

Package ‘coffee’

September 26, 2024

Type Package

Title Chronological Ordering for Fossils and Environmental Events

Version 0.4.1

Maintainer Maarten Blaauw <maarten.blaauw@qub.ac.uk>

Description While individual calibrated radiocarbon dates can span several centuries, combining multiple dates together with any chronological constraints can make a chronology much more robust and precise. This package uses Bayesian methods to enforce the chronological ordering of radiocarbon and other dates, for example for trees with multiple radiocarbon dates spaced at exactly known intervals (e.g., 10 annual rings). For methods see Christen 2003 <doi:10.11141/ia.13.2>. Another example is sites where the relative chronological position of the dates is taken into account - the ages of dates further down a site must be older than those of dates further up (Buck, Kenworthy, Litton and Smith 1991 <doi:10.1017/S0003598X00080534>; Nicholls and Jones 2001 <doi:10.1111/1467-9876.00250>). The paper accompanying this R package is Blaauw et al. 2024 <doi:10.1017/RDC.2024.56>.

License GPL (>= 2)

Encoding UTF-8

URL <https://github.com/Maarten14C/coffee>

RoxygenNote 7.3.2

NeedsCompilation no

Language en-GB

VignetteBuilder knitr

Suggests knitr, rmarkdown, utf8

Imports data.table, rintcal (>= 1.0.0)

Depends rice (>= 0.2.0)

Author Maarten Blaauw [aut, cre] (<<https://orcid.org/0000-0002-5680-1515>>),
Marco A. Aquino Lopez [aut, ctb]
(<<https://orcid.org/0000-0002-5076-7205>>),
J. Andres Christen [aut, ctb, cph]
(<<https://orcid.org/0000-0002-5795-4345>>)

Repository CRAN

Date/Publication 2024-09-25 23:20:12 UTC

Contents

ages.undated	2
draw.rings	3
draw.strat	6
IAT	9
rings	10
scissors	12
sim.rings	13
sim.strat	15
strat	16
strat.cleanup	21
thinner	21
withinhpd	22

Index	23
--------------	-----------

ages.undated	<i>Model ages between two dated levels</i>
--------------	--

Description

Model ages of undated levels, by for each MCMC iteration finding the age of the layer above and of the layer below, and sampling a random age from a uniform distribution between the age estimates of the two ages.

Usage

```
ages.undated(position, set = get("info"), draw = TRUE)
```

Arguments

position	Position of the to-be-estimated undated layer. Should be larger than the layer above but smaller than the layer below it.
set	This option reads the 'info' variable, which contains the data and the model output.
draw	Whether or not to draw the age distribution.

`draw.rings`*plot the dates and model of a wiggle-match dated tree*

Description

A plot with two panels. The top panel shows the calibrated distributions (in blue) and the wiggle-match age-modelled age estimates for each dated ring (grey). The bottom panel shows the fit of the uncalibrated radiocarbon dates (steelblue dots and lab error bars) to the calibration curve (green), as well as the age estimate of the oldest/starting ring (grey) and its hpd range (black).

Usage

```
draw.rings(  
  name = "mytree",  
  tree.dir = "trees",  
  sep = ",",  
  normal = TRUE,  
  dat = c(),  
  out = c(),  
  cc = 1,  
  postbomb = FALSE,  
  BCAD = FALSE,  
  t.a = 3,  
  t.b = 4,  
  x.lim = c(),  
  x1.axis = TRUE,  
  x1.labels = FALSE,  
  x1.lab = c(),  
  rev.x = FALSE,  
  y1.lab = c(),  
  y1.lim = c(),  
  y2.lim = c(),  
  x2.lab = c(),  
  y2.lab = c(),  
  ex = 0.05,  
  plot.cc = TRUE,  
  plot.dists = TRUE,  
  mar.1 = c(1, 3, 1, 1),  
  mar.2 = c(3, 3, 0, 1),  
  mgp = c(1.7, 0.7, 0),  
  dist.res = 500,  
  date.col = "steelblue",  
  cc.col = rgb(0, 0.5, 0, 0.5),  
  dist.col = rgb(0, 0, 0, 0.5),  
  calib.col = rgb(0, 0, 1, 0.25),  
  range.col = "black",  
  set.layout = TRUE
```

)

