate the conditional/marginal log-likelihood for an boral model with no latent variables.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

## Example 1 - model with two latent variables, site effects,
## and no environmental covariates
spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
  row.eff = "fixed", save.model = TRUE, calc.ics = FALSE)

## Extract all MCMC samples
fit.mcmc <- mcmc(spider.fit.nb$jags.model$BUGSoutput$sims.matrix)

## Find the posterior medians
coef.mat <- matrix(apply(fit.mcmc[,grep("all.params", colnames(fit.mcmc))],
  2, median), nrow=p)
site.coef <- list(ID1 = apply(fit.mcmc[,grep("row.params", colnames(fit.mcmc))],
  2, median))

## Calculate the marginal log-likelihood at the posterior median
calc.marglogLik(y, family = "negative.binomial",
  lv.coefs = coef.mat, row.eff = "fixed", row.params = site.coef,
  num.lv = 2)

## Example 2 - model with one latent variable, no site effects,
## and environmental covariates
spider.fit.nb2 <- boral(y, X = spider$x, family = "negative.binomial",
  num.lv = 2, save.model = TRUE, calc.ics = FALSE)

## Extract all MCMC samples
fit.mcmc <- mcmc(spider.fit.nb2$jags.model$BUGSoutput$sims.matrix)

## Find the posterior medians
coef.mat <- matrix(apply(fit.mcmc[,grep("all.params", colnames(fit.mcmc))],
  2, median), nrow=p)
X.coef.mat <- matrix(apply(fit.mcmc[,grep("X.params", colnames(fit.mcmc))],
  2, median), nrow=p)

## Calculate the log-likelihood at the posterior median
```

```

calc.marglogLik(y, X = spider$x, family = "negative.binomial",
lv.coefs = coef.mat, X.coefs = X.coef.mat, num.lv = 2)

## End(Not run)

```

coefspplot

Caterpillar plots of the regression coefficients from a boral model

Description

Constructs horizontal line plot (point estimate and HPD intervals), otherwise known as "caterpillar plots", for the column-specific regression coefficients corresponding to a covariate in X fitted in the boral model.

Usage

```
coefspplot(covname, x, labely = NULL, est = "median", ...)
```

Arguments

covname	The name of one of the covariates fitted in the boral model. That is, it must be a character vector corresponding to one of the elements in <code>colnames(x)\$X.coefs.median</code> .
x	An object for class "boral".
labely	Controls the labels on the y-axis for the line plot. If it is not NULL, then it must be a vector either of length 1 or the same length as the number of columns in the y in the fitted boral object. In the former, it is treated as the y-axis label. In the latter, it is used in place of the column names of y to label each line. Defaults to NULL, in which the each line in the plot is labeled according to the columns of y, or equivalently <code>rownames(x)\$X.coefs.median</code> .
est	A choice of either the posterior median (<code>est = "median"</code>) or posterior mean (<code>est = "mean"</code>), which are then used as the point estimates in the lines. Default is posterior median.
...	Additional graphical options to be included in. These include values for <code>cex</code> , <code>cex.lab</code> , <code>cex.axis</code> , <code>cex.main</code> , <code>lwd</code> , and so on.

Details

For each species (column of y), the horizontal line or "caterpillar" is constructed by first marking the point estimate (posterior mean or median) with an "x" symbol. Then the line is construed based on the lower and upper limits of the highest posterior density (HPD) intervals as found in `x$hp dintervals`. By default these intervals of 95% HPD intervals. To complete the plot, a vertical dotted line is drawn to denote the zero value. All HPD intervals that include zero are colored gray, while HPD intervals that exclude zero are colored black.

The graph is probably better explained by, well, plotting it using the toy example below =P

Thanks to Robert O'Hara for suggesting and providing the original code for this function.

Value

If SSVS was applied individually to each coefficient of X when fitting the boral model, then the posterior probabilities of including the specified covariate are printed out i.e., those from `x$ssvs.indcoefs.mean`.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

See Also

`caterplot` from the `mcmcplots` package for other, sexier caterpillar plots.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

X <- scale(spider$x)
spider.fit.nb <- boral(y, X = X, family = "negative.binomial",
num.lv = 2)

## Do separate line plots for all the coefficients of X
par(mfrow=c(2,3), mar = c(5,6,1,1))
sapply(colnames(spider.fit.nb$X), coefplot,
spider.fit.nb)

## End(Not run)
```

create.life

Simulate a Multivariate Response Matrix

Description

Simulate a multivariate response matrix, given parameters such as but not necessarily all of: family, number of latent variables and related coefficients, an matrix of explanatory variables and related coefficients, row effects, cutoffs for cumulative probit regression of ordinal responses.

Usage

```
create.life(true.lv = NULL, lv.coefs, X = NULL, X.coefs = NULL,
traits = NULL, traits.coefs = NULL, family, row.eff = "none",
row.params = NULL, row.ids = NULL, trial.size = 1, cutoffs = NULL,
```

```

powerparam = NULL, manual.dim = NULL, save.params = FALSE)

## S3 method for class 'boral'
simulate(object, nsim = 1, seed = NULL, est = "median", ...)

```

Arguments

<code>object</code>	An object of class "boral".
<code>nsim</code>	Number of multivariate response matrices to simulate. Defaults to 1.
<code>seed</code>	Seed for dataset simulation. Defaults to NULL, in which case no seed is set.
<code>est</code>	A choice of either the posterior median (<code>est == "median"</code>) or posterior mean (<code>est == "mean"</code>), which are then treated as estimates and the fitted values are calculated from. Default is posterior median.
<code>true.lv</code>	A matrix of true latent variables. With multivariate abundance data in ecology for instance, each row corresponds to the true site ordination coordinates. Defaults to NULL, in which case no latent variables are included.
<code>lv.coefs</code>	A matrix containing column-specific intercepts, latent variable coefficients relating to <code>true.lv</code> , and dispersion parameters.
<code>X</code>	An model matrix of covariates, which can be included as part of the data generation. Defaults to NULL, in which case no model matrix is used. No intercept column should be included in <code>X</code> .
<code>X.coefs</code>	The coefficients relating to the model matrix <code>X</code> . Defaults to NULL. This argument needs to be supplied if <code>X</code> is supplied and no <code>traits</code> are supplied.
<code>traits</code>	A model matrix of species covariates, which can be included as part of the data generation. Defaults to NULL, in which case no matrix is used. No intercept column should be included in <code>traits</code> , as it is included automatically.
<code>traits.coefs</code>	<p>A matrix of coefficients that are used to generate "new" column-specific intercepts and <code>X.coefs</code>. The number of rows should equal to $(\text{ncol}(X)+1)$ and the number of columns should equal to $(\text{ncol}(\text{traits})+2)$.</p> <p>How this argument works is as follows: when both <code>traits</code> and <code>traits.coefs</code> are supplied, then new column-specific intercepts (i.e. the first column of <code>lv.coefs</code> is overwritten) are generated by simulating from a normal distribution with mean equal to $\text{traits.coefs}[1,1] + \text{traits} * \text{traits.coefs}[1,2:(\text{ncol}(\text{traits.coefs})-1)]$ and standard deviation $\text{traits.coefs}[1,\text{ncol}(\text{traits.coefs})]$. In other words, the last column of <code>traits.coefs</code> provides the standard deviation of the normal distribution, with the other columns being the regression coefficients in the mean of the normal distribution. Analogously, new <code>X.coefs</code> are generated in the same manner using the remaining rows of <code>traits.coefs</code>. Please see the section on including species traits in the help file for boral for more information.</p> <p>It is important that highlight then with in this data generation mechanism, the new column-specific intercepts and <code>X.coefs</code> are now random effects, being drawn from a normal distribution.</p> <p>Defaults to NULL, in conjunction with <code>traits = NULL</code>.</p>

family	<p>Either a single element, or a vector of length equal to the number of columns in y. The former assumes all columns of y come from this distribution. The latter option allows for different distributions for each column of y. Elements can be one of "binomial" (with probit link), "poisson" (with log link), "negative.binomial" (with log link), "normal" (with identity link), "lnormal" for log-normal (with log link), "tweedie" (with log link), "exponential" (with log link), "gamma" (with log link), "beta" (with logit link), "ordinal" (cumulative probit regression).</p> <p>For the negative binomial distribution, the variance is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the column-specific dispersion parameter. For the normal distribution, the variance is parameterized as $Var(y) = \phi^2$, where ϕ is the column-specific standard deviation. For the tweedie distribution, the variance is parameterized as $Var(y) = \phi\mu^p$ where ϕ is the column-specific dispersion parameter and p is a power parameter common to all columns assumed to be tweedie, with $1 < p < 2$. For the gamma distribution, the variance is parameterized as $Var(y) = \mu/\phi$ where ϕ is the column-specific rate (henceforth referred to also as dispersion parameter). For the beta distribution, the parameterization is in terms of the mean μ and sample size ϕ (henceforth referred to also as dispersion parameter), so that the two shape parameters are given by $a = \mu\phi$ and $b = (1 - \mu)\phi$.</p> <p>All columns assumed to have ordinal responses are constrained to have the same cutoffs points, with a column-specific intercept to account for differences between the columns (please see <i>Details</i> for formulation).</p>
row.eff	<p>Single element indicating whether row effects are included as fixed effects ("fixed"), random effects ("random") or not included ("none") in the boral model. If random effects, they are drawn from a normal distribution with mean zero and standard deviation given by row.params. Defaults to "none".</p>
row.params	<p>Parameters corresponding to the row effect from the boral model. If row.eff = "fixed", then these are the fixed effects and should have length equal to the number of columns in y. If row.eff = "random", then this is the standard deviation for the random effects normal distribution. If row.eff = "none", then this argument is ignored.</p>
row.ids	<p>A matrix with the number of rows equal to the number of rows in y, and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row i in y for random effect eqnj; please see the help file for the main boral function for details. Defaults to NULL, so that if row.params = NULL then the argument is ignored, otherwise if row.params is supplied then row.ids = matrix(1:nrow(y), ncol = 1) i.e., a single, row effect unique to each row. An internal check is done to see row.params and row.ids are consistent in terms of arguments supplied.</p>
trial.size	<p>Either equal to a single element, or a vector of length equal to the number of columns in y. If a single element, then all columns assumed to be binomially distributed will have trial size set to this. If a vector, different trial sizes are allowed in each column of y. The argument is ignored for all columns not assumed to be binomially distributed. Defaults to 1, i.e. Bernoulli distribution.</p>
cutoffs	<p>A vector of common common cutoffs for proportional odds regression when any of family is ordinal. They should be increasing order. Defaults to NULL.</p>

powerparam	A common power parameter for tweedie regression when any of family is tweedie. Defaults to NULL.
manual.dim	A vector of length 2, containing the number of rows (n) and columns (p) for the multivariate response matrix. This is a "backup" argument only required when create.life can not determine how many rows or columns the multivariate response matrix should be.
save.params	If save.params = TRUE, then all parameters provided as input and/or generated (e.g. when traits and traits.coefs are supplied then X.coefs is generated internally; please see traits.coefs argument above) are returned, in addition to the simulated multivariate response matrix. Defaults to FALSE.
...	Not used.

