

Package ‘rem’

February 19, 2016

Type Package

Title Relational Event Models (REM)

Version 1.1.2

Date 2016-02-16

Author Laurence Brandenberger

Maintainer Laurence Brandenberger <laurence.brandenberger@eawag.ch>

Description Calculate endogenous network effects in event sequences and fit relational event models (REM): Using network event sequences (where each tie between a sender and a target in a network is time-stamped), REMs can measure how networks form and evolve over time. Endogenous patterns such as popularity effects, inertia, similarities, cycles or triads can be calculated and analyzed over time.

License GPL (>= 2)

Depends R (>= 2.14.0)

Imports Rcpp, flexsurv, survival

LinkingTo Rcpp

Suggests xergm, texreg, statnet, ggplot2

Repository CRAN

Repository/R-Forge/Project xergm

Repository/R-Forge/Revision 182

Repository/R-Forge/DateTimeStamp 2016-02-19 13:22:45

Date/Publication 2016-02-19 23:50:57

NeedsCompilation yes

R topics documented:

rem-package	2
degreeStat	2
eventSequence	7
fourCycleStat	10
inertiaStat	14

reciprocityStat	17
remRate	21
similarityStat	24
triadStat	29

Index	33
--------------	-----------

rem-package	<i>Fit Relational Event Models (REM)</i>
-------------	--

Description

The **rem** package uses a combination of event history and network analysis to test network dependencies in event sequences. If events in an event sequence depend on each other, network structures and patterns can be calculated and estimated using relational event models. The rem-package includes functions to calculate endogenous network statistics in (signed) one-, two- and multi-mode network event sequences. The statistics include inertia ([inertiaStat](#)), reciprocity ([reciprocityStat](#)), in- or outdegree statistics ([degreeStat](#)), closing triads ([triadStat](#)), closing four-cycles ([fourCycleStat](#)) or endogenous similarity statistics ([similarityStat](#)). The rate of event occurrence can then be tested using standard models of event history analysis, such as a weibull model ([remRate](#)).

Details

Package:	rem
Type:	Package
Version:	1.1.2
Date:	2016-02-16

Author(s)

Laurence Brandenberger <laurence.brandenberger@eawag.ch>

References

Lerner, J., Bussmann, M., Snijders, T. A., & Brandes, U. (2013). Modeling frequency and type of interaction in event networks. *Corvinus Journal of Sociology and Social Policy*, (1), 3-32.

degreeStat	<i>Calculate (in/out)-degree statistics</i>
------------	---

Description

Calculate the endogenous network statistic indegree/outdegree for relational event models. indegree/outdegree measures the senders' tendency to be involved in events (sender activity, sender out- or indegree) or the tendency of events to surround a specific target (target popularity, target in- or outdegree)

Usage

```
degreeStat(data, time, degreevar, halflife,
           weight = NULL,
           eventtypevar = NULL,
           eventtypevalue = "valuematch",
           eventattributevar = NULL,
           eventattributevalue = "valuematch",
           degree.on.other.var = NULL,
           variablename = "degree",
           returnData = FALSE,
           showprogressbar = FALSE)
```

Arguments

data	A data frame containing all the variables.
time	Numeric variable that represents the event sequence. The variable has to be sorted in ascending order.
degreevar	A string (or factor or numeric) variable that represents the sender or target of the event. The degree statistic will calculate how often in the past, a given sender or target has been active by counting the number of events in the past where the degreevar is repeated. See details for more information on which variable to chose as degreevar for one- and two-mode networks.
halflife	A numeric value that is used in the decay function. The vector of past events is weighted by an exponential decay function using the specified halflife. The halflife parameter determines after how long a period the event weight should be halved. E.g. if halflife = 5, the weight of an event that occurred 5 units in the past is halved. Smaller halflife values give more importance to more recent events, while larger halflife values should be used if time does not affect the sequence of events that much.
weight	An optional numeric variable that represents the weight of each event. If weight = NULL each event is given an event weight of 1.
eventtypevar	An optional variable that represents the type of the event. Use eventtypevalue to specify how the eventtypevar should be used to filter past events.
eventtypevalue	An optional value (or set of values) used to specify how paste events should be filtered depending on their type. eventtypevalue = "valuematch" indicates that only past events that have the same type should be used to calculate the degree statistic. eventtypevalue = "valuemix" indicates that past and present events of specific types should be used for the degree statistic. All the possible combinations of the eventtypevar-values will be used. E.g. if eventtypevar contains two unique values "a" and "b", 4 degree statistics will be calculated. The first variable calculates the degree effect where the present event is of type

"a" and all the past events are of type "b". The next variable calculates the degree statistic for present events of type "b" and past events of type "a". Additionally, a variable is calculated, where present events as well as past events are of type "a" and a fourth variable calculates the degree statistic for events with type "b" (i.e. `valuematch` on value "b"). `eventtypevalue = c("..", "..")` is similar to the "nodemix"-option, all different combinations of the values specified in `eventtypevalue` are used to create the degree statistics.

`eventattributevar`

An optional variable that represents an attribute of the event. It can be a sender attribute, a target attribute, time or dyad attribute. Use `eventattributevalue` to specify how the `eventattributevar` should be used.

`eventattributevalue`

An optional value (or set of values) used to specify how past events should be filtered depending on their attribute. Similar to `eventtypevalue`. `eventattributevalue = "valuematch"` indicates that only past events that have the same attribute should be used to calculate the degree statistic. `eventtypevalue = "valuemix"` indicates that past and present events of specific attributes should be used for the degree statistic. All the possible combinations of the `eventattributevar`-values will be used. `eventtypevalue = c("..", "..")` is similar to the "valuemix"-option, all different combinations of the values specified in `eventtypevalue` are used to create degree statistics.

`degree.on.other.var`

A string (or factor or numeric) variable that represents the sender or target of the event. It can be used to calculate target-outdegree or sender-indegree statistics in one-mode networks. For the sender indegree statistic, fill the sender variable into the `degreevar` and the target variable into the `degree.on.other.var`. For the target-outdegree statistic, fill the target variable into the `degreevar` and the sender variable into the `degree.on.other.var`.

`variablename`

An optional value (or values) with the name the degree statistic variable should be given. Default "degree" is used. To be used if `returnData = TRUE` or multiple degree statistics are calculated.

`returnData`

TRUE/FALSE. Set to FALSE by default. The new variable(s) are bound directly to the `data.frame` provided and the data frame is returned in full.

`showprogressbar`

TRUE/FALSE. To be implemented.

Details

The `degreeStat()`-function calculates an endogenous statistic that measures whether events have a tendency to include either the same sender or the same target over the entire event sequence.

The effect is calculated as follows.

$$G_t = G_t(E) = (A, B, w_t),$$

G_t represents the network of past events and includes all events E . These events consist each of a sender $a \in A$ and a target $b \in B$ (in one-mode networks $A = B$) and a weight function w_t :

$$w_t(i, j) = \sum_{e:a=i,b=j} |w_e| \cdot e^{-\frac{(t-t_e) \cdot \ln(2)}{T_{1/2}}} \cdot \frac{\ln(2)}{T_{1/2}},$$

where w_e is the event weight (usually a constant set to 1 for each event), t is the current event time, t_e is the past event time and $T_{1/2}$ is a halflife parameter.

For the degree effect, the past events G_t are filtered to include only events where the senders or targets are identical to the current sender or target.

$$\text{sender-outdegree}(G_t, a, b) = \sum_{j \in B} w_t(a, j)$$

$$\text{target-indegree}(G_t, a, b) = \sum_{i \in A} w_t(i, b)$$

$$\text{sender-indegree}(G_t, a, b) = \sum_{i \in A} w_t(i, a)$$

$$\text{target-outdegree}(G_t, a, b) = \sum_{j \in B} w_t(b, j)$$

Depending on whether the degree statistic is measured on the sender variable or the target variable, either activity or popularity effects are calculated.