Arguments

name	The name of the tree. The .csv file should be saved under a folder named exactly the same as name, and the folder should live under the tree.dir folder. Default names include Ulandryk4 and mytree.
tree.dir	The directory where the folders of the individual trees live. Defaults to treedir="trees".
sep	Separator for the fields in the .csv file. Defaults to a comma.
normal	Calculations can be done assuming that the measurements are normally distributed. By default this is set to FALSE and a student-t distribution is used (Christen and Perez 2009).
dat	If plot.rings is called from within rings(), both dat and out are provided. If an existing run has to be plotted again, dat and out are read from the files in the folder named name.
out	If plot.rings is called from within rings(), both dat and out are provided. If an existing run has to be plotted again, dat and out are read from the files in the folder named name.
cc	Calibration curve to be used. Could be 1 (IntCal20; default), 2 (Marine20), 3 (SHCal20) or 4 (custom curve).
postbomb	Negative C-14 ages should be calibrated using a postbomb curve. This could be 1 (northern-hemisphere region 1), 2 (NH region 2), 3 (NH region 3), 4 (southern hemisphere regions 1-2), or 5 (SH region 3).
BCAD	The calendar scale of graphs and age output-files is in cal BP by default, but can be changed to BC/AD using BCAD=TRUE.
t.a	First parameter for the student-t distribution (defaults to 3; higher numbers make the distribution approximate the normal distribution more).
t.b	Second parameter for the student-t distribution (defaults to 4; higher numbers make the distribution approximate the normal distribution more).
x.lim	Limits of the x-axes. Calculated automatically by default.
x1.axis	Whether or not to plot the upper x-axis (slightly redundant since the bottom axis shows the values already). Defaults to TRUE.
x1.labels	Whether or not to plot the labels of the upper x-axis. (slightly redundant since the bottom axis shows the values already). Defaults to FALSE.
x1.lab	The labels for the calendar axis of the upper panel. Defaults to empty.
rev.x	Whether or not to reverse the x-axis. Defaults to FALSE.
y1.lab	The labels for the y-axis. Defaults to "rings".
y1.lim	Limits of the top y-axis. Calculated automatically if left empty.
y2.lim	Limits of the bottom y-axis. Calculated automatically if left empty.
x2.lab	The labels for the bottom calendar axis (default age.lab="cal BP" or "BC/AD" if BCAD=TRUE).
y2.lab	The labels for the bottom y-axis. Defaults to 14C BP with superscript 14, so expression("" ¹⁴ C~BP).

ex	Exaggeration of the heights of the calibrated distributions. Defaults to 0.05 so there is plenty space to plot many distributions.
plot.cc	Whether or not to plot a panel with the calibration curve.
plot.dists	Whether or not to plot a panel with the distributions.
mar.1	Margins around the first/top plot.
mar.2	Margins around the first/bottom plot.
mgp	Axis text margins (where should titles, labels and tick marks be plotted). Defaults to <code>mgp=c(1.7, .7, .0)</code> .
dist.res	Resolution of the plot of the distribution. The default is 500, which results in smooth plots.
date.col	Colour of the uncalibrated dates when plotted on top of the calibration curve. Defaults to "steelblue".
cc.col	Colour of the calibration curve. Defaults to semi-transparent darkgreen, <code>cc.col=rgb(0.5, 0, 0.5)</code> .
dist.col	Colour of the age-modelled distribution. Defaults to semi-transparent grey, <code>dist.col=rgb(0,0,0,0.5)</code> .
calib.col	Colour of the calibrated distributions. Defaults to semi-transparent blue, <code>dist.col=rgb(0,0,1,0.5)</code> .
range.col	Colour of the hpd ranges. Defaults to "black".
set.layout	By default, the layout of the two plots is set automatically (2 plots in one column).

Value

A plot with the calibrated distributions of the individual dates and the wiggle-match distributions (top), and the dates on the calibration curve together with the age distribution for the earliest ring, 0.

Author(s)

Maarten Blaauw, J. Andres Christen

Examples

```
treedir <- tempdir()
rings("Ulandryk", tree.dir=treedir, draw=FALSE)
draw.rings("Ulandryk", tree.dir=treedir)
```

draw.strat

plot the dates and model of chronologically ordered dated depths

Description

A plot with two panels. The top panel shows the MCMC output. The bottom panel shows the individually calibrated dates (in downward light gray) as well as the modelled ages constrained by chronological ordering (upward dark-grey) and lines with the hpd ranges (black). Any similarity with swimming elephants or island chains is coincidental.

Usage

```
draw.strat(
  name = "mystrat",
  set = get("info"),
  structure = set$struc,
  y.scale = "positions",
  strat.dir = "strats",
  cc.dir = c(),
  sep = ", ",
  postbomb = FALSE,
  calibrated.ex = c(),
  calibrated.mirror = FALSE,
  calibrated.up = TRUE,
  modelled.ex = c(),
  modelled.mirror = FALSE,
  modelled.up = FALSE,
  BCAD = FALSE,
  threshold = 0.001,
  xtop.lab = c(),
  ytop.lab = c(),
  xbottom.lab = c(),
  ybottom.lab = "position",
  calibrated.col = rgb(0, 0, 0, 0.2),
  calibrated.border = NA,
  calBP.col = rgb(0, 0, 0, 0.2),
  calBP.border = NA,
  modelled.col = rgb(0, 0, 0, 0.5),
  modelled.border = rgb(0, 0, 0, 0.5),
  range.col = "black",
  block.col = rgb(0, 0, 1, 0.05),
  gap.col = "blue",
  gap.pos = 1,
  simulation = FALSE,
  simulation.col = grey(0.5),
  pos.lim = c(),
  age.lim = c(),
```

```

    mgp = c(2, 0.7, 0),
    mar.top = c(3, 3, 1, 1),
    mar.bottom = c(3, 3, 0.5, 1),
    heights = c(0.3, 0.7),
    iterations.warning = TRUE,
    min.its = 1000,
    warning.loc = "bottomleft",
    warning.col = "red"
)