Details

create.life gives the user full capacity to control the true parameters of the model from which the multivariate responses matrices are generated from.

simulate makes use of the generic function of the same name in R: it takes a fitted boral model, treats either the posterior medians and mean estimates from the model as the true parameters, and generates response matrices based off that.

Value

If create.life is used, then: 1) if save.params = FALSE, a n by p multivariate response matrix is returned only, 2) if save.params = TRUE, then a list containing the element resp which is a n times p multivariate response matrix, as well as other elements for the parameters used in the true model, e.g. true.lv, lv.coefs = lv.coefs, traits.coef, is returned.

If simulate is used, then a three dimensional array of dimension n by p by nsim is returned.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

See Also

[boral](#) for the default function for fitting a boral model.

Examples

```
## Example 1a - Simulate a response matrix of normally distributed data
library(mvtnorm)

## 30 rows (sites) with two latent variables
true.lv <- rbind(rmvnorm(n=15,mean=c(1,2)),rmvnorm(n=15,mean=c(-3,-1)))
## 30 columns (species)
lv.coefs <- cbind(matrix(runif(30*3),30,3),1)

X <- matrix(rnorm(30*4),30,4)
## 4 explanatory variables
X.coefs <- matrix(rnorm(30*4),30,4)
```

```

sim.y <- create.life(true.lv, lv.coefs, X, X.coefs, family = "normal")

## Not run:
fit.boral <- boral(sim.y, X = X, family = "normal", num.lv = 2)

summary(fit.boral)

## End(Not run)

## Example 1b - Include a nested random row effect
## 30 subregions nested within six regions
row.ids <- cbind(1:30, rep(1:6,each=5))
## Subregion has a small std deviation; region has a larger one
true.row.sigma <- list(ID1 = 0.5, ID2 = 2)

sim.y <- create.life(true.lv, lv.coefs, X, X.coefs, row.eff = "random",
row.params = true.row.sigma, row.ids = row.ids, family = "normal",
save.params = TRUE)

## Example 2 - Simulate a response matrix of ordinal data

## 30 rows (sites) with two latent variables
true.lv <- rbind(rmvnorm(15,mean=c(-2,-2)),rmvnorm(15,mean=c(2,2)))
## 10 columns (species)
true.lv.coefs <- rmvnorm(10,mean = rep(0,3));
## Impose a sum-to-zero constraint on the column effects
true.lv.coefs[nrow(true.lv.coefs),1] <- -sum(true.lv.coefs[-nrow(true.lv.coefs),1])
## Cutoffs for proportional odds regression (must be in increasing order)
true.ordinal.cutoffs <- seq(-2,10,length=10-1)

sim.y <- create.life(true.lv = true.lv, lv.coefs = true.lv.coefs,
family = "ordinal", cutoffs = true.ordinal.cutoffs, save.params = TRUE)

## Not run:
fit.boral <- boral(y = sim.y$resp, family = "ordinal", num.lv = 2)

## End(Not run)

## Not run:
## Example 3 - Simulate a response matrix of count data based off
## a fitted boral model involving traits (ants data from mvabund)
library(mvabund)
data(antTraits)

y <- antTraits$abun
X <- as.matrix(antTraits$env)
## Include only traits 1, 2, and 5, plus an intercept
traits <- as.matrix(antTraits$traits[,c(1,2,5)])
## Please see help file for boral regarding the use of which.traits
which.traits <- vector("list",ncol(X)+1)

```

```

for(i in 1:length(which.traits)) which.traits[[i]] <- 1:ncol(traits)

fit.traits <- boral(y, X = X, traits = traits, which.traits = which.traits,
family = "negative.binomial", num.lv = 2)

## The hard way
sim.y <- create.life(true.lv = NULL, lv.coefs = fit.traits$lv.coefs.median,
X = X, X.coefs = fit.traits$X.coefs.median,
traits = traits, traits.coefs = fit.traits$traits.coefs.median,
family = "negative.binomial")

## The easy way
sim.y <- simulate(object = fit.traits)

## End(Not run)

## Example 4 - simulate Bernoulli data, based on a model with two latent variables,
## no site variables, with two traits and one environmental covariates
## This example is a proof of concept that traits can used
## to explain environmental responses
library(mvtnorm)

n <- 100; s <- 50
X <- as.matrix(scale(1:n))
colnames(X) <- c("elevation")

traits <- cbind(rbinom(s,1,0.5), rnorm(s))
## one categorical and one continuous variable
colnames(traits) <- c("thorns-dummy", "SLA")

simfit <- list(true.lv = rmvnorm(n, mean = rep(0,2)),
lv.coefs = cbind(rnorm(s), rmvnorm(s, mean = rep(0,2))),
traits.coefs = matrix(c(0.1,1,-0.5,1,0.5,0,-1,1), 2, byrow = TRUE))
rownames(simfit$traits.coefs) <- c("beta0", "elevation")
colnames(simfit$traits.coefs) <- c("kappa0", "thorns-dummy", "SLA", "sigma")

simy = create.life(true.lv = simfit$true.lv, lv.coefs = simfit$lv.coefs, X = X,
traits = traits, traits.coefs = simfit$traits.coefs, family = "binomial")

## Not run:
which.traits <- vector("list",ncol(X)+1);
for(i in 1:length(which.traits)) which.traits[[i]] <- 1:ncol(traits)
fit.traits <- boral(y = simy, X = X, traits = traits, which.traits = which.traits,
family = "binomial", num.lv = 2, save.model = TRUE, calc.ics = FALSE)

## End(Not run)

```

`ds.residuals`*Dunn-Smyth Residuals for a boral model*

Description

Calculates the Dunn-Smyth residuals for a fitted boral model and, if some of the responses are ordinal, a confusion matrix between predicted and true levels.

Usage

```
ds.residuals(object, est = "median")
```

Arguments

<code>object</code>	An object for class "boral".
<code>est</code>	A choice of either the posterior median (<code>est == "median"</code>) or posterior mean (<code>est == "mean"</code>), which are then treated as parameter estimates and the residuals are calculated from. Default is posterior median.

Details

Details regarding Dunn-Smyth residuals, based on the randomized quantile residuals of Dunn and Smyth (1996), can be found in `plot.manyglm` function in the `mvabund` package (Wang et al., 2012) where they are implemented in all their glory. Due their inherent stochasticity, Dunn-Smyth residuals will be slightly different each time this function is run. As with other types of residuals, Dunn-Smyth residuals can be used in the context of residual analysis.

For ordinal responses, a single confusion matrix between the predicted levels (as based on the class with the highest probability) and true levels is also returned. The table pools the results over all columns assumed to be ordinal.

The Dunn-Smyth residuals are calculated based on a point estimate of the parameters, as determined by the argument `est`. A fully Bayesian approach would calculate the residuals by averaging over the posterior distribution of the parameters i.e., ergodically average over the MCMC samples. In general however, the results (as in the trends seen in residual analysis) from either approach should be very similar.

Value

A list containing `agree.ordinal` which is a single confusion matrix for ordinal columns, and `residuals` which contains Dunn-Smyth residuals.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Dunn, P. K., and Smyth, G. K. (1996). Randomized quantile residuals. *Journal of Computational and Graphical Statistics*, 5, 236-244.
- Wang, Y. et al. (2012). mvabund-an R package for model-based analysis of multivariate abundance data. *Methods in Ecology and Evolution*, 3, 471-474.

See Also

[plot.boral](#) for constructing residual analysis plots directly; [fitted.boral](#) which calculated fitted values from a boral model.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun

spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
  row.eff = "fixed")

ds.residuals(spider.fit.nb)

## End(Not run)
```

fitted.boral

Extract Model Fitted Values for an boral object

Description

Calculated the predicted mean responses based on the fitted boral model, by using the posterior medians or means of the parameters.

Usage

```
## S3 method for class 'boral'
fitted(object, est = "median",...)
```

Arguments

object	An object of class "boral".
est	A choice of either the posterior median (est == "median") or posterior mean (est == "mean"), which are then treated as estimates and the fitted values are calculated from. Default is posterior median.
...	Not used.

Details

This fitted values here are calculated based on a point estimate of the parameters, as determined by the argument `est`. A fully Bayesian approach would calculate the fitted values by averaging over the posterior distribution of the parameters i.e., ergodically average over the MCMC samples. For simplicity and speed though (to avoid generation of a large number of predicted values), this is not implemented.

Value

A list containing `ordinal.probs` which is an array with dimensions (no. of rows of `y`) x (no. of rows of `y`) x (no. of levels) containing the predicted probabilities for ordinal columns, and `out` which is a matrix of the same dimension as the original response matrix `y` containing the fitted values. For ordinal columns, the "fitted values" are defined as the level/class that had the highest fitted probability.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

See Also

[plot.boral](#) which uses the fitted values calculated from this function to construct plots for residual analysis; [ds.residuals](#) for calculating the Dunn-Smyth residuals for a fitted boral model.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun

spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
  row.eff = "fixed")

fitted(spider.fit.nb)

## End(Not run)
```

get.dic

Extract Deviance Information Criterion for boral model

Description

Calculates the Deviance Information Criterion (DIC) for a boral model fitted using JAGS.