For one-mode networks: Four distinct statistics can be calculated: sender-indegree, sender-outdegree, target-indegree or target-outdegree. The sender-indegree measures how often the current sender was targeted by other senders in the past (i.e. how popular were current senders). The sender-outdegree measures how often the current sender was involved in an event, where they were also marked as sender (i.e. how active the current sender has been in the past). The target-indegree statistic measures how often the current targets were targeted in the past (i.e. how popular were current targets). And the target-outdegree measures how often the current targets were senders in the past (i.e. how active were current targets in the past).

For two-mode networks: Two distinct statistics can be calculated: sender-outdegree and target-indegree. Sender-outdegree measures how often the current sender has been involved in an event in the past (i.e. how active the sender has been up until now). The target-indegree statistic measures how often the current target has been involved in an event in the past (i.e. how popular a given target has been before the current event).

An exponential decay function is used to model the effect of time on the endogenous statistics. Each past event that contains the same sender or the same target (depending on the variable specified in `degreevar`) and fulfills additional filtering options (specified via `eventtype` or `eventattributes`) is weighted with an exponential decay. The further apart the past event is from the present event, the less weight is given to this event. The halflife parameter in the `degreeStat()`-function determines at which rate the weights of past events should be reduced.

The `eventtypevar`- and `eventattributevar`-options help filter the past events more specifically. How they are filtered depends on the `eventtypevalue`- and `eventattributevalue`-option.

Author(s)

Laurence Brandenberger <laurence.brandenberger@eawag.ch>

See Also

[rem-package](#)

Examples

```
# create some data with 'sender', 'target' and a 'time'-variable
# (Note: Data used here are random events from the Correlates of War Project)
sender <- c('TUN', 'NIR', 'NIR', 'TUR', 'TUR', 'USA', 'URU',
           'IRQ', 'MOR', 'BEL', 'EEC', 'USA', 'IRN', 'IRN',
           'USA', 'AFG', 'ETH', 'USA', 'SAU', 'IRN', 'IRN',
           'ROM', 'USA', 'USA', 'PAN', 'USA', 'USA', 'YEM',
           'SYR', 'AFG', 'NAT', 'NAT', 'USA')
target <- c('BNG', 'ZAM', 'JAM', 'SAU', 'MOM', 'CHN', 'IRQ',
           'AFG', 'AFG', 'EEC', 'BEL', 'ITA', 'RUS', 'UNK',
           'IRN', 'RUS', 'AFG', 'ISR', 'ARB', 'USA', 'USA',
           'USA', 'AFG', 'IRN', 'IRN', 'IRN', 'AFG', 'PAL',
           'ARB', 'USA', 'EEC', 'BEL', 'PAK')
time <- c('800107', '800107', '800107', '800109', '800109',
         '800109', '800111', '800111', '800111', '800113',
         '800113', '800113', '800114', '800114', '800114',
         '800116', '800116', '800116', '800119', '800119',
         '800119', '800122', '800122', '800122', '800124',
         '800125', '800125', '800127', '800127', '800127',
         '800204', '800204', '800204')
type <- sample(c('cooperation', 'conflict'), 33,
              replace = TRUE)

# combine them into a data.frame
dt <- data.frame(sender, target, time, type)

# create event sequence and order the data
dt <- eventSequence(datevar = dt$time, dateformat = '%y%m%d',
                   data = dt, type = 'continuous',
                   byTime = 'daily', returnData = TRUE,
                   sortData = TRUE)

# calculate sender-outdegree statistic
dt$sender.outdegree <- degreeStat(data = dt,
                                 time = dt$time,
                                 degreevar = dt$sender,
                                 halflife = 2,
                                 returnData = FALSE)

# plot sender-outdegree over time
library('ggplot2')
ggplot(dt, aes ( event.seq.cont, sender.outdegree )) +
  geom_point()+ geom_smooth()
```

```

# calculate sender-indegree statistic
dt$sender.indegree <- degreeStat(data = dt,
                                time = dt$time,
                                degreevar = dt$sender,
                                degree.on.other.var = dt$target,
                                halflife = 2,
                                returnData = FALSE)

# calculate target-indegree statistic
dt$target.indegree <- degreeStat(data = dt,
                                time = dt$time,
                                degreevar = dt$target,
                                halflife = 2,
                                returnData = FALSE)

# calculate target-outdegree statistic
dt$target.outdegree <- degreeStat(data = dt,
                                time = dt$time,
                                degreevar = dt$target,
                                degree.on.other.var = dt$sender,
                                halflife = 2,
                                returnData = FALSE)

# calculate target-indegree with typematch
dt$target.indegree.tm <- degreeStat(data = dt,
                                    time = dt$time,
                                    degreevar = dt$target,
                                    halflife = 2,
                                    eventtypevar = dt$type,
                                    eventtypevalue = 'valuematch',
                                    returnData = FALSE)

```

eventSequence	<i>Create event sequence</i>
---------------	------------------------------

Description

Create the event sequence for relational event models. Continuous or ordinal sequences can be created. Various dates may be excluded from the sequence (e.g. special holidays, specific weekdays or longer time spans).

Usage

```

eventSequence(datevar,
              dateformat = NULL, data = NULL,
              type = "continuous", byTime = "daily",
              excludeDate = NULL, excludeTypeOfDay = NULL,
              excludeYear = NULL, excludeFrom = NULL,
              excludeTo = NULL, returnData = FALSE,
              sortData = FALSE)

```

Arguments

datevar	The variable containing the information on the date and/or time of the event.
dateformat	A character string indicating the format of the datevar. see as.Date
data	An optional data frame containing all the variables.
type	"continuous" or "ordinal". Specifies whether the event sequence is to be created as a continuous sequence or an ordinal sequence.
byTime	String value. Specifies at what interval the event sequence is created. Use "daily", "monthly" or "yearly".
excludeDate	An optional string or string vector containing one or more dates that should be excluded from the event.sequence. The dates have to be in the same format as provided in dateformat. Only valid for continuous event sequences.
excludeTypeOfDay	String value or vector naming the day(s) that should be excluded from the event sequence. Depending on the locale the weekdays may be named differently. Use Sys.getlocale("LC_TIME") to find which locale is installed.
excludeYear	A string value or vector naming the year(s) that should be excluded from the event sequence.
excludeFrom	A string value (or a vector of strings) with the start value of the date from (from-value included) which the event sequence should not be affected. The value has to be in the same format as specified in dateformat.
excludeTo	A string value (or a vector of strings) with the end value of the date to which time the event sequence should not be affected (to-value included). The value has to be in the same format as specified in dateformat.
returnData	TRUE/FALSE. Default set to FALSE. The data frame provided is returned in full, together with the new variable for the event sequence.
sortData	TRUE/FALSE. Default set to FALSE. Should only be used if returnData = TRUE. The entire data.frame will be ordered according to the event sequence.

Details

In order to estimate relational event models, the events have to be ordered, either according to an ordinal or a continuous event sequence. The ordinal event sequence simply orders the events and gives each event a place in the sequence. The continuous event sequence creates an artificial sequence ranging from `min(datevar)` to `max(datevar)` and matches each event with its place in the artificial event sequence. Dates, years or Weekdays can be excluded from the artificial event sequence. This is useful for excluding specific holidays, weekends etc..

Where two or more events occur at the same time, they are given the same value in the event sequence.