```

Arguments

name	Name of the stratigraphy dataset. Defaults to "mystrat".
set	This option reads the 'info' variable, which contains the data and the model output.
structure	Information about the structure (e.g., blocks, gaps, dates, undated levels) of the dataset. Filled automatically.
y.scale	The scale of the vertical axis of the main plot. This can be the positions of the dated levels ('positions') or their position order ('dates').
strat.dir	The directory where the folders of the individual stratigraphies live. Defaults to start.dir="strats".
cc.dir	Directory of calibration curve. Keep empty for the default value.
sep	Separator for the fields in the .csv file. Defaults to a comma.
postbomb	Negative C-14 ages should be calibrated using a postbomb curve. This could be 1 (northern-hemisphere region 1), 2 (NH region 2), 3 (NH region 3), 4 (southern hemisphere regions 1-2), or 5 (SH region 3).
calibrated.ex	Exaggeration of the heights of the calibrated distributions. Calculated automatically by default. Note that more precise dates peak higher than dates with lower precision.
calibrated.mirror	Whether or not the individually calibrated (but not the modelled) distributions should be drawn both up and down, quite a bit like fish or swans. Defaults to FALSE.
calibrated.up	Whether the calibrated distributions should be drawn upward or downward (the default, resembling the reflections of islands in the sea, or swimming animals if you wish)
modelled.ex	Exaggeration of the heights of the age-modelled distributions. Calculated automatically by default. Note that more precise ages peak higher than ages with lower precision.
modelled.mirror	Whether or not the age-modelled distributions should be drawn both up and down, quite a bit like fish or swans. Defaults to FALSE.
modelled.up	Whether the age-modelled distributions should be drawn downward or upward (the default, resembling islands in the sea)

BCAD	The calendar scale of graphs and age output-files is in cal BP by default, but can be changed to BC/AD using BCAD=TRUE.
threshold	Value below which probabilities should not be drawn any more. Defaults to 0.001 of the distribution's peak.
xtop.lab	The label for the x-axis of the top panel showing the MCMC run. Defaults to "iterations".
ytop.lab	The label for the y-axis of the top panel showing the MCMC run. Defaults to "energy".
xbottom.lab	The label for the x-axis of the bottom panel showing the age-model output. Defaults to "cal BP" or "BC/AD".
ybottom.lab	The label for the y-axis of the bottom panel showing the age-model output. Defaults to "position".
calibrated.col	Colour of the inside of the unmodelled, calibrated ages. Defaults to semi-transparent light grey, <code>rgb(0,0,0,0.5)</code> .
calibrated.border	Colour of the border of the unmodelled non-14C ages. Defaults to nothing, NA.
calBP.col	Colour of the inside of the unmodelled non-14C ages. Defaults to semi-transparent light grey, <code>rgb(0,0,0,0.5)</code> .
calBP.border	Colour of the border of the unmodelled, calibrated ages. Defaults to nothing, NA.
modelled.col	Colour of the inside of the modelled ages. Defaults to semi-transparent dark grey, <code>rgb(0,0,0,0.5)</code> .
modelled.border	Colour of the border of the modelled ages. Defaults to semi-transparent dark grey, <code>rgb(0,0,0,0.5)</code> .
range.col	Colour of the hpd ranges. Defaults to "black".
block.col	Colour of the field indicating unordered dates within a 'block'. Defaults to light-blue, <code>rgb(0,0,1,.05)</code> .
gap.col	Colour of the diagonal lines and parameters of any gaps.
gap.pos	Plotting position of the gap information.
simulation	Whether or not the data are part of a simulated stratigraphy.
simulation.col	If the data are part of a simulated stratigraphy, the 'true' ages can also be plotted.
pos.lim	Limit of the main y-axis.
age.lim	Limit of the main x-axis.
mgp	Axis text margins (where should titles, labels and tick marks be plotted). Defaults to <code>mgp=c(1.7, .7, .0)</code> .
mar.top	Margins around the top panel. Defaults to <code>mar.top=c(3,3,1,1)</code> .
mar.bottom	Margins around the bottom panel. Defaults to <code>mar.bottom=c(3,3,0.5,1)</code> .
heights	Relative heights of the two panels in the plot. Defaults to 0.3 for the top and 0.7 for the bottom panel.

<code>iterations.warning</code>	Whether or not to plot a warning if there are < 3000 iterations, too few for a reliable MCMC run.
<code>min.its</code>	The minimum amount of iterations, below which a warning is printed (if <code>iterations.warning=TRUE</code>). Defaults to 1,000.
<code>warning.loc</code>	Location of the warning
<code>warning.col</code>	Colour of the warning - defaults to red.

Value

A plot with two panels showing the MCMC run and the calibrated and modelled ages.