Usage

```
get.dic(jagsfit)
```

Arguments

`jagsfit` The `jags.model` component of the output, from a model fitted using `boral` with `save.model = TRUE`.

Details

Details regarding the Deviance Information Criterion may be found in (Spiegelhalter et al., 2002; Ntzoufras, 2011; Gelman et al., 2013). The DIC here is based on the conditional log-likelihood i.e., the latent variables (and row effects if applicable) are treated as "fixed effects". A DIC based on the marginal likelihood is obtainable from [get.more.measures](#), although this requires a much longer time to compute. For models with overdispersed count data, conditional DIC may not perform as well as marginal DIC (Millar, 2009)

Value

DIC value for the jags model.

Note

This function and consequently the DIC value is automatically returned when a `boral` model is fitted using `boral` with `calc.ics = TRUE`.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Gelman et al. (2013). Bayesian data analysis. CRC press.
- Millar, R. B. (2009). Comparison of hierarchical Bayesian models for overdispersed count data using DIC and Bayes' factors. *Biometrics*, 65, 962-969.
- Ntzoufras, I. (2011). Bayesian modeling using WinBUGS (Vol. 698). John Wiley & Sons.
- Spiegelhalter, et al. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64, 583-639.

See Also

[get.measures](#) and [get.more.measures](#) for other information criteria which could potentially be used for variable selection.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
```



```

save.model = TRUE, calc.ics = TRUE)

spider.fit.nb$ics ## DIC returned as one of several information criteria.

## End(Not run)

```

get.enviro.cor	<i>Extract covariances and correlations due to shared environmental responses from boral models</i>
----------------	---

Description

Calculates the correlation between columns of the response matrix, due to similarities in the response to explanatory variables (i.e., shared environmental response)

Usage

```
get.enviro.cor(object, est = "median", prob = 0.95)
```

Arguments

object	An object for class "boral".
est	A choice of either the posterior median (est = "median") or posterior mean (est = "mean"), which are then treated as estimates and the fitted values are calculated from. Default is posterior median.
prob	A numeric scalar in the interval (0,1) giving the target probability coverage of the intervals, by which to determine whether the correlations are "significant". Defaults to 0.95.

Details

In both independent response and correlated response models, where each of the columns of the response matrix y are fitted to a set of explanatory variables given by X , the covariance and thus between two columns j and j' due to similarities in their response to the model matrix is calculated based on the linear predictors $x_i^T \beta_j$ and $x_i^T \beta_{j'}$, where β_j are column-specific coefficients relating to the explanatory variables (see also the help file for [boral](#)).

For multivariate abundance data, the correlation calculated by this function can be interpreted as the correlation attributable to similarities in the environmental response between species. Such correlation matrices are discussed and found in Ovaskainen et al., (2010), Pollock et al., 2014.

Value

A list with the following components:

cor	A $p \times p$ correlation matrix based on model matrix and the posterior or mean estimators of the associated regression coefficients.
-----	---

sig.cor	A $p \times p$ correlation matrix containing only the “significant” correlations whose 95% highest posterior interval does not contain zero. All non-significant correlations are zero to zero.
cov	A $p \times p$ covariance matrix based on model matrix and the posterior or mean estimators of the associated regression coefficients.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Ovaskainen et al. (2010). Modeling species co-occurrence by multivariate logistic regression generates new hypotheses on fungal interactions. *Ecology*, 91, 2514-2521.
- Pollock et al. (2014). Understanding co-occurrence by modelling species simultaneously with a Joint Species Distribution Model (JSDM). *Methods in Ecology and Evolution*, 5, 397-406.

See Also

[get.residual.cor](#), which calculates the residual correlation matrix for boral models involving latent variables.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
library(corrplot) ## For plotting correlations
data(spider)
y <- spider$abun
X <- scale(spider$x)
n <- nrow(y); p <- ncol(y);

spider.fit.nb <- boral(y, X = X, family = "negative.binomial",
  save.model = TRUE)

enviro.cors <- get.enviro.cor(spider.fit.nb)

corrplot(enviro.cors$sig.cor, title = "Shared response correlations",
  type = "lower", diag = FALSE, mar = c(3,0.5,2,1), tl.srt = 45)

## End(Not run)
```

get.hpdistervals

Highest posterior density intervals for an boral model

Description

Calculates the lower and upper bounds of the highest posterior density intervals for parameters and latent variables in a fitted boral model.

Usage

```
get.hpdintervals(y, X = NULL, traits = NULL, row.ids = NULL,
  fit.mcmc, num.lv, prob = 0.95)
```

Arguments

<code>y</code>	The response matrix that the boral model was fitted to.
<code>X</code>	The model matrix used in the boral model. Defaults to NULL, in which case it is assumed no model matrix was used.
<code>traits</code>	The matrix of species traits used in the boral model. Defaults to NULL, in which case it is assumed no traits were included.
<code>row.ids</code>	A matrix with the number of rows equal to the number of rows in <code>y</code> , and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row i in <code>y</code> for random effect eqnj; please see help file for the main boral function for details. Defaults to NULL, in which case it is assumed no random effects were included in the model.
<code>fit.mcmc</code>	All MCMC samples for the fitted boral model, as obtained from JAGS. These can be extracted by fitting an boral model using <code>boral</code> with <code>save.model = TRUE</code> , and then accessing the <code>jags.model</code> component of the output.
<code>num.lv</code>	The number of latent variables used in the boral model. If zero, then HPD intervals will not be produced for latent variables.
<code>prob</code>	A numeric scalar in the interval (0,1) giving the target probability coverage of the intervals. Defaults to 0.95.

Details

The function uses the `HPDinterval` function from the `coda` package to obtain the HPD intervals. See `HPDinterval` for details regarding the definition of the HPD interval.

Value

<code>lv.coefs.hpd.lower/upper</code>	Two matrices corresponding to the lower and upper bounds of the HPD intervals for the column-specific intercepts, latent variable coefficients, and dispersion parameters if appropriate.
<code>lv.hpd.lower/upper</code>	Two matrices corresponding to the lower and upper bounds of the HPD intervals for the latent variables.
<code>row.coefs.lower/upper</code>	A list with each element containing two vectors corresponding to the lower and upper bounds of the HPD intervals for row effects. The number of elements in the list should equal the number of row effects included in the model i.e., <code>ncol(row.ids)</code> .
<code>row.sigma.lower/upper</code>	A list with each element containing two scalars corresponding to the lower and upper bounds of the HPD interval for the standard deviation of the normal distribution for the row effects, if they were assumed to be random. The number

of elements in the list should equal the number of row effects included in the model i.e., `ncol(row.ids)`.

`X.coefs.hpd.lower/upper`

Two matrices corresponding to the lower and upper bounds of the HPD intervals for coefficients relating to the model matrix `X`.

`traits.coefs.hpd.lower/upper`

Two matrices corresponding to the lower and upper bounds of the HPD intervals for coefficients and standard deviation relating to the traits matrix `traits`.

`cutoffs.hpd.lower/upper`

Two vectors corresponding to the lower and upper bounds of the HPD intervals for common cutoffs in proportional odds regression.

`powerparam.hpd.lower/upper`

Two scalars corresponding to the lower and upper bounds of the HPD interval for common power parameter in tweedie regression.

Warnings

- HPD intervals tend to be quite wide, and inference is somewhat tricky with them. This is made more difficult by the multiple comparison problem due to the construction one interval for each parameter!
- Be very careful with interpretation of coefficients and HPD intervals if different columns of `y` have different distributions!
- HPD intervals for the cutoffs in proportional odds regression may be poorly estimated for levels with few data.

Note

`boral` fits the boral model and returns the HPD intervals by default.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

## Example 1 - model with two latent variables, site effects,
## and no environmental covariates
spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
  row.eff = "fixed", save.model = TRUE)

## Returns a list with components corresponding to values described above.
spider.fit.nb$hpdintervals
```

```
## Example 2 - model with two latent variable, site effects,
## and environmental covariates
spider.fit.nb2 <- boral(y, X = spider$x, family = "negative.binomial",
num.lv = 2, row.eff = "fixed", save.model = TRUE)

## Returns a list with components corresponding to values described above.
spider.fit.nb2$hpdiintervals

## End(Not run)
```

get.measures

*Information Criteria for boral models***Description**

Calculates some information criteria for an boral model, which could be used for model selection.

Usage

