

# Package ‘brea’

October 6, 2016

**Version** 0.1.0

**Date** 2016-09-30

**Title** Bayesian Recurrent Event Analysis

**Description** A function to produce MCMC samples for posterior inference in semiparametric Bayesian discrete time competing risks recurrent events models.

**Author** Adam J King

**Maintainer** Adam J King <king@cpp.edu>

**License** GPL-3

**NeedsCompilation** no

**Repository** CRAN

**Date/Publication** 2016-10-06 13:06:07

## R topics documented:

brea_mcmc	1
<b>Index</b>	<b>5</b>

---

brea_mcmc	<i>Bayesian Discrete Survival Inference</i>
-----------	---

---

## Description

This function performs Metropolis-Hastings exploration of the posterior distribution for Bayesian discrete time survival models. The models may have competing risks, semiparametric covariate effects (like arbitrary baseline hazards), and random effects shared across repeated events and correlated across competing risks.

## Usage

```
brea_mcmc(x, y, S = 1000L, priors = NULL, n = NULL, K = NULL, store_re = FALSE)
```

## Arguments

x	an integer matrix or dataframe all columns of which are factors specifying the values of the (discretized) predictors at each person-time point; specifically, the (i,j) entry is the value of predictor j at discrete person-time observation i
y	an integer matrix whose (i,j) entry counts the number of events of type j occurring at discrete person-time observation i
S	the number of MCMC iterations to perform
priors	a list with one element for each predictor variable (column of x) specifying the prior type to use for that predictor; see Details for more information
n	a vector of positive integers with length equal to the number of person-time observations whose entries equal the number of replicated observations that row stands for (defaults to 1 for each observation)
K	a vector of positive integers with length equal to the number of predictors giving the number of distinct values each discretized predictor may assume; this is not used if a dataframe of factors is provided for x
store_re	if TRUE, the random effects are stored from each MCMC iteration, and if FALSE they are not stored

## Details

The data provided to the `brea_mcmc` function is specified at the person-time level: there is one row in `x` and `y` for each discrete time point each person or thing was at risk for event occurrence. All predictors in `x` must be encoded as factors (or their corresponding integer codes in the case that `x` is an integer matrix). The underlying type of predictor is specified in the `priors` argument, which is a list with one element for each predictor variable which specifies both the type of that predictor and the prior distribution to use. The allowed predictor types are:

- "cat" for categorical variables. The first element of the prior specification list is the string "cat", and second element is a positive conditional prior standard deviation parameter.
- "gmrfr" for underlying continuous predictors; continuous predictors should be `cut()` before being included into `x`; Gaussian Markov random field (GMRF) priors are then used to smooth the effects of adjacent categories of the discretized continuous predictor. The first element of the prior specification list is the string "gmrfr", the second element is a prior degrees of freedom for the scaled inverse chi-squared prior for the random walk increment variance, and the third element is a prior scale for the scaled inverse chi-squared.
- "re" for variables (like subject id numbers) that represent random effects. The first element of the prior specification list is the string "re", the second element is a prior degrees of freedom for an inverse Wishart prior for the random effects covariance, and the third element is a prior scale matrix for the random effects covariance.
- "zero" for predictors that are not used in the current MCMC run. This is provided as a convenient way to exclude predictors from certain runs. The first and only element of the prior specification list is the string "zero".



```
y[next_row:(next_row + study_time[i] - 1),] <- c(rep(0L,study_time[i] - 1),
                                                as.integer(cens[i]))
next_row <- next_row + study_time[i]
}

# group the time variable into 6-week intervals:
x[,1] <- cut(x[,1],seq(0,36,6),labels=FALSE)

# use GMRF prior for time, and categorical prior for treatment group:
priors <- list(list("gmrf",3,.01),list("cat",2))

# run 1,100 MCMC iterations:
fit <- brea_mcmc(x, y, 1100, priors)

# look at the structure of the returned posterior samples and acceptance counts:
str(fit)

# approximate posterior mean hazard ratio:
ss <- 101:1100 # use last 1,000 samples, discarding first 100 as burn-in
exp(2*mean(fit$b_m_s[[2]][1,2,ss]))
# hazard ratio of approximately 5
```

# Index

\*Topic **regression**

bre\_a\_mcmc, 1

\*Topic **survival**

bre\_a\_mcmc, 1

bre\_a\_mcmc, 1