Author(s)

Maarten Blaauw

IAT *calculate the Integrated Autocorrelation Time*

Description

Calculate the Integrated Autocorrelation Time, which gives the proposed value for thinning. E.g., if the IAT is 80, it is good to thin the MCMC run by storing only every 80 iterations. This method is slower than `GetAutoCov`, but much better.

Usage

```
IAT(set, par = 0, from = 1, to)
```

Arguments

<code>set</code>	This option reads the 'info' variable, which contains the data and the model output.
<code>par</code>	The parameter to test. Defaults to 0.
<code>from</code>	The first of the iterations of the MCMC run to check. Defaults to the first one.
<code>to</code>	The last of the iterations of the MCMC run to check. Defaults to the last one.

Value

The IAT value

Author(s)

Andres Christen

rings

*wiggle-match C-14 dating of a tree***Description**

Produce a Bayesian wiggle-match date of a tree dated with multiple C-14 dates at exactly known spacing (e.g., every 10 tree-ring years).

Usage

```
rings(
  name = "Ulandryk",
  tree.dir = "trees",
  sep = ",",
  normal = FALSE,
  delta.R = 0,
  delta.STD = 0,
  t.a = 3,
  t.b = 4,
  ask = TRUE,
  age.steps = 1,
  cutoff = 1e-06,
  prob = 0.95,
  cc = 1,
  postbomb = FALSE,
  BCAD = FALSE,
  times = 3,
  talk = TRUE,
  draw = TRUE,
  ...
)
```

Arguments

name	Name of the tree. The .csv file should be saved under a folder named exactly the same as name, and the folder should live under the <code>tree.dir</code> folder. The default is Ulandryk.
tree.dir	The directory where the folders of the individual trees live. Defaults to <code>tree.dir="trees"</code> .
sep	Separator for the fields in the .csv file. Defaults to a comma.
normal	Calculations can be done assuming that the measurements are normally distributed. By default this is set to FALSE and a student-t distribution is used (Christen and Perez 2009)
delta.R	The ages can be modelled to have an offset. The mean is 0 by default.
delta.STD	The error of the offset. Set to 0 by default.
t.a	First parameter for the student-t distribution (defaults to 3; higher numbers make the distribution approximate the normal distribution more).

t.b	Second parameter for the student-t distribution (defaults to 4; higher numbers make the distribution approximate the normal distribution more).
ask	Whether or not to ask if new folders should be written (if required)
age.steps	Steps in years for the calculations. Defaults to 1, every year.
cutoff	Value below which probabilities are no longer taken into account. Defaults to 0.000001.
prob	After the run, a fit of the model with the dates is calculated, as the ratio of model iterations that fit the hpd ranges of the dates. Defaults to the prob=0.95 hpd ranges.
cc	Calibration curve to be used, for glueing to a postbomb curve. Could be 1 (IntCal20; default), 2 (Marine20), 3 (SHCal20) or 4 (custom curve). Normally not used, except in the case where there are postbomb dates (requiring the 'glueing' of pre- and postbomb curves).
postbomb	Negative C-14 ages should be calibrated using a postbomb curve. This could be 1 (northern-hemisphere region 1), 2 (NH region 2), 3 (NH region 3), 4 (southern hemisphere regions 1-2), or 5 (SH region 3).
BCAD	The calendar scale of graphs and age output-files is in cal BP by default, but can be changed to BC/AD using BCAD=TRUE.
times	The range of years to be calculated, as multiples of the uncertainties of the data points. E.g. if the lab error of the oldest date is 20 years, and times is set to 5, the calculation range will be extended by 20*5 years.
talk	Whether or not to provide feedback on folders written into.
draw	Whether or not to draw plots.
...	Options for the plot. See draw.rings.

Details

The calculations are based on Bwigg (Christen and Litton 1995; Christen 2003). In OxCal, this is called a D_Sequence (Bronk Ramsey et al. 2001).

Since only one parameter has to be estimated (the age of the earliest, innermost ring), a MCMC approach is not necessary nor recommended, and results are calculated analytically.

The tree files should be in plain-text and fields separated by commas, and the file's extension should be ".csv". The files should start with a line contain the following headers: "lab ID", "C-14 age", "error", "ring", "cc", separated by commas. Then each row should have the corresponding values, also separated by commas. Rings are counted from the inner ring (0 year old) outwards, so, forward in time. The file should start with the youngest rings, then work downward until reaching the oldest, bottommost dated rings. cc should either be 1 (IntCal20; northern hemisphere terrestrial), 2 (Marine20, though we've never heard of marine trees), 3 (SHCal20; southern hemisphere) or 4 (custom curve).

The default tree is called Ulandryk (Kuzman et al. 2004). As an alternative, a tree can be simulated (see sim.tree()).

By default, the data are calibrated assuming a student-t distribution, which has wider tails than the normal distribution and deals well with scatter and outliers.

Value

the probabilities for the relevant calendar years.

Author(s)

Maarten Blaauw, J. Andres Christen

References

Bronk Ramsey C, van der Plicht J, Weninger B, 2001. 'Wiggle matching' radiocarbon dates. Radiocarbon 43, 381–389.