```
get.measures(y, X = NULL, family, trial.size = 1, row.eff = "none",
row.ids = NULL, num.lv, fit.mcmc)
```

Arguments

y	The response matrix that the boral model was fitted to.
X	The model matrix used in the boral model. Defaults to NULL, in which case it is assumed no model matrix was used.
family	Either a single element, or a vector of length equal to the number of columns in y. The former assumes all columns of y come from this distribution. The latter option allows for different distributions for each column of y. Elements can be one of "binomial" (with probit link), "poisson" (with log link), "negative.binomial" (with log link), "normal" (with identity link), "lnormal" for log-normal (with log link), "tweedie" (with log link), "exponential" (with log link), "gamma" (with log link), "beta" (with logit link), "ordinal" (cumulative probit regression).

For the negative binomial distribution, the variance is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the column-specific dispersion parameter. For the normal distribution, the variance is parameterized as $Var(y) = \phi^2$, where ϕ is the column-specific standard deviation. For the tweedie distribution, the variance is parameterized as $Var(y) = \phi\mu^p$ where ϕ is the column-specific dispersion parameter and p is a power parameter common to all columns assumed to be tweedie, with $1 < p < 2$. For the gamma distribution, the variance is parameterized as $Var(y) = \mu/\phi$ where ϕ is the column-specific rate (henceforth referred to also as dispersion parameter). For the beta distribution, the parameterization is in terms of the mean μ and sample size ϕ (henceforth referred to also as dispersion parameter), so that the two shape parameters are given by $a = \mu\phi$ and $b = (1 - \mu)\phi$.

	All columns assumed to have ordinal responses are constrained to have the same cutoffs points, with a column-specific intercept to account for differences between the columns (please see <i>Details</i> for formulation).
trial.size	Either equal to a single element, or a vector of length equal to the number of columns in <i>y</i> . If a single element, then all columns assumed to be binomially distributed will have trial size set to this. If a vector, different trial sizes are allowed in each column of <i>y</i> . The argument is ignored for all columns not assumed to be binomially distributed. Defaults to 1, i.e. Bernoulli distribution.
row.eff	Single element indicating whether row effects are included as fixed effects ("fixed"), random effects ("random") or not included ("none") in the boral model. If random effects, they are drawn from a normal distribution with mean zero and unknown standard deviation. Defaults to "none".
row.ids	A matrix with the number of rows equal to the number of rows in <i>y</i> , and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row <i>i</i> in <i>y</i> for random effect eqnj; please see the help file for the main boral function for details. Defaults to NULL, so that if row.eff = "none" then the argument is ignored, otherwise if row.eff = "fixed" or "random", then row.ids = matrix(1:nrow(y), ncol = 1) i.e., a single, row effect unique to each row.
num.lv	The number of latent variables used in the fitted boral model.
fit.mcmc	All MCMC samples for the fitted boral model, as obtained from JAGS. These can be extracted by fitting an boral model using <code>boral</code> with <code>save.model = TRUE</code> , and then accessing the <code>jags.model</code> component of the output.

Details

The following information criteria are currently calculated, when permitted: 1) Widely Applicable Information Criterion (WAIC, Watanabe, 2010) based on the conditional log-likelihood; 2) expected AIC (EAIC, Carlin and Louis, 2011); 3) expected BIC (EBIC, Carlin and Louis, 2011); 4) AIC (using the marginal likelihood) evaluated at the posterior median; 5) BIC (using the marginal likelihood) evaluated at the posterior median.

1) WAIC has been argued to be more natural and extension of AIC to the Bayesian and hierarchical modeling context (Gelman et al., 2013), and is based on the conditional log-likelihood calculated at each of the MCMC samples.

2 & 3) EAIC and EBIC were suggested by (Carlin and Louis, 2011). Both criteria are of the form $-2 * \text{mean}(\text{conditional log-likelihood}) + \text{penalty} * (\text{no. of parameters in the model})$, where the mean is averaged all the MCMC samples. EAIC applies a penalty of 2, while EBIC applies a penalty of $\log(n)$.

4 & 5) AIC and BIC take the form $-2 * (\text{marginal log-likelihood}) + \text{penalty} * (\text{no. of parameters in the model})$, where the log-likelihood is evaluated at the posterior median. If the parameter-wise posterior distributions are unimodal and approximately symmetric, these will produce similar results to an AIC and BIC where the log-likelihood is evaluated at the posterior mode. EAIC applies a penalty of 2, while EBIC applies a penalty of $\log(n)$.

Intuitively, comparing boral models with and without latent variables (using information criteria such as those returned) amounts to testing whether the columns of the response matrix *y* are correlated. With multivariate abundance data for example, where *y* is a matrix of *n* sites and *p* species,

comparing models with and without latent variables tests whether there is any evidence of correlation between species.

Please note that criteria 4 and 5 are not calculated all the time. In models where traits are included in the model (such that the regression coefficients β_{0j}, β_j are random effects), more than two columns are ordinal responses (such that the intercepts β_{0j} for these columns are random effects), or more than one random row effect is included (such that `row.ids` contains more than one column), then criteria 4 and 5 are will not calculated. This is because the calculation of the marginal log-likelihood in such cases currently fail to marginalize over such random effects; see the details in the help files for `calc.logLik.lv0` and `calc.marglogLik`.

Value

A list with the following components:

<code>waic</code>	WAIC based on the conditional log-likelihood.
<code>eaic</code>	EAIC based on the mean of the conditional log-likelihood.
<code>ebic</code>	EBIC based on the mean of the conditional log-likelihood.
<code>all.cond.logLik</code>	The conditional log-likelihood evaluated at all MCMC samples. This is done via repeated application of <code>calc.condlogLik</code> .
<code>cond.num.params</code>	Number of estimated parameters used in the fitted model, when all parameters are treated as "fixed" effects.
<code>do.marglik.ics</code>	A boolean indicating whether marginal log-likelihood based information criteria are calculated.

If `do.marglik.ics = TRUE`, then we also have:

<code>aic.median</code>	AIC (using the marginal log-likelihood) evaluated at the posterior median.
<code>bic.median</code>	BIC (using the marginal log-likelihood) evaluated at the posterior median.
<code>marg.num.params</code>	Number of estimated parameters used in the fitted model, when all parameters are treated as "fixed" effects.

Warning

Using information criterion for variable selection should be done with extreme caution, for two reasons: 1) The implementation of these criteria are both *heuristic* and experimental. 2) Deciding what model to fit for ordination purposes should be driven by the science. For example, it may be the case that a criterion suggests a model with 3 or 4 latent variables. However, if we interested in visualizing the data for ordination purposes, then models with 1 or 2 latent variables are far more appropriate. As an another example, whether or not we include row effects when ordinating multivariate abundance data depends on if we are interested in differences between sites in terms of relative species abundance (`row.eff = FALSE`) or in terms of species composition (`row.eff = "fixed"`).

Also, the use of information criterion in the presence of variable selection using SSVS is questionable.

Note

When a boral model is fitted using `boral` with `calc.ics = TRUE`, then this function is applied and the information criteria are returned as part of the model output.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Carlin, B. P., and Louis, T. A. (2011). Bayesian methods for data analysis. CRC Press.
- Gelman et al. (2013). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 1-20.
- Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *The Journal of Machine Learning Research*, 11, 3571-3594.

See Also

[get.dic](#) for calculating the Deviance Information Criterion (DIC) based on the conditional log-likelihood; [get.more.measures](#) for even more information criteria.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

spider.fit.pois <- boral(y, family = "poisson",
num.lv = 2, row.eff = "random")

spider.fit.pois$ics ## Returns information criteria

spider.fit.nb <- boral(y, family = "negative.binomial",
num.lv = 2, row.eff = "random")

spider.fit.nb$ics ## Returns the information criteria

## End(Not run)
```

get.more.measures *Additional Information Criteria for boral models*

Description

Calculates some information criteria beyond those from `get.measures` for a boral model, although this set of criteria takes much longer to compute!

Usage

```
get.more.measures(y, X = NULL, family, trial.size = 1,
  row.eff = "none", row.ids = NULL, num.lv, fit.mcmc,
  verbose = TRUE)
```

Arguments

- | | |
|------------|--|
| y | The response matrix that the boral model was fitted to. |
| X | The model matrix used in the boral model. Defaults to NULL, in which case it is assumed no model matrix was used. |
| family | <p>Either a single element, or a vector of length equal to the number of columns in <i>y</i>. The former assumes all columns of <i>y</i> come from this distribution. The latter option allows for different distributions for each column of <i>y</i>. Elements can be one of "binomial" (with probit link), "poisson" (with log link), "negative.binomial" (with log link), "normal" (with identity link), "lnormal" for log-normal (with log link), "tweedie" (with log link), "exponential" (with log link), "gamma" (with log link), "beta" (with logit link), "ordinal" (cumulative probit regression).</p> <p>For the negative binomial distribution, the variance is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the column-specific dispersion parameter. For the normal distribution, the variance is parameterized as $Var(y) = \phi^2$, where ϕ is the column-specific standard deviation. For the tweedie distribution, the variance is parameterized as $Var(y) = \phi\mu^p$ where ϕ is the column-specific dispersion parameter and p is a power parameter common to all columns assumed to be tweedie, with $1 < p < 2$. For the gamma distribution, the variance is parameterized as $Var(y) = \mu/\phi$ where ϕ is the column-specific rate (henceforth referred to also as dispersion parameter). For the beta distribution, the parameterization is in terms of the mean μ and sample size ϕ (henceforth referred to also as dispersion parameter), so that the two shape parameters are given by $a = \mu\phi$ and $b = (1 - \mu)\phi$.</p> <p>All columns assumed to have ordinal responses are constrained to have the same cutoffs points, with a column-specific intercept to account for differences between the columns (please see <i>Details</i> for formulation).</p> |
| trial.size | <p>Either equal to a single element, or a vector of length equal to the number of columns in <i>y</i>. If a single element, then all columns assumed to be binomially distributed will have trial size set to this. If a vector, different trial sizes are allowed in each column of <i>y</i>. The argument is ignored for all columns not assumed to be binomially distributed. Defaults to 1, i.e. Bernoulli distribution.</p> |

row.eff	Single element indicating whether row effects are included as fixed effects ("fixed"), random effects ("random") or not included ("none") in the boral model. If random effects, they are drawn from a normal distribution with mean zero and unknown standard deviation. Defaults to "none".
row.ids	A matrix with the number of rows equal to the number of rows in y , and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row i in y for random effect eqnj; please see the help file for the main boral function for details. Defaults to NULL, so that if row.eff = "none" then the argument is ignored, otherwise if row.eff = "fixed" or "random", then row.ids = matrix(1:nrow(y), ncol = 1) i.e., a single, row effect unique to each row.
num.lv	The number of latent variables used in the fitted boral model.
fit.mcmc	All MCMC samples for the fitted boral model, as obtained from JAGS. These can be extracted by fitting an boral model using <code>boral</code> with <code>save.model = TRUE</code> , and then accessing the <code>jags.model</code> component of the output.
verbose	If TRUE, a notice is printed every 100 samples indicating progress in calculation of the marginal log-likelihood. Defaults to TRUE.

Details

Currently, four information criteria are calculated using this function, when permitted: 1) AIC (using the marginal likelihood) evaluated at the posterior mode; 2) BIC (using the marginal likelihood) evaluated at the posterior mode; 3) Deviance information criterion (DIC) based on the marginal log-likelihood; 4) Widely Applicable Information Criterion (WAIC, Watanabe, 2010) based on the marginal log-likelihood. Since flat priors are used in fitting boral models, then the posterior mode should be approximately equal to the maximum likelihood estimates.

All four criteria require computing the marginal log-likelihood across all MCMC samples. This takes a very long time to run, since Monte Carlo integration needs to be performed for all MCMC samples. Consequently, this function is currently not implemented as an argument in main `boral` fitting function, unlike `get.measures` which is available via the `calc.ics = TRUE` argument.

