Package 'EDE'

October 31, 2026

Title Extinction Date Estimation from Sighting Records

Version 0.3.0

Maintainer Rodrigo Fonseca Villa <rodrigo03.villa@gmail.com>

Description Infers the extinction of a species from a time-ordered record of sightings, collecting methods from the sighting-record and stratigraphic-range literature under one interface. Includes endpoint estimators (Robson & Whitlock, 1964; Strauss & Sadler, 1989; Roberts & Solow, 2003), frequentist persistence tests (Solow, 1993a, 1993b; Burgman, Grimson & Ferson, 1995; Solow & Roberts, 2003; Solow, 2005; McInerny, Roberts, Davy & Cribb, 2006; Jarić & Ebenhard, 2010), Bayesian models (Rout, Salomon & McCarthy, 2009; Caley & Barry, 2014), methods for sightings of uncertain reliability (Thompson et al., 2013; Jarić & Roberts, 2014; Lee et al., 2014; Brook, Buettel & Jarić, 2019; Kodikara et al., 2020, 2021), models combining records and surveys (Thompson et al., 2017, 2019), and confidence intervals on the end-points of fossil ranges (Marshall, 1994, 1997; Weiss & Marshall, 1999). Every method is tested against the published examples of its source paper and returns a common result class; persistence tests are inverted to give one-sided upper confidence bounds.

Language en-US

License GPL (>= 3)

URL https://github.com/rodrigosqrt3/EDE

BugReports https://github.com/rodrigosqrt3/EDE/issues

Encoding UTF-8

LazyData true

RoxygenNote 7.3.3

Depends R (>= 4.1.0)

Imports stats

Suggests testthat (>= 3.0.0), knitr, rmarkdown, ggplot2

VignetteBuilder knitr

Config/testthat/edition 3

NeedsCompilation no

Author Rodrigo Fonseca Villa [aut, cre] (ORCID:
<https://orcid.org/0009-0005-2938-2270>)

Repository CRAN

Date/Publication 2026-10-31 08:20:55 UTC

Contents

brook2019 . . . 3
burgman1995 . . . 5
caley_barry2014 . . . 7
dodo . . . 9
ede_compare . . . 10
ede_estimate . . . 11
extinct_birds . . . 11
ferret . . . 12
jaric2010 . . . 13
jaric_roberts2014 . . . 14
kodikara2020 . . . 15
kodikara2021 . . . 17
lee2014 . . . 20
marshall1994 . . . 22
marshall1994_trend . . . 23
marshall1997 . . . 24
mcinerny2006 . . . 25
monk_seal . . . 26
ole . . . 27
orchid . . . 28
robson1964 . . . 28
rout2009 . . . 29
rout2009_absences . . . 30
sighting_data . . . 31
solow1993 . . . 32
solow1993b . . . 33
solow1993_bayes . . . 34
solow2005 . . . 35
solow_roberts2003 . . . 36
strauss1989 . . . 37
strauss1989_bayes . . . 37
strauss1989_curve . . . 39
thompson2013 . . . 39
thompson2017 . . . 41
thompson2019 . . . 43
weiss_marshall1999 . . . 45
weiss_marshall1999_prior . . . 47

Index **48**

brook2019 *Reliability-weighted resampling of a sighting record (Brook et al., 2019)*

Description

Applies any extinction-date estimator to many random sub-records drawn from a sighting record of mixed reliability, and summarises the resulting empirical distribution (the "BBJ" weighted-resampling approach of Brook, Buettel & Jaric, 2019).

Usage

```
brook2019(
  sd,
  reliability,
  method = strauss1989,
  n_resamples = 1000,
  statistic = c("estimate", "upper"),
  alpha = 0.05,
  persist_year = NULL,
  min_sightings = 3,
  seed = NULL,
  ...
)
```

Arguments

sd	A sighting_data object.
reliability	Sighting reliabilities in [0, 1]. Either a numeric vector with one value per row of sd or per positive-count row (taken as the combined reliability of that time unit), or a list with one numeric vector per positive-count row giving the reliabilities of the individual sightings in that time unit (its length must equal the count); these are combined as in step 1. Values for zero-count rows are ignored.
method	An estimator function whose first argument is a sighting_data object, such as ole(), strauss1989(), robson1964() or solow1993(). It may return an ede_estimate object or a single number. The "S93" estimator of the authors' code, t[n] + (t[n] - t[1]) / (n - 1) with the first sighting as origin, is the point estimate of strauss1989() (the default).
n_resamples	Number of random sub-records (the paper suggests about 1,000 or more; the authors' code uses 10,000).
statistic	Component of the ede_estimate returned by method to collect: "estimate" (default) or "upper". Ignored if method returns a number.
alpha	Level for the upper quantile, in (0, 1).
persist_year	Optional time at which to report the proportion of resampled values later than it.
min_sightings	Minimum number of kept records for a sub-record to be analysed.

seed	Optional random seed. The global random number state is restored on exit.
...	Further arguments passed to method (e.g. test_year).