Christen JA, Litton CD, 1995. A Bayesian approach to wiggle-matching. Journal of Archaeological Science 22, 719-725 doi:10.1016/03054403(95)900020

Christen JA, 2003. Bwigg: An Internet facility for Bayesian radiocarbon wiggle-matching. Internet Archaeology 13. doi:10.11141/ia.13.2

Christen JA, Perez S, 2009. A new robust statistical model for radiocarbon data. Radiocarbon 51, 1047-1059

Kuzmin Y, Slusarenko I, Hajdas I, Bonani G, Christen JA. 2004. The comparison of 14C wiggle-matching results for the 'floating' tree-ring chronology of the Ulandryk-4 Burial Ground (Altai Mountains, Siberia). Radiocarbon 46, 943–948.

Examples

```
rings("Ulandryk", tree.dir=tempdir())
```

scissors

Remove the first n iterations.

Description

Removes iterations of the MCMC time series, and then updates the output files.

Usage

```
scissors(burnin, set = get("info"), write = TRUE, save.info = TRUE)
```

Arguments

burnin Number of iterations to remove of the iterative time series. If this value is higher than the amount of remaining iterations, a warning is given and the iterations are not removed. If the provided number is negative, the iterations will be removed from the end of the run, not from the start. If a range is given, this range of iterations is removed.

set Detailed information of the current run, stored within this session's memory as variable info.

write	Whether or not to write the changes to the output file. Defaults to TRUE.
save.info	Whether or not to store a variable 'info' in the session which contains the run input, output and settings. Defaults to save.info=TRUE.

Details

The strat function will perform thousands millions of MCMC iterations, although usually only a fraction of these will be stored. The remaining MCMC iterations should be well mixed (the upper panel of the fit of the iterations should show no undesirable features such as trends or sudden systematic drops or rises). If the run has a visible remaining burn-in, scissors can be used to remove them. To remove, e.g., the first 300 iterations, type `scissors(300)`. To remove the last 300 iterations, type `scissors(-300)`. To remove iterations 300 to 600, type `scissors(300:600)`.

Value

NA

sim.rings	<i>Simulate the radiocarbon dating of tree-rings</i>
-----------	--

Description

Simulate the dense radiocarbon dating of a tree or other deposit with exactly known yearly rings, and thus with gaps of exactly known age. The radiocarbon dates are assumed to have a degree of lab error and scatter. A (constant) offset can also be modelled.

Usage

```
sim.rings(
  name = "mytree",
  age.start = 1000,
  length = 400,
  gaps = 20,
  offset = 0,
  scatter = 1.05,
  error = 0.03,
  min.error = 15,
  tree.dir = "trees",
  sep = ", ",
  cc = 1,
  postbomb = FALSE,
  ask = TRUE
)
```

Arguments

name	Name of the simulated tree-ring set. Defaults to "mytree".
age.start	Starting age of the simulated tree.
length	Length of the sequence (if gaps are given as constant). Could be set at e.g 400 for an oak, but many trees will not live as long so shorter sequences could also make sense.
gaps	How many calendar years there are between the dated rings. Can be set constant (one value, e.g. 20), or alternatively the gaps can be provided as a list of values.
offset	The ages could be offset by some constant value. Defaults to 0.
scatter	There is always a degree of scatter between measurements, and the amount of scatter can be modelled using, e.g., $scatter=2*error$. Set at 0 to model radiocarbon dates that are 100% faithful to the calibration curve (very unlikely!).
error	Laboratory error of the radiocarbon dates as percentage of the mean. Defaults to 0.02.
min.error	Minimum laboratory to be reported. Defaults to 10 (C-14 year).
tree.dir	The directory where the folders of the individual trees live. Defaults to <code>tree.dir="trees"</code> .
sep	Separator for the fields in the .csv file. Defaults to a comma.
cc	Calibration curve to be used. Could be 1 (IntCal20; default), 2 (Marine20), 3 (SHCal20) or 4 (custom curve).
postbomb	Negative C-14 ages (younger than 0 cal BP or AD 1950) should be modelled using a postbomb curve. This could be 1 (northern-hemisphere region 1), 2 (NH region 2), 3 (NH region 3), 4 (southern hemisphere regions 1-2), or 5 (SH region 3).
ask	Ask if a folder may be made and files written into it

Value

A file containing 5 columns: the simulated calendar ages, the radiocarbon ages, their errors, the rings (starting with the youngest year and working backward in time), and the calibration curve to be used.

Author(s)

Maarten Blaauw, J. Andres Christen

Examples

```
treedir <- tempdir()
sim.rings("manyrings", age.start=1000, length=400, gaps=10, tree.dir=treedir)
rings("manyrings", tree.dir=treedir)
```

sim.strat	<i>Simulate the radiocarbon dating of random depths of a sediment which has accumulated over time.</i>
-----------	--

Description

Simulate the radiocarbon dating (or with dates that are already on the cal BP scale) of a deposit that is known to have accumulated over time, and for which therefore the dated depths can be safely assumed to be in chronological order.