Moreover, note these criteria are not calculated all the time. In models where traits are included in the model (such that the regression coefficients β_{0j}, β_j are random effects), more than two columns are ordinal responses (such that the intercepts β_{0j} for these columns are random effects), or more than one random row effect is included (such that row.ids contains more than one column), then these extra information criteria are will not calculated, and the function returns nothing except a simple message. This is because the calculation of the marginal log-likelihood in such cases currently fail to marginalize over such random effects; see the details in the help files for `calc.logLik.lv0` and `calc.marglogLik`.

The two main differences between the criteria and those returned from `get.measures` are:

- The AIC and BIC computed here are based on the log-likelihood evaluated at the posterior mode, whereas the AIC and BIC from `get.measures` are evaluated at the posterior median. The posterior mode and median will be quite close to one another if the component-wise posterior distributions are unimodal and symmetric. Furthermore, given uninformative priors are used, then both will be approximate maximum likelihood estimators.
- The DIC and WAIC computed here are based on the marginal log-likelihood, whereas the DIC and WAIC from `get.measures` are based on the conditional log-likelihood. Criteria based on

the two types of log-likelihood are equally valid, and to a certain extent, which one to use depends on the question being answered i.e., whether to condition on the latent variables or treat them as "random effects" (see discussions in Spiegelhalter et al. 2002, and Vaida and Blanchard, 2005).

Value

If calculated, then a list with the following components:

<code>marg.aic</code>	AIC (using on the marginal log-likelihood) evaluated at posterior mode.
<code>marg.bic</code>	BIC (using on the marginal log-likelihood) evaluated at posterior mode.
<code>marg.dic</code>	DIC based on the marginal log-likelihood.
<code>marg.waic</code>	WAIC based on the marginal log-likelihood.
<code>all.marg.logLik</code>	The marginal log-likelihood evaluated at all MCMC samples. This is done via repeated application of <code>calc.marglogLik</code> .
<code>num.params</code>	Number of estimated parameters used in the fitted model.

Warning

Using information criterion for variable selection should be done with extreme caution, for two reasons: 1) The implementation of these criteria are both *heuristic* and experimental. 2) Deciding what model to fit for ordination purposes should be driven by the science. For example, it may be the case that a criterion suggests a model with 3 or 4 latent variables. However, if we interested in visualizing the data for ordination purposes, then models with 1 or 2 latent variables are far more appropriate. As an another example, whether or not we include row effects when ordinating multivariate abundance data depends on if we are interested in differences between sites in terms of relative species abundance (`row.eff = FALSE`) or in terms of species composition (`row.eff = "fixed"`).

Also, the use of information criterion in the presence of variable selection using SSVS is questionable.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Spiegelhalter et al. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64, 583-639.
- Vaida, F., and Blanchard, S. (2005). Conditional Akaike information for mixed-effects models. *Biometrika*, 92, 351-370.
- Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *The Journal of Machine Learning Research*, 11, 3571-3594.

See Also

[get.measures](#) for several information criteria which take considerably less time to compute, and are automatically implemented in [boral](#) with `calc.ics = TRUE`.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
  row.eff = "fixed", save.model = TRUE, calc.ics = TRUE)

## Extract MCMC samples
fit.mcmc <- mcmc(spider.fit.nb$jags.model$BUGSoutput$sims.matrix)

## WATCH OUT! The following takes a very long time to run!
get.more.measures(y, family = "negative.binomial",
  num.lv = spider.fit.nb$num.lv, fit.mcmc = fit.mcmc,
  row.eff = "fixed", row.ids = spider.fit.nb$row.ids)

## Illustrating what happens in a case where these criteria will
## not be calculated.
data(antTraits)
y <- antTraits$abun
X <- as.matrix(scale(antTraits$env))
## Include only traits 1, 2, and 5
traits <- as.matrix(antTraits$traits[,c(1,2,5)])
which.traits <- vector("list",ncol(X)+1)
for(i in 1:length(which.traits)) which.traits[[i]] <- 1:ncol(traits)

fit.traits <- boral(y, X = X, traits = traits, num.lv = 2,
  which.traits = which.traits, family = "negative.binomial",
  calc.ics = FALSE, save.model = TRUE)

## Extract MCMC samples
fit.mcmc <- mcmc(fit.traits$jags.model$BUGSoutput$sims.matrix)

get.more.measures(y, X = X, family = "negative.binomial",
  num.lv = fit.traits$num.lv, fit.mcmc = fit.mcmc)

## End(Not run)
```

Description

Calculates the residual correlation from models that include latent variables.

Usage

```
get.residual.cor(object, est = "median", prob = 0.95)
```

Arguments

object	An object for class "boral".
est	A choice of either the posterior median (est = "median") or posterior mean (est = "mean"), which are then treated as estimates and the fitted values are calculated from. Default is posterior median.
prob	A numeric scalar in the interval (0,1) giving the target probability coverage of the intervals, by which to determine whether the correlations are "significant". Defaults to 0.95.

Details

In models with latent variables, the residual covariance matrix is calculated based on the matrix of latent variables regression coefficients formed by stacking the rows of θ_j . That is, if we denote $\Theta = (\theta_1 \dots \theta_p)'$, then the residual covariance and hence residual correlation matrix is calculated based on $\Theta\Theta'$ (see also the help file for [boral](#)).

For multivariate abundance data, the inclusion of latent variables provides a parsimonious method of accounting for correlation between species. Specifically, the linear predictor,

$$\beta_{0j} + \mathbf{x}_i^T \beta_j + \mathbf{z}_i^T \theta_j$$

is normally distributed with a residual covariance matrix given by $\Theta\Theta'$. A strong residual covariance/correlation matrix between two species can then be interpreted as evidence of species interaction (e.g., facilitation or competition), missing covariates, as well as any additional species correlation not accounted for by shared environmental responses (see also Pollock et al., 2014, for residual correlation matrices in the context of Joint Species Distribution Models).

In addition to the residual correlation matrix, the median or mean point estimator of trace of the residual covariance matrix is returned, $\sum_{j=1}^p [\Theta\Theta']_{jj}$. Often used in other areas of multivariate statistics, the trace may be interpreted as the amount of covariation explained by the latent variables. One situation where the trace may be useful is when comparing a pure LVM versus a model with latent variables and some predictors (correlated response models) – the proportional difference in trace between these two models may be interpreted as the proportion of covariation between species explained by the predictors. Of course, the trace itself is random due to the MCMC sampling, and so it is not always guaranteed to produce sensible answers =P

Value

A list with the following components:

cor	A $p \times p$ residual correlation matrix based on posteriori median or mean estimators of the latent variables and coefficients.
sig.cor	A $p \times p$ correlation matrix containing only the “significant” correlations whose 95% highest posterior interval does not contain zero. All non-significant correlations are zero to zero.
cov	A $p \times p$ covariance correlation matrix based on posteriori median or mean estimators of the latent variables and coefficients.
trace	The median/mean point estimator of the trace (sum of the diagonal elements) of the residual covariance matrix.

Note

Residual correlation matrices are reliably modeled only with two or more latent variables i.e., `num.lv > 1` when fitting the model using `boral`.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Pollock et al. (2014). Understanding co-occurrence by modelling species simultaneously with a Joint Species Distribution Model (JSDM). *Methods in Ecology and Evolution*, 5, 397-406.

See Also

[get.enviro.cor](#), which calculates the correlation matrix due to similarities in the response to the explanatory variables (i.e., similarities due to a shared environmental response).

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
library(corrplot) ## For plotting correlations
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

spider.fit.nb <- boral(y, X = spider$x, family = "negative.binomial",
num.lv = 2, save.model = TRUE)

res.cors <- get.residual.cor(spider.fit.nb)

corrplot(res.cors$sig.cor, title = "Residual correlations",
type = "lower", diag = FALSE, mar = c(3,0.5,2,1), tl.srt = 45)

## End(Not run)
```

lvplot	<i>Plot the latent variables from an boral model</i>
--------	--

Description

Construct a 1-D index plot or 2-D scatterplot of the latent variables, and their corresponding coefficients i.e., a biplot, from a fitted boral model.

Usage

```
lvplot(x, jitter = FALSE, biplot = TRUE, ind.spp = NULL, alpha = 0.5,
main = NULL, est = "median", which.lvs = c(1,2), return.vals = FALSE, ...)
```

Arguments

x	An object for class "boral".
jitter	If <code>jitter = TRUE</code> , then some jittering is applied so that points on the plots do not overlap exactly (which can often occur with discrete data, small sample sizes, and if some sites are identical in terms species co-occurrence). Please see jitter for its implementation. Defaults to FALSE.
biplot	If <code>biplot = TRUE</code> , then a biplot is construct such that both the latent variables <i>and</i> their corresponding coefficients are plotted. Otherwise, only the latent variable scores are plotted. Defaults to TRUE.
ind.spp	Controls the number of latent variable coefficients to plot if <code>biplot = TRUE</code> . If <code>ind.spp</code> is an integer, then only the first <code>ind.spp</code> "most important" latent variable coefficients are included in the biplot, where "most important" means the latent variable coefficients with the largest L2-norms. Defaults to NULL, in which case all latent variable coefficients are included in the biplot.
alpha	A numeric scalar between 0 and 1 that is used to control the relative scaling of the latent variables and their coefficients, when constructing a biplot. Defaults to 0.5, and we typically recommend between 0.45 to 0.55 so that the latent variables and their coefficients are on roughly the same scale.
main	Title for resulting ordination plot. Defaults to NULL, in which case a "standard" title is used.
est	A choice of either the posterior median (<code>est = "median"</code>) or posterior mean (<code>est = "mean"</code>), which are then treated as estimates and the ordinations based off. Default is posterior median.
which.lvs	A vector of length two, indicating which latent variables (ordination axes) to plot which x is an object with two or more latent variables. The argument is ignored if x only contains one latent variables. Defaults to <code>which.lvs = c(1,2)</code> .
return.vals	If <code>return.vals = TRUE</code> , then the <i>scaled</i> latent variables scores and corresponding scaled coefficients are returned (based on the value of alpha used). This is useful if the user wants to construct their own custom model-based ordinations. Defaults to FALSE.