Author(s)

Laurence Brandenberger <laurence.brandenberger@eawag.ch>

See Also

[rem-package](#)

Examples

```

# create some data with 'sender', 'target' and a 'time'-variable
# (Note: Data used here are random events from the Correlates of War Project)
sender <- c('TUN', 'NIR', 'NIR', 'TUR', 'TUR', 'USA', 'URU',
           'IRQ', 'MOR', 'BEL', 'EEC', 'USA', 'IRN', 'IRN',
           'USA', 'AFG', 'ETH', 'USA', 'SAU', 'IRN', 'IRN',
           'ROM', 'USA', 'USA', 'PAN', 'USA', 'USA', 'YEM',
           'SYR', 'AFG', 'NAT', 'NAT', 'USA')
target <- c('BNG', 'ZAM', 'JAM', 'SAU', 'MOM', 'CHN', 'IRQ',
           'AFG', 'AFG', 'EEC', 'BEL', 'ITA', 'RUS', 'UNK',
           'IRN', 'RUS', 'AFG', 'ISR', 'ARB', 'USA', 'USA',
           'USA', 'AFG', 'IRN', 'IRN', 'IRN', 'AFG', 'PAL',
           'ARB', 'USA', 'EEC', 'BEL', 'PAK')
time <- c('800107', '800107', '800107', '800109', '800109',
         '800109', '800111', '800111', '800111', '800113',
         '800113', '800113', '800114', '800114', '800114',
         '800116', '800116', '800116', '800119', '800119',
         '800119', '800122', '800122', '800122', '800124',
         '800125', '800125', '800127', '800127', '800127',
         '800204', '800204', '800204')

# combine them into a data.frame
dt <- data.frame(sender, target, time)

# create continuous event sequence: return the data with the
# event sequence and sort the data according to the event sequence.
dt <- eventSequence(datevar = dt$time, dateformat = '%y%m%d',
                  data = dt, type = 'continuous',
                  byTime = 'daily', returnData = TRUE,
                  sortData = TRUE)

# alternative : create variable with the continuous event
# sequence, unsorted
dt$eventSeq <- eventSequence(datevar = dt$time,
                          dateformat = '%y%m%d',
                          data = dt, type = 'continuous',
                          byTime = 'daily',
                          returnData = FALSE,
                          sortData = FALSE)

# manually sort the data set
dt <- dt[order(dt$eventSeq), ]

# create the sequence by month
dt$eventSeqMonthly <- eventSequence(datevar = dt$time,
                                   dateformat = '%y%m%d',
                                   data = dt,
                                   type = 'continuous',
                                   byTime = 'monthly',
                                   returnData = FALSE,
                                   sortData = FALSE)

# create the sequence by year

```

```

dt$eventSeqYearly <- eventSequence(datevar = dt$time,
                                  dateformat = '%y%m%d',
                                  data = dt,
                                  type = 'continuous',
                                  byTime = 'yearly',
                                  returnData = FALSE,
                                  sortData = FALSE)

# create an ordinal event sequence
dt$eventSeqOrdinal <- eventSequence(datevar = dt$time,
                                    dateformat = '%y%m%d',
                                    data = dt,
                                    type = 'ordinal',
                                    byTime = 'daily',
                                    returnData = FALSE,
                                    sortData = FALSE)

# exclude certain dates
dt$eventSeqEx <- eventSequence(datevar = dt$time,
                                dateformat = '%y%m%d',
                                data = dt, type = 'continuous',
                                byTime = 'daily',
                                excludeDate = c('800108', '800112'),
                                returnData = FALSE,
                                sortData = FALSE)

```

fourCycleStat

Calculate four cycle statistics

Description

Calculate the endogenous network statistic fourCycle that measures the tendency for events to close four cycles in two-mode event sequences.

Usage

```

fourCycleStat(data, time, sender, target, halflife,
              weight = NULL,
              eventtypevar = NULL,
              eventtypevalue = 'standard',
              eventattributevar = NULL,
              eventattributeAB = NULL, eventattributeAJ = NULL,
              eventattributeIB = NULL, eventattributeIJ = NULL,
              variablename = 'fourCycle',
              returnData = FALSE,
              showprogressbar = FALSE)

```

Arguments

data	A data frame containing all the variables.
time	Numeric variable that represents the event sequence. The variable has to be sorted in ascending order.
sender	A string (or factor or numeric) variable that represents the sender of the event.
target	A string (or factor or numeric) variable that represents the target of the event.
halflife	A numeric value that is used in the decay function. The vector of past events is weighted by an exponential decay function using the specified halflife. The halflife parameter determines after how long a period the event weight should be halved. E.g. if halflife = 5, the weight of an event that occurred 5 units in the past is halved. Smaller halflife values give more importance to more recent events, while larger halflife values should be used if time does not affect the sequence of events that much.
weight	An optional numeric variable that represents the weight of each event. If weight = NULL each event is given an event weight of 1.
eventtypevar	An optional variable that represents the type of the event. Use eventtypevalue to specify how the eventtypevar should be used to filter past events.
eventtypevalue	An optional value (or set of values) used to specify how past events should be filtered depending on their type. 'standard', 'positive' or 'negative' may be used. Default set to 'standard'. 'standard' refers to closing four cycles where the type of the events is irrelevant. 'positive' closing four cycles can be classified as reciprocity via the second mode. It indicates whether senders have a tendency to reciprocate or show support by engaging in targets that close a four cycle between two senders. 'negative' closing four cycles represent opposition between two senders, where the current event is more likely if the two senders have opposed each other in the past. Support or opposition is represented by the eventtypevar value for each event.
eventattributevar	An optional variable that represents an attribute of the event. It can be a sender attribute, a target attribute, time or dyad attribute. Use eventattributeAB, eventattributeAJ, eventattributeIB or eventattributeIJ to specify how the eventattributevar should be used.
eventattributeAB	An optional value used to specify how past events should be filtered depending on their attribute. Each distinct edge that form a four cycle can be filtered. eventattributeAB refers to the current event. eventattributeAJ refers to the event involving the current sender and target j that has been used by the current as well as the second actor in the past. eventattributeIB refers to the event involving the second sender and the current target. eventattributeIJ filters events that involve the second sender and the second target. See the four cycle formula in the details section for more information.
eventattributeAJ	see eventattributeAB.
eventattributeIB	see eventattributeAB.

eventattributeIJ	see eventattributeAB.
variablename	An optional value (or values) with the name the four cycle statistic variable should be given. To be used if returnData = TRUE.
returnData	TRUE/FALSE. Set to FALSE by default. The new variable(s) are bound directly to the data.frame provided and the data frame is returned in full.
showprogressbar	TRUE/FALSE. To be implemented.

Details

The `fourCycleStat()`-function calculates an endogenous statistic that measures whether events have a tendency to form four cycles.

The effect is calculated as follows:

$$G_t = G_t(E) = (A, B, w_t),$$

G_t represents the network of past events and includes all events E . These events consist each of a sender $a \in A$ and a target $b \in B$ and a weight function w_t :

$$w_t(i, j) = \sum_{e:a=i,b=j} |w_e| \cdot e^{-\frac{(t-t_e) \cdot \ln(2)}{T_{1/2}}} \cdot \frac{\ln(2)}{T_{1/2}},$$

where w_e is the event weight (usually a constant set to 1 for each event), t is the current event time, t_e is the past event time and $T_{1/2}$ is a halflife parameter.