Usage

```
sim.strat(
  name = "mystrat",
  age.min = 4321,
  length = 800,
  n = 5,
  offset = 0,
  scatter = 2 * error,
  error = 0.02,
  min.error = 10,
  rounded = 0,
  strat.dir = "strats",
  sep = ", ",
  cc = 1,
  postbomb = FALSE,
  ask = TRUE
)
```

Arguments

name	Name of the simulated stratigraphy. Defaults to "mystrat".
age.min	Minimum age of the simulation.
length	Length of the sequence.
n	The amount of dated depths.
offset	The ages could be offset by some constant value. Defaults to 0.
scatter	There is always a degree of scatter between measurements, and the amount of scatter can be modelled using, e.g., scatter=2*error. Set at 0 to model radiocarbon dates that are 100% faithful to the calibration curve (very unlikely!).
error	Laboratory error of the radiocarbon dates as percentage of the mean. Defaults to 0.02.
min.error	Minimum laboratory to be reported. Defaults to 10 (C-14 year).
rounded	Rounding of the simulated calendar years. Rounds to single years (0 decimals) by default.

strat.dir	The directory where the folders of the individual trees live. Defaults to tree.dir="trees".
sep	Separator for the fields in the .csv file. Defaults to a comma.
cc	Calibration curve to be used. Could be 1 (IntCal20; default), 2 (Marine20), 3 (SHCal20) or 4 (custom curve).
postbomb	Negative C-14 ages (younger than 0 cal BP or AD 1950) should be modelled using a postbomb curve. This could be 1 (northern-hemisphere region 1), 2 (NH region 2), 3 (NH region 3), 4 (southern hemisphere regions 1-2), or 5 (SH region 3).
ask	Ask if a folder may be made and files written into it

Details

Dates further down the sequence should have older ages than dates further up, even though owing to scatter, the dates themselves might not be in exact chronological order. The modelling is performed in a Bayesian framework (see strat). The amount of scatter, the laboratory error and an offset can also be modelled.

Value

A file containing 5 columns: the simulated calendar ages, the radiocarbon ages, their errors, their relative positions (starting with the youngest one at top, and counting upward going down the sequence), and the calibration curve to be used.

Author(s)

Maarten Blaauw, J. Andres Christen

Examples

```
stratdir <- tempdir()
sim.strat("ordered.mud", age.min=1000, length=5000, n=10, strat.dir=stratdir)
```

strat	<i>Model chronologically ordered dates</i>
-------	--

Description

Model radiocarbon dates (or dates that are already on the cal BP scale) of a deposit that is known to have accumulated over time, and for which therefore the dated positions can be safely assumed to be in chronological order.

Usage

```

strat(
  name = "mystrat",
  strat.dir = "strats",
  run = TRUE,
  its = 50000,
  burnin = 100,
  thinning = c(),
  internal.thinning = c(),
  min.its = 1000,
  write.MCMC = FALSE,
  MCMC.dir = tempdir(),
  remove.tmp = TRUE,
  init.ages = c(),
  ballpark.method = 2,
  y.scale = "dates",
  showrun = FALSE,
  sep = ", ",
  normal = FALSE,
  delta.R = 0,
  delta.STD = 0,
  t.a = 3,
  t.b = 4,
  cc = 1,
  cc.dir = c(),
  prob = 0.95,
  postbomb = FALSE,
  BCAD = FALSE,
  ask = FALSE,
  talk = TRUE,
  show.progress = TRUE,
  clean.garbage = TRUE,
  save.info = TRUE,
  age.span = c(),
  ...
)

```

Arguments

<code>name</code>	Name of the stratigraphy dataset. Defaults to "mystrat".
<code>strat.dir</code>	The directory where the folders of the individual stratigraphies live. Defaults to <code>strat.dir="strats"</code> .
<code>run</code>	Whether or not to run the data. If set to <code>FALSE</code> , will try to load existing run data.
<code>its</code>	Amount of iterations to be run. Setting this to low numbers (e.g., 1000) will result in fast but less stable and less reliable runs. Higher values will take longer but result in more stable and robust runs. Defaults to 50000. Aim to set this to

such values that at least 3000 iterations remain after removing the burnin and thinning.