... Additional graphical options to be included in. These include values for `cex`, `cex.lab`, `cex.axis`, `cex.main`, `lwd`, and so on.

Details

This function allows an ordination plot to be constructed, based on either the posterior medians and posterior means of the latent variables respectively depending on the choice of `est`. The latent variables are labeled using the row index of the response matrix `y`. If the fitted model contains more than two latent variables, then one can specify which latent variables i.e., ordination axes, to plot based on the `which.lvs` argument. This can prove useful (to check) if certain sites are outliers on one particular ordination axes.

If the fitted model did not contain any covariates, the ordination plot can be interpreted in the exactly same manner as unconstrained ordination plots constructed from methods such as Nonmetric Multi-dimensional Scaling (NMDS, Kruskal, 1964) and Correspondence Analysis (CA, Hill, 1974). With multivariate abundance data for instance, where the response matrix `y` consists of n sites and p species, the ordination plots can be studied to look for possible clustering of sites, location and/or dispersion effects, an arch pattern indicative of some sort species succession over an environmental gradient, and so on.

If the fitted model did include covariates, then a "residual ordination" plot is produced, which can be interpreted as offering a graphical representation of the (main patterns of) residual covariations, i.e. covariations after accounting for the covariates. With multivariate abundance data for instance, these residual ordination plots represent could represent residual species co-occurrence due to phylogeny, species competition and facilitation, missing covariates, and so on (Warton et al., 2015)

If `biplot = TRUE`, then a biplot is constructed so that both the latent variables and their corresponding coefficients are included in their plot (Gabriel, 1971). The latent variable coefficients are shown in red, and are indexed by the column names of `y`. The number of latent variable coefficients to plot is controlled by `ind.spp`. In ecology for example, often we are only be interested in the "indicator" species, e.g. the species with most represent a particular set of sites or species with the strongest covariation (see Chapter 9, Legendre and Legendre, 2012, for additional discussion). In such case, we can then biplot only the `ind.spp` "most important" species, as indicated by the the L2-norm of their latent variable coefficients.

As with correspondence analysis, the relative scaling of the latent variables and the coefficients in a biplot is essentially arbitrary, and could be adjusted to focus on the sites, species, or put even weight on both (see Section 9.4, Legendre and Legendre, 2012). In `lvsp1ot`, this relative scaling is controlled by the `alpha` argument, which basically works by taking the latent variables to a power `alpha` and the latent variable coefficients to a power `1-alpha`.

For latent variable models, we are generally interested in "symmetric plots" that place the latent variables and their coefficients on the same scale. In principle, this is achieved by setting `alpha = 0.5`, the default value, although sometimes this needs to be tweaked slightly to a value between 0.45 and 0.55 (see also the `corresp` function in the MASS package that also produces symmetric plots, as well as Section 5.4, Borcard et al., 2011 for more details on scaling).

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Borcard et al. (2011). Numerical Ecology with R. Springer.
- Gabriel, K. R. (1971). The biplot graphic display of matrices with application to principal component analysis. *Biometrika*, 58, 453-467.
- Hill, M. O. (1974). Correspondence analysis: a neglected multivariate method. *Applied statistics*, 23, 340-354.
- Kruskal, J. B. (1964). Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, 29, 115-129.
- Legendre, P. and Legendre, L. (2012). Numerical ecology, Volume 20. Elsevier.
- Warton et al. (2015). So Many Variables: Joint Modeling in Community Ecology. *Trends in Ecology and Evolution*, in review.

Examples

```
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

## NOTE: The values below MUST NOT be used in a real application;
## they are only used here to make the examples run quick!!!
example.mcmc.control <- list(n.burnin = 10, n.iteration = 100,
                             n.thin = 1)

spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
                      row.eff = "fixed", calc.ics = FALSE,
                      mcmc.control = example.mcmc.control)

lvplot(spider.fit.nb)
```

make.jagsboralmode1 *Write a text file containing an boral model for use into JAGS*

Description

This function is designed to write boral models with one or more latent variables.

Usage

```
make.jagsboralmode1(family, num.X = 0, num.traits = 0,
                    which.traits = NULL, num.lv = 2, row.eff = "none", row.ids = NULL,
                    trial.size = 1, n, p, model.name = NULL,
                    prior.control = list(type = c("normal", "normal", "normal", "uniform"),
                                         hypparams = c(100, 20, 100, 50), ssvs.index = -1, ssvs.g = 1e-6))
```

Arguments

family	<p>Either a single element, or a vector of length equal to the number of columns in y. The former assumes all columns of y come from this distribution. The latter option allows for different distributions for each column of y. Elements can be one of "binomial" (with probit link), "poisson" (with log link), "negative.binomial" (with log link), "normal" (with identity link), "lnormal" for log-normal (with log link), "tweedie" (with log link), "exponential" (with log link), "gamma" (with log link), "beta" (with logit link), "ordinal" (cumulative probit regression).</p> <p>For the negative binomial distribution, the variance is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the column-specific dispersion parameter. For the normal distribution, the variance is parameterized as $Var(y) = \phi^2$, where ϕ is the column-specific standard deviation. For the tweedie distribution, the variance is parameterized as $Var(y) = \phi\mu^p$ where ϕ is the column-specific dispersion parameter and p is a power parameter common to all columns assumed to be tweedie, with $1 < p < 2$. For the gamma distribution, the variance is parameterized as $Var(y) = \mu/\phi$ where ϕ is the column-specific rate (henceforth referred to also as dispersion parameter). For the beta distribution, the parameterization is in terms of the mean μ and sample size ϕ (henceforth referred to also as dispersion parameter), so that the two shape parameters are given by $a = \mu\phi$ and $b = (1 - \mu)\phi$.</p> <p>All columns assumed to have ordinal responses are constrained to have the same cutoffs points, with a column-specific intercept to account for differences between the columns (please see <i>Details</i> for formulation).</p>
num.X	Number of columns in the model matrix X . Defaults to 0, in which case it is assumed that no covariates are included in the model. Recall that no intercept should be included in X .
num.traits	Number of columns in the model matrix <code>traits</code> . Defaults to 0, in which case it is assumed no traits are included in model. Recall that no intercept should be included in <code>traits</code> .
which.traits	<p>A list of length equal to (number of columns in $X + 1$), informing which columns of <code>traits</code> the column-specific intercepts and each of the column-specific regression coefficients should be regressed against. The first element in the list applies to the column-specific intercept, while the remaining elements apply to the regression coefficients. Each element of <code>which.traits</code> is a vector indicating which traits are to be used.</p> <p>For example, if <code>which.traits[[2]] = c(2, 3)</code>, then the regression coefficients corresponding to the first column in X are regressed against the second and third columns of <code>traits</code>. If <code>which.traits[[2]] = 0</code>, then the regression coefficients for each column are treated as independent. Please see help file below for more details.</p> <p>Defaults to NULL, in conjunction with <code>traits = NULL</code>.</p>
num.lv	Number of latent variables to fit. Can take any non-negative integer value. Defaults to 2.
row.eff	Single element indicating whether row effects are included as fixed effects ("fixed"), random effects ("random") or not included ("none") in the boralm model. If ran-

dom effects, they are drawn from a normal distribution with mean zero and unknown standard deviation. Defaults to "none".

row.ids	A matrix with the number of rows equal to the number of rows in <i>y</i> , and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row <i>i</i> in <i>y</i> for random effect eqnj; please see the help file for the main boral function for details. Defaults to NULL, so that if row.eff = "none" then the argument is ignored, otherwise if row.eff = "fixed" or "random", then row.ids = matrix(1:nrow(y), ncol = 1) i.e., a single, row effect unique to each row.
trial.size	Either equal to a single element, or a vector of length equal to the number of columns in <i>y</i> . If a single element, then all columns assumed to be binomially distributed will have trial size set to this. If a vector, different trial sizes are allowed in each column of <i>y</i> . The argument is ignored for all columns not assumed to be binomially distributed. Defaults to 1, i.e. Bernoulli distribution.
n	The number of rows in the response matrix <i>y</i> .
p	The number of columns in the response matrix <i>y</i> .
model.name	Name of the text file that the JAGS model is written to. Defaults to NULL, in which case the default of "jagsboralmode.txt" is used.
prior.control	A list of parameters for controlling the prior distributions. These include: <ul style="list-style-type: none"> • <i>type</i>: Vector of four strings indicating the type of prior distributions to use. In order, these are: 1) priors for all column-specific intercepts, row effects, and cutoff points for ordinal data; 2) priors for the latent variable coefficients (ignored if num.lv = 0); 3) priors for all column-specific coefficients relating to the model matrix <i>X</i> (ignored if <i>X</i> = NULL). When traits are included in the model, this is also the prior for the trait regression coefficients (please see section <i>Including species traits</i> in the help file for boral for more information); 4) priors for any dispersion parameters. For elements 1-3, the prior distributions currently available include: I) "normal", which is normal prior with the variance controlled by the hypparams argument; II) "cauchy", which is a Cauchy prior with variance controlled by the hypparams argument. Gelman, et al. (2008) considers using Cauchy priors with variance 2.5^2 as weakly informative priors for regression coefficients; III) "uniform", which is uniform prior with minimum values given by -hypparams and maximum values given by +hypparams. For element 4, the prior distributions currently available include: I) "uniform", which is uniform prior with minimum zero and maximum controlled by hypparams[4]; II) "halfnormal", which is half-normal prior with variance controlled by hypparams[4]; III) "halfcauchy", which is a half-Cauchy prior with variance controlled by the hypparams[4] argument. Defaults to the vector c("normal", "normal", "normal", "uniform"). • <i>hypparams</i>: Vector of four hyperparameters used in the set up of prior distributions. In order, these are: 1) affects the prior distribution for all column-specific intercepts, row effects, and cutoff points for ordinal data. If row.eff = "random", this also controls the maximum of the uniform prior for the standard deviation of the random effects normal distribution. Also, if more than two of the columns are ordinal, then this also controls

the maximum of the uniform prior for the standard deviation of the column-specific random intercepts for these columns; 2) affects the prior distribution for all latent variable coefficients (ignored if `num.lv = 0`); 3) affects the prior distribution for column-specific coefficients relating to the model matrix `X` (ignored if `X = NULL`). When traits are included in the model, it also affects the prior distribution for the trait regression coefficients, and controls the maximum of the uniform prior for the standard deviation of the normally distributed random effects; 4) affects the prior distribution used for any dispersion parameters.