For the four-cycle effect, the past events G_t are filtered to include only events where the current event closes an open four-cycle in the past.

$$fourCycle(G_t, a, b) = \sqrt[3]{\sum_{i \in A \& j \in B} w_t(a, j) \cdot w_t(i, b) \cdot w_t(i, j)}$$

An exponential decay function is used to model the effect of time on the endogenous statistics. The further apart the past event is from the present event, the less weight is given to this event. The halflife parameter in the `fourCycleStat()`-function determines at which rate the weights of past events should be reduced. Therefore, if the one (or more) of the three events in the four cycle have occurred further in the past, less weight is given to this four cycle because it becomes less likely that the two senders reacted to each other in the way the four cycle assumes.

The `eventtypevar`- and `eventattributevar`-options help filter the past events more specifically. How they are filtered depends on the `eventtypevalue`- and `eventattributevalue`-option.

Author(s)

Laurence Brandenberger <laurence.brandenberger@eawag.ch>

See Also

[rem-package](#)


```
# calculate negative closing four-cycles: general opposition
dt$fourCycle.neg <- fourCycleStat(data = dt,
                                time = dt$event.seq.cont,
                                sender = dt$sender,
                                target = dt$target,
                                halflife = 20,
                                eventtypevar = dt$type,
                                eventtypevalue = 'negative')
```

inertiaStat

Calculate inertia statistics

Description

Calculate the endogenous network statistic *inertia* for relational event models. *inertia* measures the tendency for events to consist of the same sender and target (i.e. repeated events).

Usage

```
inertiaStat(data, time, sender, target, halflife,
            weight = NULL,
            eventtypevar = NULL,
            eventtypevalue = "valuematch",
            eventattributevar = NULL,
            eventattributevalue = "valuematch",
            variablename = "inertia",
            returnData = FALSE,
            showprogressbar = FALSE)
```

Arguments

<code>data</code>	A data frame containing all the variables.
<code>time</code>	Numeric variable that represents the event sequence. The variable has to be sorted in ascending order.
<code>sender</code>	A string (or factor or numeric) variable that represents the sender of the event.
<code>target</code>	A string (or factor or numeric) variable that represents the target of the event.
<code>halflife</code>	A numeric value that is used in the decay function. The vector of past events is weighted by an exponential decay function using the specified halflife. The halflife parameter determines after how long a period the event weight should be halved. E.g. if <code>halflife = 5</code> , the weight of an event that occurred 5 units in the past is halved. Smaller halflife values give more importance to more recent events, while larger halflife values should be used if time does not affect the sequence of events that much.
<code>weight</code>	An optional numeric variable that represents the weight of each event. If <code>weight = NULL</code> each event is given an event weight of 1.

- eventtypevar** An optional variable that represents the type of the event. Use `eventtypevalue` to specify how the `eventtypevar` should be used to filter past events.
- eventtypevalue** An optional value (or set of values) used to specify how paste events should be filtered depending on their type. `eventtypevalue = "valuematch"` indicates that only past events that have the same type as the current event should be used to calculate the inertia statistic. `eventtypevalue = "valuemix"` indicates that past and present events of specific types should be used for the inertia statistic. All the possible combinations of the `eventtypevar`-values will be used. E.g. if `eventtypevar` contains two unique values "a" and "b", 4 inertia statistics will be calculated. The first variable calculates the inertia effect where the present event is of type "a" and all the past events are of type "b". The next variable calculates inertia for present events of type "b" and past events of type "a". Additionally, a variable is calculated, where present events as well as past events are of type "a" and a fourth variable calculates inertia for events with type "b" (i.e. `valuematch` on value "b"). `eventtypevalue = c(., .)` is similar to the "nodmix"-option, all different combinations of the values specified in `eventtypevalue` are used to create inertia statistics.
- eventattributevar** An optional variable that represents an attribute of the event. It can be a sender attribute, a target attribute, time or dyad attribute. Use `eventattributevalue` to specify how the `eventattributevar` should be used.
- eventattributevalue** An optional value (or set of values) used to specify how paste events should be filtered depending on their attribute. Similar to `eventtypevalue`. `eventattributevalue = "valuematch"` indicates that only past events that have the same attribute should be used to calculate the inertia statistic. `eventattributevalue = "valuemix"` indicates that past and present events of specific attributes should be used for the inertia statistic. All the possible combinations of the `eventattributevar`-values will be used. `eventattributevalue = c(., .)` similar to the "valuemix"-option, all different combinations of the values specified in `eventtypevalue` are used to create inertia statistics.
- variablename** An optional value (or values) with the name the inertia statistic variable should be given. To be used if `returnData = TRUE` or multiple inertia statistics are calculated.
- returnData** TRUE/FALSE. Set to FALSE by default. The new variable(s) are bound directly to the `data.frame` provided and the data frame is returned in full.
- showprogressbar** TRUE/FALSE. To be implemented.

Details

The `inertiaStat()`-function calculates an endogenous statistic that measures whether events have a tendency to be repeated with the same sender and target over the entire event sequence.

The effect is calculated as follows.

$$G_t = G_t(E) = (A, B, w_t),$$

G_t represents the network of past events and includes all events E . These events consist each of a sender $a \in A$ and a target $b \in B$ and a weight function w_t :

$$w_t(i, j) = \sum_{e:a=i,b=j} |w_e| \cdot e^{-(t-t_e) \cdot \frac{\ln(2)}{T_{1/2}}} \cdot \frac{\ln(2)}{T_{1/2}},$$

where w_e is the event weight (usually a constant set to 1 for each event), t is the current event time, t_e is the past event time and $T_{1/2}$ is a half-life parameter.

For the inertia effect, the past events G_t are filtered to include only events where the senders and targets are identical to the current sender and target.

$$\text{inertia}(G_t, a, b) = w_t(a, b)$$

An exponential decay function is used to model the effect of time on the endogenous statistics. Each past event that contains the same sender and target and fulfills additional filtering options specified via event type or event attributes is weighted with an exponential decay. The further apart the past event is from the present event, the less weight is given to this event. The half-life parameter in the `inertiaStat()`-function determines at which rate the weights of past events should be reduced.

The `eventtypevar`- and `eventattributevar`-options help filter the past events more specifically. How they are filtered depends on the `eventtypevalue`- and `eventattributevalue`-option.

Author(s)

Laurence Brandenberger <laurence.brandenberger@eawag.ch>

See Also

[rem-package](#)