Details

The algorithm follows the paper and the authors' R code (Data S1, Zenodo doi:10.5281/zenodo.3247189, bbj.2018()):

1. Sightings in the same time unit are combined into one record with reliability p[j] = 1 - prod(1 - p[ij]) (the unnumbered equation in the Methods of Brook et al., 2019).
2. Each record is kept independently with probability equal to its reliability (resampling "without replacement", using reliability as an inclusion probability). Records with reliability 1 are therefore always kept; there is no other rule that forces the most recent certain sighting to be retained.
3. method is applied to each sub-record and the chosen statistic is stored. Sub-records with fewer than min_sightings records (the authors' code requires at least 3), and sub-records for which method fails or returns NA, are discarded; their number and proportion are reported.
4. The kept values form the empirical distribution of the statistic.

Summaries: estimate is the median of the distribution, the point estimate named in the paper's Methods and in its Fig. 1 (the authors' code instead reports the mean, which is returned as mean). upper is the 1 - alpha quantile of the distribution, the "upper 95% confidence bound" of the paper (its text calls it the "2.5th percentile", apparently in error; the authors' code takes the last year at which the proportion of values not yet reached exceeds alpha, i.e. the 1 - alpha quantile). The extinction years in Table 1(b) of the paper are these upper bounds. If persist_year is given, persistence is the proportion of values later than persist_year, the "probability of persistence" in Table 1. Quantiles use type = 1 (inverse empirical distribution), so infinite values (e.g. a non-existent upper bound) are handled.

Kept records enter the sub-record with count 1. Zero-count rows of sd (for example a leading row that fixes the observation origin) are kept in every sub-record; otherwise the earliest kept sighting becomes the origin for estimators that use one, as in the authors' code.

Status of the method: Solow (2023) criticised the statistical properties of the extinction date and interval produced this way, noting in particular that the estimated extinction time exceeds the true one when reliabilities are high. Jaric et al. (2023) replied that the procedure is not itself an extinction-date estimator but a pre-processing step for the input data, whose result is passed to any estimator (McCrea et al., 2024). The summaries returned here are descriptive statistics of that resampling distribution, not confidence intervals with established coverage.

Value

An ede_estimate object (interval_type = "one-sided") with estimate (median), upper (1 - alpha quantile) and additional components mean, values (the kept resampled values), n_resamples, n_failed, prop_failed, statistic, combined_reliability and, if persist_year is given, persistence and persist_year.

References

Brook, B. W., Buettel, J. C., & Jaric, I. (2019). A fast re-sampling method for using reliability ratings of sightings with extinction-date estimators. Ecology, 100(9), e02787.

Jaric, I., Buettel, J. C., & Brook, B. W. (2023). A fast re-sampling method for using reliability ratings of sightings with extinction-date estimators: reply. Ecology, 104(9), e4124.

Solow, A. R. (2023). A fast re-sampling method for using reliability ratings of sightings with extinction-date estimators: comment. Ecology, 104(9), e4123.

McCrea, R. S., Cheale, T., Campillo-Funollet, E., & Roberts, D. L. (2024). Inferring species extinction from sighting data. Cambridge Prisms: Extinction, 2, e19.

`burgman1995` *Burgman, Grimson & Ferson (1995) tests for frequency sighting data*

Description

Exact tests for sighting records divided into `CT` equal periods, where `N` is the total number of sightings (all records counted, not only occupied periods).

Usage

```
burgman1995(sd, alpha = 0.05, test_year, data_out = FALSE, equation = 2)
```

Arguments

`sd`	A sighting_data object.
`alpha`	Significance level, in (0, 1). The returned upper bound is the earliest candidate time at which persistence is rejected at this level.
`test_year`	Latest time to test. Must be later than the last sighting. A bound beyond `test_year` is reported as `NA` with a warning.
`data_out`	If `TRUE`, return the p-value curve at integer steps after the last sighting instead of the upper bound.
`equation`	Which equation of the source paper to use: `2` for extinction (the default), `3` (binary runs test) or `4` (frequency runs test) for a change in conservation status.

Details

- `equation = 2` (the default) tests extinction. The p-value `(Ce / CT)^N`, where `Ce` is the period of the last sighting, is the probability that the species is not extinct; it is the empty-cells test restricted to the run of empty periods at the end of the record.
- `equation = 3` is the runs test of Bradley (1968) on the binary (occupied / empty) sequence of periods, so counts are ignored. With `n1` occupied periods out of `CT` and `r` the longest run of empty periods anywhere in the record, the probability of an empty run at least `r` long is `sum_{k=1}^{[n0/r]} (-1)^(k+1) C(n1+1, k) C(CT - r k, n1) / C(CT, n1)`, with `n0 = CT - n1` empty periods (inclusion-exclusion over the `n1 + 1` gaps that hold the empty periods).