Defaults to the vector `c(100, 20, 100, 50)`.

- `ssvs.index`: Indices to be used for stochastic search variable selection (SSVS, George and McCulloch, 1993). Either a single element or a vector with length equal to the number of columns in the implied model matrix `X`. Each element can take values of -1 (no SSVS is performed on this covariate), 0 (SSVS is performed on individual coefficients for this covariate), or any integer exceeding 1 (SSVS is performed on collectively all coefficients on this covariate/s.)

This argument is only read if `X.eff = TRUE`. Please see the [boral](#) help file for more information regarding the implementation of SSVS. Defaults to -1, in which case no model selection is performed on the fitted model at all.

- `ssvs.g`: Multiplicative, shrinkage factor for SSVS, which controls the strength of the "spike" in the SSVS mixture prior. In summary, if the coefficient is included in the model, the "slab" prior is a normal distribution with mean zero and variance given by `hypparams`, while if the coefficient is not included in the model, the "spike" prior is normal distribution with mean zero and variance given by `hypparams*ssvs.g`. Please see the [boral](#) help file for more information regarding the implementation of SSVS. Defaults to `1e-6`.

Details

This function is automatically executed inside [boral](#), and therefore does not need to be run separately before fitting the boral model. It can however be run independently if one is: 1) interested in what the actual JAGS file for a particular boral model looks like, 2) wanting to modify a basic JAGS model file to construct more complex model e.g., include environmental variables.

Please note that [boral](#) currently does not allow the user to manually enter a script to be run.

When running the main function [boral](#), setting `save.model = TRUE` which automatically save the JAGS model file as a text file (with name based on the `model.name`) in the current working directory.

Value

A text file is created, containing the JAGS model to be called by the [boral](#) function for entering into `jags`. This file is automatically deleted once [boral](#) has finished running `save.model = TRUE`.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Gelman, et al. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2, 1360-1383.

See Also

[make.jagsboralnullmodel](#) for writing boral models JAGS scripts with no latent variables (so-called "null models").

Examples

```
library(mvtnorm)
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

## Example 1 - Create a boral model JAGS script, where distributions alternative
## between Poisson and negative binomial distributions
## across the rows of y.
make.jagsboralmodel(family = rep(c("poisson", "negative.binomial"), length=p),
row.eff = "fixed", num.X = 0, n = n, p = p)

## Example 2 - Create a boral model JAGS script, where distributions are all
## negative binomial distributions and covariates will be included.
make.jagsboralmodel(family = "negative.binomial", num.X = ncol(spider$x),
n = n, p = p)

## Example 3 - Simulate some ordinal data and create a JAGS model script
## 30 rows (sites) with two latent variables
true.lv <- rbind(rmvnorm(15, mean=c(-2, -2)), rmvnorm(15, mean=c(2, 2)))
## 10 columns (species)
true.lv.coefs <- rmvnorm(10, mean = rep(0, 3));
true.lv.coefs[nrow(true.lv.coefs), 1] <- -sum(true.lv.coefs[-nrow(true.lv.coefs), 1])
## Impose a sum-to-zero constraint on the column effects
true.ordinal.cutoffs <- seq(-2, 10, length=10-1)

sim.y <- create.life(true.lv = true.lv, lv.coefs = true.lv.coefs,
family = "ordinal", cutoffs = true.ordinal.cutoffs)

make.jagsboralmodel(family = "ordinal", num.X = 0,
row.eff = FALSE, n=30, p=10, model.name = "myawesomeordmodel.txt")

## Have a look at the JAGS model file for a boral model involving traits,
## based on the ants data from mvabund.
library(mvabund)
data(antTraits)

y <- antTraits$abun
X <- as.matrix(antTraits$env)
```

```
## Include only traits 1, 2, and 5, plus an intercept
traits <- as.matrix(antTraits$traits[,c(1,2,5)])
## Please see help file for boral regarding the use of which.traits
which.traits <- vector("list",ncol(X)+1)
for(i in 1:length(which.traits)) which.traits[[i]] <- 1:ncol(traits)

## Not run:
fit.traits <- boral(y, X = X, traits = traits, which.traits = which.traits,
family = "negative.binomial", num.lv = 2, model.name = "anttraits.txt")

## End(Not run)
```

```
make.jagsboralnullmodel
```

Write a text file containing an boral model for use into JAGS

Description

This function is designed to write boral models with no latent variables (so-called "null" models).

Usage

```
make.jagsboralnullmodel(family, num.X = 0, num.traits = 0,
  which.traits = NULL, row.eff = "none", row.ids = NULL,
  trial.size = 1, n, p, model.name = NULL,
  prior.control = list(type = c("normal", "normal", "normal", "uniform"),
  hypparams = c(100, 20, 100, 50), ssvs.index = -1, ssvs.g = 1e-6))
```

Arguments

family Either a single element, or a vector of length equal to the number of columns in y . The former assumes all columns of y come from this distribution. The latter option allows for different distributions for each column of y . Elements can be one of "binomial" (with probit link), "poisson" (with log link), "negative.binomial" (with log link), "normal" (with identity link), "lnormal" for log-normal (with log link), "tweedie" (with log link), "exponential" (with log link), "gamma" (with log link), "beta" (with logit link), "ordinal" (cumulative probit regression).

For the negative binomial distribution, the variance is parameterized as $Var(y) = \mu + \phi\mu^2$, where ϕ is the column-specific dispersion parameter. For the normal distribution, the variance is parameterized as $Var(y) = \phi^2$, where ϕ is the column-specific standard deviation. For the tweedie distribution, the variance is parameterized as $Var(y) = \phi\mu^p$ where ϕ is the column-specific dispersion parameter and p is a power parameter common to all columns assumed to be

tweedie, with $1 < p < 2$. For the gamma distribution, the variance is parameterized as $Var(y) = \mu/\phi$ where ϕ is the column-specific rate (henceforth referred to also as dispersion parameter). For the beta distribution, the parameterization is in terms of the mean μ and sample size ϕ (henceforth referred to also as dispersion parameter), so that the two shape parameters are given by $a = \mu\phi$ and $b = (1 - \mu)\phi$.

All columns assumed to have ordinal responses are constrained to have the same cutoffs points, with a column-specific intercept to account for differences between the columns (please see *Details* for formulation).

num.X	Number of columns in the model matrix X . Defaults to 0, in which case it is assumed that no covariates are included in the model. Recall that no intercept should be included in X .
num.traits	Number of columns in the model matrix <code>traits</code> . Defaults to 0, in which case it is assumed no traits are included in model. Recall that no intercept should be included in <code>traits</code> .
which.traits	<p>A list of length equal to (number of columns in $X + 1$), informing which columns of <code>traits</code> the column-specific intercepts and each of the column-specific regression coefficients should be regressed against. The first element in the list applies to the column-specific intercept, while the remaining elements apply to the regression coefficients. Each element of <code>which.traits</code> is a vector indicating which traits are to be used.</p> <p>For example, if <code>which.traits[[2]] = c(2, 3)</code>, then the regression coefficients corresponding to the first column in X are regressed against the second and third columns of <code>traits</code>. If <code>which.traits[[2]] = 0</code>, then the regression coefficients for each column are treated as independent. Please see help file below for more details.</p> <p>Defaults to NULL, in conjunction with <code>traits = NULL</code>.</p>
row.eff	Single element indicating whether row effects are included as fixed effects ("fixed"), random effects ("random") or not included ("none") in the boral model. If random effects, they are drawn from a normal distribution with mean zero and unknown standard deviation. Defaults to "none".
row.ids	A matrix with the number of rows equal to the number of rows in y , and the number of columns equal to the number of row effects to be included in the model. Element (i, j) indicates to the cluster ID of row i in y for random effect eqnj; please see the help file for the main boral function for details. Defaults to NULL, so that if <code>row.eff = "none"</code> then the argument is ignored, otherwise if <code>row.eff = "fixed"</code> or <code>"random"</code> , then <code>row.ids = matrix(1:nrow(y), ncol = 1)</code> i.e., a single, row effect unique to each row.
trial.size	Either equal to a single element, or a vector of length equal to the number of columns in y . If a single element, then all columns assumed to be binomially distributed will have trial size set to this. If a vector, different trial sizes are allowed in each column of y . The argument is ignored for all columns not assumed to be binomially distributed. Defaults to 1, i.e. Bernoulli distribution.
n	The number of rows in the response matrix y .
p	The number of columns in the response matrix y .