burnin	Amount of iterations to remove at the start of the run. Defaults to 100.
thinning	After running all iterations, only some will be stored. For example, if thinning is set at the default 50, only every 50th MCMC iteration will be stored, and the others will be discarded. This is to remove the dependence between neighbouring MCMC iterations. Defaults to a value calculated from the MCMC run itself.
internal.thinning	Does internal thinning during the MCMC process, storing only every 'internal.thinning' MCMC iterations.
min.its	Minimum number of (remaining) iterations, below which a warning is given. Defaults for now to 1,000 iterations.
write.MCMC	Especially longer runs or sites with many dates can take up lots of memory. For such cases, MCMC iterations are stored in temporary files (and in the session's tempdir() folder) rather than in memory - however this could impact your computer as the files have to be written to repeatedly and they can become huge. Defaults to FALSE.
MCMC.dir	Directory where temporary MCMC files will be saved. If left empty, defaults to tempdir().
remove.tmp	If write.MCMC=TRUE, then by default the (often huge) temporary files are deleted after the run. Set to remove.tmp=FALSE to keep these temporary files. The files will be placed in the temporary folder of this R session - where this is can be checked by providing the command tempdir().
init.ages	By default, the ballpark age estimates to feed the MCMC are calculated automatically, however they can also be provided manually, as two rows of values (for x0 and xp0) which have to obey the assumptions (e.g., chronological ordering).
ballpark.method	Initial, ballpark values for the initial ages (if not provided by the option 'init.ages'). Can be 1 (based on a linear model) or 2 (based on sorted random draws).
y.scale	The scale of the vertical axis of the main plot. This can be the positions of the dated levels ('positions') or their position order ('dates').
showrun	Whether or not to show how the MCMC process is progressing during the run. Defaults to FALSE.
sep	Separator for the fields in the .csv file. Defaults to a comma.
normal	Calculations can be done assuming that the measurements are normally distributed. By default this is set to FALSE and a student-t distribution is used (Christen and Perez 2009)
delta.R	The ages can be modelled to have an offset. The mean is 0 by default.
delta.STD	The error of the offset. Set to 0 by default.
t.a	First parameter for the student-t distribution (defaults to 3; higher numbers make the distribution approximate the normal distribution more).
t.b	Second parameter for the student-t distribution (defaults to 4; higher numbers make the distribution approximate the normal distribution more).

cc	Calibration curve to be used, for glueing to a postbomb curve. Could be 1 (Int-Cal20; default), 2 (Marine20), 3 (SHCal20) or 4 (custom curve). Normally not used (since this information is best provided within the .csv files), except in the case where there are postbomb dates (requiring the 'gluing' of pre- and post-bomb curves).
cc.dir	Directory of calibration curve(s). Keep empty for the default value.
prob	After the run, a fit of the model with the dates is calculated, as the ratio of model iterations that fit the hpd ranges of the dates. Defaults to the prob=0.95 hpd ranges.
postbomb	Negative C-14 ages should be calibrated using a postbomb curve. This could be 1 (northern-hemisphere region 1), 2 (NH region 2), 3 (NH region 3), 4 (southern hemisphere regions 1-2), or 5 (SH region 3). Defaults to no postbomb curve (postbomb=FALSE).
BCAD	The calendar scale of graphs and age output-files is in cal BP by default, but can be changed to BC/AD using BCAD=TRUE. Needs more work probably.
ask	Whether or not to ask if a folder should be made (if required).
talk	Whether or not to provide feedback on folders written into and on what is happening.
show.progress	Whether or not to provide feedback on progress made with the run.
clean.garbage	Whether or not to clean up the memory 'garbage collection' after a run. Recommendable if you have many dates or long runs.
save.info	Whether or not to store a variable 'info' in the session which contains the run input, output and settings. Defaults to save.info=TRUE.
age.span	Expected age span. Defaults to run from the current year in AD to 55e3 which is the current cal BP limit for C-14 dates. If older, non-14C dates are present, age.span is set to the larger of the radiocarbon limit or twice the age of the oldest non-radiocarbon age.
...	Options for the plot. See draw.strat.

Details

Dates further down the sequence should have older ages than dates further up, even though owing to scatter, the dates themselves might not be in exact chronological order. The amount of scatter, the laboratory error and an offset can also be modelled. The function reads in a comma-separated-values (.csv) file of a specific format. The first column contains the names of the dates/information, the second column has the age(s) (uncalibrated for radiocarbon dates, as they will be calibrated during the modelling), the third column their errors, the fourth column their position (see below), and the fifth column cc, the calibration curve information. Additional columns for the reservoir effect (delta.R and delta.STD) and the student-t model (t.a and t.b) can be added, much like rbacon .csv files. The file should contain a header as the first row, named "lab ID", "age", "error", "position", and "cc" (with additional fields as per below if required). The extension of the file should be ".csv". The positions of the dates (column 4) should be entered with the topmost, youngest levels first, and then working downward toward the oldest levels. The topmost position gets the lowest number (e.g., 0), and each subsequent entry should have a higher position number to ensure that the levels are ordered in time. Dates in 'blocks' where there is no known age ordering between the dates in