Examples

```
# create some data with 'sender', 'target' and a 'time'-variable
# (Note: Data used here are random events from the Correlates of War Project)
sender <- c('TUN', 'NIR', 'NIR', 'TUR', 'TUR', 'USA', 'URU',
           'IRQ', 'MOR', 'BEL', 'EEC', 'USA', 'IRN', 'IRN',
           'USA', 'AFG', 'ETH', 'USA', 'SAU', 'IRN', 'IRN',
           'ROM', 'USA', 'USA', 'PAN', 'USA', 'USA', 'YEM',
           'SYR', 'AFG', 'NAT', 'NAT', 'USA')
target <- c('BNG', 'ZAM', 'JAM', 'SAU', 'MOM', 'CHN', 'IRQ',
           'AFG', 'AFG', 'EEC', 'BEL', 'ITA', 'RUS', 'UNK',
           'IRN', 'RUS', 'AFG', 'ISR', 'ARB', 'USA', 'USA',
           'USA', 'AFG', 'IRN', 'IRN', 'IRN', 'AFG', 'PAL',
           'ARB', 'USA', 'EEC', 'BEL', 'PAK')
time <- c('800107', '800107', '800107', '800109', '800109',
         '800109', '800111', '800111', '800111', '800113',
         '800113', '800113', '800114', '800114', '800114',
         '800116', '800116', '800116', '800119', '800119',
         '800119', '800122', '800122', '800122', '800124',
         '800125', '800125', '800127', '800127', '800127',
```



```

      '800204', '800204', '800204')
type <- sample(c('cooperation', 'conflict'), 33,
              replace = TRUE)

# combine them into a data.frame
dt <- data.frame(sender, target, time, type)

# create event sequence and order the data
dt <- eventSequence(datevar = dt$time, dateformat = '%y%m%d',
                   data = dt, type = 'continuous',
                   byTime = "daily", returnData = TRUE,
                   sortData = TRUE)

# calculate inertia statistics
dt$inertia <- inertiaStat(data = dt, time = dt$event.seq.cont,
                        sender = dt$sender, target = dt$target,
                        halflife = 2, returnData = FALSE)

# plot inertia over time
library('ggplot2')
ggplot(dt, aes ( event.seq.cont, inertia) ) +
  geom_point()+ geom_smooth()

# inertia with typematch (e.g. for 'cooperation' events only
# count past 'cooperation' events)
dt$inertia.tm <- inertiaStat(data = dt,
                            time = dt$event.seq.cont,
                            sender = dt$sender,
                            target = dt$target,
                            halflife = 2,
                            eventtypevar = dt$type,
                            eventtypevalue = 'valuematch',
                            returnData = FALSE)

# inertia with valuemix: for each combination of types
# in the eventtypevar, create a variable
dt <- inertiaStat(data = dt, time = dt$event.seq.cont,
                 sender = dt$sender, target = dt$target,
                 halflife = 2, eventtypevar = dt$type,
                 eventtypevalue = 'valuemix',
                 returnData = TRUE)

```

reciprocityStat

Calculate reciprocity statistics

Description

Calculate the endogenous network statistic reciprocity for relational event models. reciprocity measures the tendency for senders to reciprocate prior events where they were targeted by other senders. One-mode network statistic only.

Usage

```
reciprocityStat(data, time, sender, target, halflife,
  weight = NULL,
  eventtypevar = NULL,
  eventtypevalue = "valuematch",
  eventattributevar = NULL,
  eventattributevalue = "valuematch",
  variablename = "reciprocity",
  returnData = FALSE,
  showprogressbar = FALSE)
```

Arguments

data	A data frame containing all the variables.
time	Numeric variable that represents the event sequence. The variable has to be sorted in ascending order.
sender	A string (or factor or numeric) variable that represents the sender of the event.
target	A string (or factor or numeric) variable that represents the target of the event.
halflife	A numeric value that is used in the decay function. The vector of past events is weighted by an exponential decay function using the specified halflife. The halflife parameter determines after how long a period the event weight should be halved. E.g. if halflife = 5, the weight of an event that occurred 5 units in the past is halved. Smaller halflife values give more importance to more recent events, while larger halflife values should be used if time does not affect the time between events that much.
weight	An optional numeric variable that represents the weight of each event. If weight = NULL each event is given an event weight of 1.
eventtypevar	An optional variable that represents the type of the event. Use eventtypevalue to specify how the eventtypevar should be used to filter past events.
eventtypevalue	An optional value (or set of values) used to specify how past events should be filtered depending on their type. eventtypevalue = "valuematch" indicates that only past events that have the same type as the current event should be used to calculate the reciprocity statistic. eventtypevalue = "valuemix" indicates that past and present events of specific types should be used for the reciprocity statistic. All the possible combinations of the eventtypevar-values will be used. E.g. if eventtypevar contains three unique values "a" and "b", 4 reciprocity statistics will be calculated. The first variable calculates the reciprocity effect where the present event is of type "a" and all the past events are of type "b". The next variable calculates reciprocity for present events of type "b" and past events of type "a". Additionally, a variable is calculated, where present events as well as past events are of type "a" and a fourth variable calculates reciprocity for events with type "b" (i.e. valuematch on value "b"). eventtypevalue = c(.., ..), similar to the "nodmix"-option, all different combinations of the values specified in eventtypevalue are used to create reciprocity statistics.

eventattributevar

An optional variable that represents an attribute of the event. It can be a sender attribute, a target attribute, time or dyad attribute. Use eventattributevalue to specify how the eventattributevar should be used.

eventattributevalue

An optional value (or set of values) used to specify how past events should be filtered depending on their attribute. Similar to eventtypevalue. eventattributevalue = "valuematch" indicates that only past events that have the same attribute should be used to calculate the reciprocity statistic. eventtypevalue = "valuemix" indicates that past and present events of specific attributes should be used for the reciprocity statistic. All the possible combinations of the eventattributevar-values will be used. eventtypevalue = c(., .) similar to the "valuemix"-option, all different combinations of the values specified in eventtypevalue are used to create reciprocity statistics.

variablename

An optional value (or values) with the name the reciprocity statistic variable should be given. To be used if returnData = TRUE or multiple reciprocity statistics are calculated.

returnData

TRUE/FALSE. Set to FALSE by default. The new variable(s) are bound directly to the data.frame provided and the data frame is returned in full.

showprogressbar

TRUE/FALSE. To be implemented.

Details

The reciprocityStat()-function calculates an endogenous statistic that measures whether senders have a tendency to reciprocate events.

The effect is calculated as follows:

$$G_t = G_t(E) = (A, B, w_t),$$

G_t represents the network of past events and includes all events E . These events consist each of a sender $a \in A$ and a target $b \in B$ and a weight function w_t :

$$w_t(i, j) = \sum_{e:a=i,b=j} |w_e| \cdot e^{-\frac{(t-t_e) \cdot \ln(2)}{T_{1/2}}} \cdot \frac{\ln(2)}{T_{1/2}},$$

where w_e is the event weight (usually a constant set to 1 for each event), t is the current event time, t_e is the past event time and $T_{1/2}$ is a halflife parameter.

For the reciprocity effect, the past events G_t are filtered to include only events where the senders are the present targets and the targets are the present senders:

$$\text{reciprocity}(G_t, a, b) = w_t(b, a)$$

An exponential decay function is used to model the effect of time on the endogenous statistics. Each past event that involves the sender as target and the target as sender, and fulfills additional filtering options specified via event type or event attributes, is weighted with an exponential decay.

The further apart the past event is from the present event, the less weight is given to this event. The halflife parameter in the `reciprocityStat()`-function determines at which rate the weights of past events should be reduced.

The `eventtypevar-` and `eventattributevar-`options help filter the past events more specifically. How they are filtered depends on the `eventtypevalue-` and `eventattributevalue-`option.

Author(s)

Laurence Brandenberger <laurence.brandenberger@eawag.ch>

See Also

[rem-package](#)