- `model.name` Name of the text file that the JAGS model is written to. Defaults to NULL, in which case the default of "jagsboralmodel.txt" is used.
- `prior.control` A list of parameters for controlling the prior distributions. These include:
- *type*: Vector of four strings indicating the type of prior distributions to use. In order, these are: 1) priors for all column-specific intercepts, row effects, and cutoff points for ordinal data; 2) priors for the latent variable coefficients (ignored if `num.lv = 0`); 3) priors for all column-specific coefficients relating to the model matrix X (ignored if `X = NULL`). When traits are included in the model, this is also the prior for the trait regression coefficients (please see section *Including species traits* in the help file for [boral](#) for more information); 4) priors for any dispersion parameters. For elements 1-3, the prior distributions currently available include: I) "normal", which is normal prior with the variance controlled by the `hypparams` argument; II) "cauchy", which is a Cauchy prior with variance controlled by the `hypparams` argument. Gelman, et al. (2008) considers using Cauchy priors with variance 2.5^2 as weakly informative priors for regression coefficients; III) "uniform", which is uniform prior with minimum values given by `-hypparams` and maximum values given by `+hypparams`. For element 4, the prior distributions currently available include: I) "uniform", which is uniform prior with minimum zero and maximum controlled by `hypparams[4]`; II) "halfnormal", which is half-normal prior with variance controlled by `hypparams[4]`; III) "halfcauchy", which is a half-Cauchy prior with variance controlled by the `hypparams[4]` argument. Defaults to the vector `c("normal", "normal", "normal", "uniform")`.
 - *hypparams*: Vector of four hyperparameters used in the set up of prior distributions. In order, these are: 1) affects the prior distribution for all column-specific intercepts, row effects, and cutoff points for ordinal data. If `row.eff = "random"`, this also controls the maximum of the uniform prior for the standard deviation of the random effects normal distribution. Also, if more than two of the columns are ordinal, then this also controls the maximum of the uniform prior for the standard deviation of the column-specific random intercepts for these columns; 2) affects the prior distribution for all latent variable coefficients (ignored if `num.lv = 0`); 3) affects the prior distribution for column-specific coefficients relating to the model matrix X (ignored if `X = NULL`). When traits are included in the model, it also affects the prior distribution for the trait regression coefficients, and controls the maximum of the uniform prior for the standard deviation of the normally distributed random effects; 4) affects the prior distribution used for any dispersion parameters. Defaults to the vector `c(100, 20, 100, 50)`.
 - *ssvs.index*: Indices to be used for stochastic search variable selection (SSVS, George and McCulloch, 1993). Either a single element or a vector with length equal to the number of columns in the implied model matrix X . Each element can take values of -1 (no SSVS is performed on this covariate), 0 (SSVS is performed on individual coefficients for this covariate), or any integer exceeding 1 (SSVS is performed on collectively all coefficients on this covariate/s.) This argument is only read if `X.eff = TRUE`. Please see the [boral](#) help file

for more information regarding the implementation of SSVS. Defaults to -1, in which case no model selection is performed on the fitted model at all.

- `ssvs.g`: Multiplicative, shrinkage factor for SSVS, which controls the strength of the "spike" in the SSVS mixture prior. In summary, if the coefficient is included in the model, the "slab" prior is a normal distribution with mean zero and variance given by `hypparams`, while if the coefficient is not included in the model, the "spike" prior is normal distribution with mean zero and variance given by `hypparams*ssvs.g`. Please see the [boral](#) help file for more information regarding the implementation of SSVS. Defaults to 1e-6.

Details

This function is automatically executed inside [boral](#), and therefore does not need to be run separately before fitting the boral model. It can however be run independently if one is: 1) interested in what the actual JAGS file for a particular boral model looks like, 2) wanting to modify a basic JAGS model file to construct more complex model e.g., include environmental variables.

Please note that [boral](#) currently does not allow the user to manually enter a script to be run.

When running the main function [boral](#), setting `save.model = TRUE` which automatically save the JAGS model file as a text file (with name based on the `model.name`) in the current working directory.

Value

A text file is created, containing the JAGS model to be called by the boral function for entering into jags. This file is automatically deleted once boral has finished running unless `save.model = TRUE`.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

References

- Gelman, et al. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2, 1360-1383.

See Also

[make.jagsboralmodel](#) for writing boral model JAGS scripts with one or more latent variables.

Examples

```
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun
n <- nrow(y); p <- ncol(y);

## Create a boral "null" model JAGS script, where distributions alternative
## between Poisson and negative distributions
## across the rows of y.
```

```

make.jagsboralnullmodel(family = rep(c("poisson", "negative.binomial"), length=p),
  num.X = ncol(spider$x), row.eff = "fixed", n = n, p = p)

## Create a boral "null" model JAGS script, where distributions are all negative
## binomial distributions and covariates will be included!
make.jagsboralnullmodel(family = "negative.binomial",
  num.X = ncol(spider$x), n = n, p = p, model.name = "myawesomeordnullmodel.txt")

## Have a look at the JAGS model file for a boral model involving traits,
## based on the ants data from mvabund.
library(mvabund)
data(antTraits)

y <- antTraits$abun
X <- as.matrix(antTraits$env)
## Include only traits 1, 2, and 5, plus an intercept
traits <- as.matrix(antTraits$traits[,c(1,2,5)])
## Please see help file for boral regarding the use of which.traits
which.traits <- vector("list", ncol(X)+1)
for(i in 1:length(which.traits)) which.traits[[i]] <- 1:ncol(traits)

## Not run:
fit.traits <- boral(y, X = X, traits = traits, which.traits = which.traits,
  family = "negative.binomial", num.lv = 0, model.name = "anttraits.txt")

## End(Not run)

```

plot.boral

Plots of a fitted boral object

Description

Produces four plots relating to the fitted boral object, which can be used for residual analysis. If some of the columns are ordinal, then a single confusion matrix is also produced.

Usage

```

## S3 method for class 'boral'
plot(x, est = "median", jitter = FALSE, ...)

```

Arguments

x	An object of class "boral".
est	A choice of either the posterior median (<code>est == "median"</code>) or posterior mean (<code>est == "mean"</code>) of the parameters, which are then treated as parameter estimates and the fitted values/residuals used in the plots are calculated from. Default is posterior median.

jitter	If jitter = TRUE, then some jittering is applied so that points on the plots do not overlap exactly (which can often occur with discrete data). Please see jitter for its implementation.
...	Additional graphical options to be included in. These include values for cex, cex.lab, cex.axis, cex.main, lwd, and so on.

Details

Four plots are provided:

1. Plot of Dunn-Smyth residuals against the linear predictors. This can be useful to assess whether the assumed mean-variance relationship is adequately satisfied, as well as to look for particular outliers. For ordinal responses things are more ambiguous due to the lack of single definition for "linear predictor". Therefore, instead of linear predictors the Dunn-Smyth residuals are plotted against the fitted values (defined as the level with the highest fitted probability). It is fully acknowledged that this makes things VERY hard to interpret if only some of your columns are ordinal.
2. Plot of Dunn-Smyth residuals against the row index/row names.
3. Plot of Dunn-Smyth residuals against the column index/column names. Both this and the previous plot are useful for assessing how well each row/column of the response matrix is being modeled.
4. A normal quantile plot of the Dunn-Smyth residuals, which can be used to assess the normality assumption and overall goodness of fit.

For ordinal responses, a single confusion matrix between the predicted levels (as based on the class with the highest probability) and true levels is also returned. The table pools the results over all columns assumed to be ordinal.

Note

Due the inherent stochasticity, Dunn-Smyth residuals and consequently the plots will be slightly different time this function is run. Note also the fitted values and residuals are calculated from point estimates of the parameters, as opposed to a fully Bayesian approach (please see details in [fitted.boral](#) and [ds.residuals](#)). Consequently, it is recommended that this function is run several times to ensure that any trends observed in the plots are consistent throughout the runs.

As mentioned above, for ordinal responses things are much more challenging as there is no single definition for "linear predictor". Instead of linear predictors then, for the first plot the Dunn-Smyth residuals are plotted against the fitted values, defined as the level with the highest fitted probability. It is fully acknowledged that this makes things VERY hard to interpret if only some of your columns are ordinal though. Suggestions to improve this are welcome!!!

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

See Also

[fitted.boral](#) to obtain the fitted values, [ds.residuals](#) to obtain Dunn-Smyth residuals and details as to what they are.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun

spider.fit.p <- boral(y, family = "poisson", num.lv = 2,
row.eff = "fixed")

par(mfrow = c(2,2))
plot(spider.fit.p)
dev.off()
## A distinct fan pattern is observed in the plot of residuals
## versus linear predictors plot.

spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
row.eff = "fixed")

par(mfrow = c(2,2))
plot(spider.fit.nb)
dev.off()
## The fan shape is not as clear now,
## and the normal quantile plot also suggests a better fit to the data

## End(Not run)
```

summary.boral

Summary of fitted boral object

Description

A summary of the fitted boral objects including the type of model fitted e.g., error distribution, number of latent variables parameter estimates, values of the information criteria (if applicable), and so on.

Usage

```
## S3 method for class 'boral'
summary(object, est = "median", ...)

## S3 method for class 'summary.boral'
print(x,...)
```

Arguments

```
object      An object of class "boral".
x           An object of class "boral".
```

est	A choice of either whether to print the posterior median (est == "median") or posterior mean (est == "mean") of the parameters.
...	Not used.

Value

Attributes of the model fitted, parameter estimates, and values of the information criteria if `calc.ics = TRUE` in the `boral` object, and posterior probabilities of including individual and/or grouped coefficients in the model based on SSVS if appropriate.

Author(s)

Francis K.C. Hui <fhui28@gmail.com>

See Also

[boral](#) for the fitting function on which `summary` is applied, [get.measures](#) for details regarding the information criteria returned.

Examples

```
## Not run:
library(mvabund) ## Load a dataset from the mvabund package
data(spider)
y <- spider$abun

spider.fit.nb <- boral(y, family = "negative.binomial", num.lv = 2,
row.eff = "fixed")

summary(spider.fit.nb)

## End(Not run)
```

Index

`boral`, 3, 5, 6, 32, 34, 40, 41, 43, 44, 46, 48,
50, 52, 53, 59, 60, 64, 65, 69
`boral-package`, 2

`calc.condlogLik`, 19, 25, 29, 47
`calc.logLik.lv0`, 21, 22, 27, 29
`calc.marglogLik`, 21, 25, 26, 51
`coefspplot`, 16, 30
`create.life`, 31

`ds.residuals`, 37, 39, 67

`fitted.boral`, 38, 38, 67

`get.dic`, 39, 48
`get.enviro.cor`, 41, 54
`get.hpdintervals`, 11, 42
`get.measures`, 5, 12, 16, 21, 28, 29, 40, 45,
49, 50, 52, 69
`get.more.measures`, 16, 28, 29, 40, 48, 49
`get.residual.cor`, 8, 16, 42, 52

`jitter`, 55, 67

`lvspplot`, 8, 16, 55

`make.jagsboralmodel`, 57, 65
`make.jagsboralnullmodel`, 61, 62

`plot.boral`, 38, 39, 66
`print.boral (boral)`, 3
`print.summary.boral (summary.boral)`, 68

`simulate.boral (create.life)`, 31
`summary.boral`, 16, 68