the block (but where that block is known to be older than the level above it and younger than the level below it) should all get the same position in column 4. The function does not only deal with dates (radiocarbon or otherwise), but can also model undated levels and a range of gaps between the dated levels. This is done mostly through column 5 in the .csv files, where a 0 is for dates on the cal BP scale, 1 for radiocarbon dates that require calibration with IntCal20, 2 with Marine20, 3 with SHCal20, 4 a custom calibration curve; additional information can be provided by adding entries for undated levels (cc=10), gaps of exactly known length (cc=11), normally distributed gap lengths (cc=12), or gamma distributed gap lengths (cc=13). The age estimates are obtained through a t-walk MCMC run (Christen and Fox 2010). In this process, initial ball-park point estimates for the ages of each dated depth are given, checked for their chronological ordering (and for the sizes of any gaps) and then modified through many iterations. For each iteration, a random dated depth is chosen and its age changed by just a little nudge, a check is performed to ensure that all age estimates remain in chronological order (and that gap sizes remain obeyed), and the 'energy' or likelihood of the age estimates is calculated (iterations where all ages fit well within the calibrated distributions receive a higher energy; see `l.calib`). Then this iteration with the updated group of age estimates is either accepted or rejected. The acceptance probability depends on the iteration's energy; if its energy is higher than that of the previous iteration it is accepted, but if it is lower, it is accepted with a probability proportional to its relative energy. Therefore, over many iterations the process will 'learn' from the data and find high-energy combinations of parameter values that fit with the prior constraints that the ages should be ordered chronologically. Because the iterations are based on a process of modifying values of one parameter each iteration, and because some iterations will not be accepted, the MCMC output will often have a large degree of dependence between neighbouring iterations. Therefore, some thinning will have to be done, by storing only one every few iterations (default 20). Also, since the initial ball-park estimates could be quite wrong, the first 100 or so iterations should also be discarded (burnin). It is thus important to check the time-series of the energy after the run. We don't want to see a remaining burn-in at the start, and we don't want to see a noticeable 'structure' where iterations remain in approximately or entirely the same spot for a long time. Instead, an ideal run will look like white noise. By default, the model output and the settings are stored in a list called 'info' which is placed into R's session as a list. For example, to retrieve the model output, type `info$output`. This has a column for each date, and a row for each stored MCMC iteration. The MCMC's energy can be found in `info$Us`. The model's 'structure' such as blocks or gaps can be found in `info$struc`. To check which dates were used, type `info$dat` or `info$dat` (the latter will include all information, including any gaps).

Value

a variable 'info' which contains the dating and modelling information to produce a plot (see details). Also calls the function `draw.strat` to produce a plot of the results.

References

- Bronk Ramsey C, 1995. Radiocarbon calibration and analysis of stratigraphy: The OxCal program. *Radiocarbon* 37, 425 – 430.
- Buck CE, Kenworthy JB, Litton CD, Smith AFM, 1991. Combining archaeological and radiocarbon information: a Bayesian approach to calibration. *Antiquity* 65, 808-821.
- Buck et al. 1999. BCal: an on-line Bayesian radiocarbon calibration tool. *Internet Archaeology* 7.
- Christen JA, Fox C 2010. A general purpose sampling algorithm for continuous distributions (the t-walk). *Bayesian Analysis* 5, 263-282.

Nicholls G, Jones M 2001. Radiocarbon dating with temporal order constraints. Journal of the Royal Statistical Society: Series C (Applied Statistics) 50, 503-521.

Examples

```
## Not run:
tmp <- tempdir()
strat(, strat.dir=tmp, its=1000, thinning=1, internal.thinning=1)

## End(Not run)
```

strat.cleanup	<i>Remove .out files from a run</i>
---------------	-------------------------------------

Description

Remove all .out files in the folder of a strat file.

Usage

```
strat.cleanup(set = get("info"))
```

Arguments

set	Detailed information of the current run, stored within this session's memory as variable info.
-----	--

Value

NA

thinner	<i>Thin iterations.</i>
---------	-------------------------

Description

Randomly thin iterations by a given proportion, for example if autocorrelation is visible within the MCMC series.

Usage

```
thinner(proportion = 0.1, set = get("info"), write = TRUE, save.info = TRUE)
```

Arguments

proportion	Proportion of iterations to remove. Should be between 0 and 1. Default proportion=0.1.
set	Detailed information of the current run, stored within this session's memory as variable info.
write	Whether or not to write the changes to the output file. Defaults to TRUE.
save.info	Whether or not to store a variable 'info' in the session which contains the run input, output and settings. Defaults to save.info=TRUE.

Details

From all iterations, a proportion is removed with to-be-removed iterations sampled randomly among all iterations.

Value

NA

withinhpd

Calculate the fit of a modelled age with the corresponding date

Description

This is a measure of the fit of the modelled age to that of the date. If many of the modelled age iterations fall within any of the highest posterior density (hpd) range of a date, the model fits the date well. The values can range from 0

Usage

```
withinhpd(calibs, probs, modelled, prob = 0.95)
```

Arguments

calibs	The calibrated ages of the distributions (dates)
probs	The probabilities associated with the calibrated ages (dates)
modelled	The modelled ages
prob	Probability range for the hpd ranges

Value

a list of fits (values between 0 and 100)

Index

ages.undated, [2](#)

draw.rings, [3](#)

draw.strat, [6](#)

IAT, [9](#)

rings, [10](#)

scissors, [12](#)

sim.rings, [13](#)

sim.strat, [15](#)

strat, [16](#)

strat.cleanup, [21](#)

thinner, [21](#)

withinhpd, [22](#)