Examples

```
# create some data with 'sender', 'target' and a 'time'-variable
# (Note: Data used here are random events from the Correlates of War Project)
sender <- c('TUN', 'NIR', 'NIR', 'TUR', 'TUR', 'USA', 'URU',
           'IRQ', 'MOR', 'BEL', 'EEC', 'USA', 'IRN', 'IRN',
           'USA', 'AFG', 'ETH', 'USA', 'SAU', 'IRN', 'IRN',
           'ROM', 'USA', 'USA', 'PAN', 'USA', 'USA', 'YEM',
           'SYR', 'AFG', 'NAT', 'NAT', 'USA')
target <- c('BNG', 'ZAM', 'JAM', 'SAU', 'MOM', 'CHN', 'IRQ',
           'AFG', 'AFG', 'EEC', 'BEL', 'ITA', 'RUS', 'UNK',
           'IRN', 'RUS', 'AFG', 'ISR', 'ARB', 'USA', 'USA',
           'USA', 'AFG', 'IRN', 'IRN', 'IRN', 'AFG', 'PAL',
           'ARB', 'USA', 'EEC', 'BEL', 'PAK')
time <- c('800107', '800107', '800107', '800109', '800109',
         '800109', '800111', '800111', '800111', '800113',
         '800113', '800113', '800114', '800114', '800114',
         '800116', '800116', '800116', '800119', '800119',
         '800119', '800122', '800122', '800122', '800124',
         '800125', '800125', '800127', '800127', '800127',
         '800204', '800204', '800204')
type <- sample(c('cooperation', 'conflict'), 33,
              replace = TRUE)
important <- sample(c('important', 'not important'), 33,
                  replace = TRUE)

# combine them into a data.frame
dt <- data.frame(sender, target, time, type, important)

# create event sequence and order the data
dt <- eventSequence(datevar = dt$time, dateformat = '%y%m%d',
                  data = dt, type = "continuous",
                  byTime = 'daily', returnData = TRUE,
                  sortData = TRUE)

# calculate reciprocity statistic
dt$recip <- reciprocityStat(data = dt,
                          time = dt$time,
```

```

        sender = dt$sender,
        target = dt$target,
        halflife = 2)

# plot sender-outdegree over time
library('ggplot2')
ggplot(dt, aes ( event.seq.cont, recip) ) +
  geom_point()+ geom_smooth()

# calculate reciprocity statistic with typematch
# if a cooperated with b in the past, does
# b cooperate with a now?
dt$recip.typematch <- reciprocityStat(data = dt,
                                     time = dt$time,
                                     sender = dt$sender,
                                     target = dt$target,
                                     eventtypevar = dt$type,
                                     eventtypevalue = 'valuematch',
                                     halflife = 2)

# calculate reciprocity with valuemix on type
dt <- reciprocityStat(data = dt,
                     time = dt$time,
                     sender = dt$sender,
                     target = dt$target,
                     eventtypevar = dt$type,
                     eventtypevalue = 'valuemix',
                     halflife = 2,
                     variablename = 'recip',
                     returnData = TRUE)

# calculate reciprocity and count important events only
dt$recip.im <- reciprocityStat(data = dt,
                              time = dt$time,
                              sender = dt$sender,
                              target = dt$target,
                              eventattributevar = dt$important,
                              eventattributevalue = 'important',
                              halflife = 2)

```

remRate

Rate function for relational event models

Description

Estimate the rate function for relational event models.

Usage

```
remRate(formula,
```

```

dist = 'weibull',
var.names = NULL,
event.sequence.type = 'continuous', ...)

```

Arguments

formula	A formula expression in the conventional R syntax used for linear modelling. The response (i.e. the dependent variable) must be the event sequence. On the right-hand side, covariates are listed with a +. It is possible to use one of the endogenous statistics functions in the rem-package as independent (or control) variable. flexsurvreg is called to estimate the parameters of the model for now.
dist	The distribution specified for the parametric model. See flexsurvreg for more details.
var.names	An optional string vector containing variable names for the proposed independent variables in the formula. To be used if endogenous network effects are calculated within the remRate-function.
event.sequence.type	'continuous' or 'ordinal'. Only 'continuous' is implemented as of now.
...	Optional arguments for flexsurvreg .

Details

The function `remRate()` allows users to estimate the rate function of the relational event model. After the variables are prepared, a [flexsurvreg](#) is called to estimate the parameters of the model.

Author(s)

Laurence Brandenberger <laurence.brandenberger@eawag.ch>

See Also

[rem-package](#)

Examples

```

# create some data two-mode network event sequence data with
# a 'sender', 'target' and a 'time'-variable
sender <- c('A', 'B', 'A', 'C', 'A', 'D', 'F', 'G', 'A', 'B',
           'B', 'C', 'D', 'E', 'F', 'B', 'C', 'D', 'E', 'C',
           'A', 'F', 'E', 'B', 'C', 'E', 'D', 'G', 'A', 'G',
           'F', 'B', 'C')
target <- c('T1', 'T2', 'T3', 'T2', 'T1', 'T4', 'T6', 'T2',
           'T4', 'T5', 'T5', 'T5', 'T1', 'T6', 'T7', 'T2',
           'T3', 'T1', 'T1', 'T4', 'T5', 'T6', 'T8', 'T2',
           'T7', 'T1', 'T6', 'T7', 'T3', 'T4', 'T7', 'T8', 'T2')
time <- c('03.01.15', '04.01.15', '10.02.15', '28.02.15', '01.03.15',
         '07.03.15', '07.03.15', '12.03.15', '04.04.15', '28.04.15',
         '06.05.15', '11.05.15', '13.05.15', '17.05.15', '22.05.15',
         '09.08.15', '09.08.15', '14.08.15', '16.08.15', '29.08.15',
         '05.09.15', '25.09.15', '02.10.15', '03.10.15', '11.10.15',

```

```

      '18.10.15', '20.10.15', '28.10.15', '04.11.15', '09.11.15',
      '10.12.15', '11.12.15', '12.12.15')
type <- sample(c('con', 'pro'), 33, replace = TRUE)
important <- sample(c('important', 'not important'), 33,
  replace = TRUE)
sender.country <- sample(c('country1', 'country2'), 33,
  replace = TRUE)

# combine them into a data.frame
dt <- data.frame(sender, target, time, type, important, sender.country)

# create event sequence and order the data
dt <- eventSequence(datevar = dt$time, dateformat = '%d.%m.%y',
  data = dt, type = 'continuous',
  byTime = 'daily', returnData = TRUE,
  sortData = TRUE)

# create some endogenous variables:
dt$inertia <- inertiaStat(dt, dt$event.seq.cont, dt$sender, dt$target, 30)
dt$activitySender <- degreeStat(dt, dt$event.seq.cont, dt$sender, 30)
dt$targetPopularity <- degreeStat(dt, dt$event.seq.cont, dt$target, 30)

# check out their correlation coefficient
cor(subset(dt, select = c(inertia, activitySender,
  targetPopularity)), method="pearson")

# fit event history model using inertia, actor activity and
# target popularity
fit1 <- remRate(
  dt$event.seq.cont ~
  inertiaStat(dt, dt$event.seq.cont, dt$sender, dt$target, 30) +
  degreeStat(dt, dt$event.seq.cont, dt$sender, 30) +
  degreeStat(dt, dt$event.seq.cont, dt$target, 30),
  var.names = c('inertia', 'actorAct', 'targetPop'))
fit1
## Interpretation of the weibull-model coefficients:
# For positive coefficients, if the independent variable increases
# in value, the survival time increases, meaning the event will occur
# at a lower rate. The coefficients are log odds. Calculate the odds
# ratios by running exp(coefficient). This will give you the factor
# by which the risk of the event occurring increases
# (negative coefficient, OR < 1) or decreases
# (positive coefficient, OR > 1).

# calculate closing four-cycles
dt$fourCyc <- fourCycleStat(dt, dt$event.seq.cont,
  dt$sender, dt$target, 30)
dt$fourCyc.support <- fourCycleStat(dt, dt$event.seq.cont,
  dt$sender, dt$target, 30,
  eventtypevar = dt$type,
  eventtypevalue = 'positive')

# do actors from country1 engage more often in positive

```

```

# closing four-cycles?
library('ggplot2')
ggplot(dt, aes (event.seq.cont, fourCyc.support,
               color = sender.country)) +
  geom_point()+ geom_smooth()

# test above idea of closing four-cycles with an
# interaction term with 'sender.country'
fit2 <- remRate(
  dt$event.seq.cont ~
    inertiaStat(dt, dt$event.seq.cont, dt$sender, dt$target, 30) +
    degreeStat(dt, dt$event.seq.cont, dt$sender, 30) +
    degreeStat(dt, dt$event.seq.cont, dt$target, 30) +
    dt$fourCyc.support*dt$sender.country,
  var.names = c('inertia', 'actorActivity',
                'targetPopularity',
                'fourCyc.support', 'country',
                'fourCyc.supportXcountry'))

fit2

# sender similarity: how likely is it that two
# senders are alike?
fit3 <- remRate(
  dt$event.seq.cont ~
    inertiaStat(dt, dt$event.seq.cont, dt$sender, dt$target, 30) +
    degreeStat(dt, dt$event.seq.cont, dt$sender, 30) +
    degreeStat(dt, dt$event.seq.cont, dt$target, 30) +
    similarityStat(dt, dt$event.seq.cont, dt$sender, dt$target,
                  halflife.last.event = 10),
  var.names = c('inertia', 'actorActivity',
                'targetPopularity', 'simSender'))

fit3

# target similarity: how likely is it that two
# targets are used together (e.g. the current actor
# copies others/learns from others)
fit4 <- remRate(
  dt$event.seq.cont ~
    inertiaStat(dt, dt$event.seq.cont, dt$sender, dt$target, 30) +
    degreeStat(dt, dt$event.seq.cont, dt$sender, 30) +
    degreeStat(dt, dt$event.seq.cont, dt$target, 30) +
    similarityStat(dt, dt$event.seq.cont, dt$sender, dt$target,
                  senderOrTarget = 'target',
                  halflife.last.event = 10),
  var.names = c('inertia', 'actorActivity',
                'targetPopularity', 'simTarget'))

fit4

```


Description

Calculate the endogenous network statistic similarity for relational event models. `similarity` measures the tendency for senders to adapt their behavior to that of their peers.

Usage

```
similarityStat(data, time, sender, target,
  senderOrTarget = 'sender',
  whichSimilarity = NULL,
  halflife.last.event = NULL,
  halflife.time.between.events = NULL,
  eventtypevar = NULL,
  eventattributevar = NULL,
  eventattributevalue = NULL,
  variablename = 'similarity',
  returnData = FALSE,
  showprogressbar = FALSE)
```

Arguments

<code>data</code>	A data frame containing all the variables.
<code>time</code>	Numeric variable that represents the event sequence. The variable has to be sorted in ascending order.
<code>sender</code>	A string (or factor or numeric) variable that represents the sender of the event.
<code>target</code>	A string (or factor or numeric) variable that represents the target of the event.
<code>senderOrTarget</code>	sender or target. Indicates on which variable (sender or target) the similarity should be calculated on. Sender similarity measures how many targets the current sender has in common with other senders who used the same targets in the past. Target similarity measures how many senders have used the current target as well as another target that the current sender used in the past.
<code>whichSimilarity</code>	"total" or "average". Indicates how the variable should be aggregated. "total" counts the number of similar events there are in the past event history. "average" divides the count of similar events by the number of senders or the number of targets, depending on which mode of similarity is chosen.
<code>halflife.last.event</code>	A numeric value that is used in the decay function. The vector of past events is weighted by an exponential decay function using the specified halflife. The halflife parameter determines after how long a period the event weight should be halved. For sender similarity: The halflife determines the weight of the count of targets that two actors have in common. The further back the second sender was active, the less weight is given the similarity between this sender and the current sender. For target similarity: The halflife determines the weight of the count of targets that have used both been used by other senders in the past. The longer ago the current sender engaged in an event with the other target, the less weight is given the count.

halflife.time.between.events	A numeric value that is used in the decay function. Instead of counting each past event for the similarity statistic, each event is reduced depending on the time that passed between the current event and the past event. For sender similarity: Each target that two actors have in common is weighted by the time that passed between the two events. For target similarity: Each sender that two targets have in common is weighted by the time that passed between the two events.
eventtypevar	An optional dummy variable that represents the type of the event. If specified, only past events are considered for the count that reflect the same type as the current event.
eventattributevar	An optional variable that filters past events by the eventattributevalue specified.
eventattributevalue	A string that represents an event attribute by which all past events have to be filtered by.
variablename	An optional value (or values) with the name the similarity statistic variable should be given. To be used if returnData = TRUE.
returnData	TRUE/FALSE. Set to FALSE by default. The new variable(s) are bound directly to the data.frame provided and the data frame is returned in full.
showprogressbar	TRUE/FALSE. To be implemented.

Details

The `similarityStat()`-function calculates an endogenous statistic that measures whether sender (or targets) have a tendency to cluster together. Two distinct types of similarity measures can be calculated: sender similarity or target similarity.

Sender similarity: How many targets does the current sender have in common with senders who used the current target in the past? How likely is it that two senders are alike?

The function proceeds as follows:

1. First it filters out all the targets that the present sender a used in the past
2. Next it filters out all the senders that have also used the current target b
3. For each of the senders found in (2) it compiles a list of targets that this sender has used in the past
4. For each of the senders found in (2) it cross-checks the two lists generated in (1) and (3) and count how many targets the two senders have in common.

Target similarity: How many senders have used the same two concepts that the current sender has used (in the past and is currently using)? For each target that the current sender has used in the past, how many senders have also used these past targets as well as the current target? How likely is it that two targets are used together?

The function proceeds as follows:

1. First filter out all the targets that the current sender a has used in the past
2. Next it filters out all the senders that have also used the current target b

3. For each target found in (1) it compiles a list of senders that have also used this target in the past
4. For each target found in (1) it cross-checks the list of senders that have used b (found under (2)) and the list of senders that also used one other target that a used (found under (3))

Two decay functions may be used in the calculation of the similarity score for each event.

Author(s)

Laurence Brandenberger <laurence.brandenberger@eawag.ch>

See Also

[rem-package](#)

Examples

```
# create some data with 'sender', 'target' and a 'time'-variable
# (Note: Data used here are random events from the Correlates of War Project)
sender <- c('TUN', 'NIR', 'NIR', 'TUR', 'TUR', 'USA', 'URU',
           'IRQ', 'MOR', 'BEL', 'EEC', 'USA', 'IRN', 'IRN',
           'USA', 'AFG', 'ETH', 'USA', 'SAU', 'IRN', 'IRN',
           'ROM', 'USA', 'USA', 'PAN', 'USA', 'USA', 'YEM',
           'SYR', 'AFG', 'NAT', 'NAT', 'USA')
target <- c('BNG', 'ZAM', 'JAM', 'SAU', 'MOM', 'CHN', 'IRQ',
           'AFG', 'AFG', 'EEC', 'BEL', 'ITA', 'RUS', 'UNK',
           'IRN', 'RUS', 'AFG', 'ISR', 'ARB', 'USA', 'USA',
           'USA', 'AFG', 'IRN', 'IRN', 'IRN', 'AFG', 'PAL',
           'ARB', 'USA', 'EEC', 'BEL', 'PAK')
time <- c('800107', '800107', '800107', '800109', '800109',
         '800109', '800111', '800111', '800111', '800113',
         '800113', '800113', '800114', '800114', '800114',
         '800116', '800116', '800116', '800119', '800119',
         '800119', '800122', '800122', '800122', '800124',
         '800125', '800125', '800127', '800127', '800127',
         '800204', '800204', '800204')
type <- sample(c('cooperation', 'conflict'), 33,
              replace = TRUE)
important <- sample(c('important', 'not important'), 33,
                  replace = TRUE)

# combine them into a data.frame
dt <- data.frame(sender, target, time, type, important)

# create event sequence and order the data
dt <- eventSequence(datevar = dt$time, dateformat = '%y%m%d',
                  data = dt, type = 'continuous',
                  byTime = 'daily', returnData = TRUE,
                  sortData = TRUE)

# average sender similarity
dt$s.sim.av <- similarityStat(data = dt,
```

triadStat	<i>Calculate triad statistics</i>
-----------	-----------------------------------

Description

Calculate the endogenous network statistic triads that measures the tendency for events to close open triads.

Usage

```
triadStat(data, time, sender, target, halflife,
          weight = NULL,
          eventtypevar = NULL,
          eventtypevalues = NULL,
          eventattributevar = NULL,
          eventattributeAI = NULL,
          eventattributeBI = NULL,
          eventattributeAB = NULL,
          variablename = 'triad',
          returnData = FALSE,
          showprogressbar = FALSE)
```

Arguments

data	A data frame containing all the variables.
time	Numeric variable that represents the event sequence. The variable has to be sorted in ascending order.
sender	A string (or factor or numeric) variable that represents the sender of the event.
target	A string (or factor or numeric) variable that represents the target of the event.
halflife	A numeric value that is used in the decay function. The vector of past events is weighted by an exponential decay function using the specified halflife. The halflife parameter determines after how long a period the event weight should be halved. E.g. if halflife = 5, the weight of an event that occurred 5 units in the past is halved. Smaller halflife values give more importance to more recent events, while larger halflife values should be used if time does not affect the sequence of events that much.
weight	An optional numeric variable that represents the weight of each event. If weight = NULL each event is given an event weight of 1.
eventtypevar	An optional dummy variable that represents the type of the event. Use eventtypevalues to specify how the eventtypevar should be used to filter past events. Specifying the eventtypevar is needed to calculate effects of social balance theory, such as 'friend-of-friend' or 'enemy-of-enemy' statistics.
eventtypevalues	Two string values that represent the type of the past events. The first string value represents the eventtype that exists for all past events that include the current

sender (either as sender or target) and a third actor. The second value represents the eventtype for all past events that include the target (either as sender or target) as well as the third actor. An example: Let the eventtypevar indicate whether an event is of cooperative or hostile nature. To test whether the hypothesis 'the friend of my friend is my friend' holds, both eventtypevalues must be the same and point to the cooperative type (e.g. eventtypevalues = c("cooperation", "cooperation")) depending on how the eventtypevar is coded. To test whether the hypothesis 'the friend of my enemy is my enemy' holds, the first value in eventtypevalues represents the hostile event between current sender and a third actor and the second value represents the cooperative event between the third actor and the target. To test the hypothesis 'the enemy of my enemy is my friend', the first value represents the hostile events between current sender and a third actor and the second value represents the hostile event between the current target and the third actor. For the fourth hypothesis, to test social balance theory 'the enemy of my friend is my enemy', the first value represents a cooperative event between the current sender and a third actor and the second value represents a hostile event between the current target and the third actor.

eventattributevar

An optional string (or factor or numeric) variable that can be used to filter past and current events. Use eventattributeAI, eventattributeBI or eventattributeAB to specify which past events should be filtered and by what value.

eventattributeAI

An optional value used to specify how past events should be filtered depending on their attribute. Each distinct edge that form a triad can be filtered. eventattributeAI refers to the past event involving the current sender (a) and a third actor (i). eventattributeBI refers to past events involving target (b) and the third actor (i). eventattributeAB refers to the current event involving sender (a) and target (b).

eventattributeBI

see eventattributeAI.

eventattributeAB

see eventattributeAI.

variablename

An optional value (or values) with the name the triad statistic variable should be given. To be used if returnData = TRUE.

returnData

TRUE/FALSE. Set to FALSE by default. The new variable is bound directly to the data.frame provided and the data frame is returned in full.

showprogressbar

TRUE/FALSE. To be implemented.

Details

The triadStat()-function calculates an endogenous statistic that measures whether events have a tendency to form closing triads.

The effect is calculated as follows:

$$G_t = G_t(E) = (A, B, w_t),$$

G_t represents the network of past events and includes all events E . These events consist each of a sender $a \in A$ and a target $b \in B$ and a weight function w_t :

$$w_t(i, j) = \sum_{e:a=i,b=j} |w_e| \cdot e^{-\frac{(t-t_e) \cdot \ln(2)}{T_{1/2}}} \cdot \frac{\ln(2)}{T_{1/2}},$$

where w_e is the event weight (usually a constant set to 1 for each event), t is the current event time, t_e is the past event time and $T_{1/2}$ is a half-life parameter.

For the triad effect, the past events G_t are filtered to include only events where the current event closes an open triad in the past.

$$triad(G_t, a, b) = \sqrt{\sum_{i \in A} w_t(a, i) \cdot w_t(i, b)}$$

An exponential decay function is used to model the effect of time on the endogenous statistics. The further apart the past event is from the present event, the less weight is given to this event. The half-life parameter in the `triadStat()`-function determines at which rate the weights of past events should be reduced. Therefore, if the one (or more) of the two events in the triad have occurred further in the past, less weight is given to this triad because it becomes less likely that the sender and target actors reacted to each other in the way the triad assumes.

The `eventtypevar`- and `eventattributevar`-options help filter the past events more specifically. How they are filtered depends on the `eventtypevalue`- and `eventattributevalue`-option.

Author(s)

Laurence Brandenberger <laurence.brandenberger@eawag.ch>

See Also

[rem-package](#)

Examples

```
# create some data with 'sender', 'target' and a 'time'-variable
sender <- c('TUN', 'UNK', 'NIR', 'TUR', 'TUR', 'USA', 'URU',
           'IRQ', 'MOR', 'BEL', 'EEC', 'USA', 'IRN', 'IRN',
           'USA', 'AFG', 'ETH', 'USA', 'SAU', 'IRN', 'IRN',
           'ROM', 'USA', 'USA', 'PAN', 'USA', 'USA', 'YEM',
           'SYR', 'AFG', 'NAT', 'UNK', 'IRN')
target <- c('BNG', 'RUS', 'JAM', 'SAU', 'MOM', 'CHN', 'IRQ',
           'AFG', 'AFG', 'EEC', 'BEL', 'ITA', 'RUS', 'UNK',
           'IRN', 'RUS', 'AFG', 'ISR', 'ARB', 'USA', 'USA',
           'USA', 'AFG', 'IRN', 'IRN', 'IRN', 'AFG', 'PAL',
           'ARB', 'USA', 'EEC', 'IRN', 'CHN')
time <- c('800107', '800107', '800107', '800109', '800109',
         '800109', '800111', '800111', '800111', '800113',
         '800113', '800113', '800114', '800114', '800114',
         '800116', '800116', '800116', '800119', '800119',
         '800119', '800122', '800122', '800122', '800124',
```


Index

`as.Date`, 8

`degree (degreeStat)`, 2

`degreeStat`, 2, 2

`event sequence (eventSequence)`, 7

`event.sequence (eventSequence)`, 7

`eventSequence`, 7

`flexsurvreg`, 22

`fourCycle (fourCycleStat)`, 10

`fourCycleStat`, 2, 10

`indegree (degreeStat)`, 2

`inertia (inertiaStat)`, 14

`inertiaStat`, 2, 14

`outdegree (degreeStat)`, 2

`reciprocity (reciprocityStat)`, 17

`reciprocityStat`, 2, 17

`relational event model (rem-package)`, 2

`relational-event-model (rem-package)`, 2

`rem (rem-package)`, 2

`rem-package`, 2, 6, 8, 12, 16, 20, 22, 27, 31

`remRate`, 2, 21

`similarity (similarityStat)`, 24

`similarityStat`, 2, 24

`triad (triadStat)`, 29

`triadStat`, 2, 29