The equation is illegible in the source, so this form was reconstructed from its definitions and checked against the paper's numbers: it gives the Fig. 1b value 0.02885 and the Table 2 values for the black-footed ferret (0.003) and both stick-nest rats (0.139, 0.350). It does not give the Table 2 value for the monk seal (0.244 here against 0.258 in the paper, which equals the paper's equation 4 value; see Details). Like equation 4, it addresses a change in status rather than extinction.

- `equation = 4` tests for a change in conservation status (Grimson et al., 1992). It computes, by inclusion-exclusion over the occupancy of periods, the probability of a run of empty periods at least as long as the longest observed run anywhere in the record. Burgman et al. (1995, p. 927) note that it addresses a different question from extinction, and that a significant result may be discounted if the longest run is not near the end of the record.

As in the Caribbean monk seal example of the source paper (Fig. 1a and Tables 1-2), the earliest supplied time defines the observation origin and is not one of the `CT` periods; sightings recorded in that period are not counted. To include an earlier start date, prepend a zero-count row. The paper's black-footed ferret values (observation from January 1972) correspond to a zero-count row for 1971, so that 1972 is a period. Burgman et al. (2000, Table 2.1) used the same convention for Western Australian *Acacia*: the year of the first collection is the first period and 1993 the last; equations 2 and 4 reproduce their published values.

For the monk seal, the origin convention above reproduces the paper's equation 2 (0.053) and equation 4 (0.258) values (`CT = 77`, `n1 = N = 4`, `r = 40`), but equation 3 gives 0.244. No choice of `CT` and `r` near these values gives 0.258 exactly (`CT = 78` with `n1 = 4`, i.e. counting 1915 as an empty period, gives 0.2587). Since all seal counts are 1, the paper's equal entries for equations 3 and 4 suggest that the equation 4 value was reported for both.

The paper defines only the tests. Since the periods are discrete, the returned upper bound is the first period after the last sighting at which the null hypothesis is rejected. Under `equation = 3` or `4` the p-value curve need not be monotone, so this bound should be read with the caveat above.

Value

An ede_estimate object with `estimate = NA` and the one-sided upper bound in `upper`, or (if `data_out = TRUE`) a data frame with columns `time` and `chance` (p-values).

References

Burgman, M. A., Grimson, R. C., & Ferson, S. (1995). Inferring threat from scientific collections. Conservation Biology, 9(4), 923-928.

Burgman, M. A., Maslin, B. R., Andrewartha, D., Keatley, M. R., Boek, C., & McCarthy, M. (2000). Inferring threat from scientific collections: power tests and an application to Western Australian Acacia species. In S. Ferson & M. Burgman (Eds.), Quantitative Methods for Conservation Biology (pp. 7-26). Springer, New York.

Bradley, J. V. (1968). Distribution-free statistical tests. Prentice-Hall, Englewood Cliffs, New Jersey.

caley_barry2014 *Caley & Barry (2014) Bayesian extinction probability from a detection record*

Description

Posterior probability that a population is extant at the end of `test_year`, from a record of detections (`y[t] = 1`) and non-detections (`y[t] = 0`) in each period `t = 1, ..., T` (Caley & Barry, 2014). The population is extant at the start of period 1. The latent variable `z` is the period in which it goes extinct (`z > T` if it persists beyond the record). If it goes extinct in period `z`, it cannot be detected from period `z` on.

Usage

```
caley_barry2014(
  sd,
  test_year,
  model = c("constant", "nonconstant"),
  formulation = c("published", "equations"),
  prior_phi = c(1.1, 1.1),
  prior_lambda = c(1.1, 1.1),
  prior_growth = c(-2.3, 0.69),
  prior_delta = c(0.01, 4.6),
  prior_eps0 = c(-20, 20),
  prior_eps1 = c(0, 20),
  n_iter = 1e+05,
  burn_in = n_iter/10,
  data_out = FALSE
)
```

Arguments

`sd`	A sighting_data object with whole-number times.
`test_year`	Last period of observation. Must be later than the last sighting. The posterior refers to persistence beyond the end of this period, given no detections after the last sighting.
`model`	`"constant"` (exact computation) or `"nonconstant"` (MCMC).
`formulation`	For the constant model: `"published"` (Code S1, which reproduces the paper) or `"equations"` (likelihood as printed).
`prior_phi`, `prior_lambda`	Beta prior parameters `c(a, b)` for the constant model.
`prior_growth`, `prior_delta`, `prior_eps0`, `prior_eps1`	Uniform prior bounds for the non-constant model.
`n_iter`, `burn_in`	Total MCMC iterations and initial iterations discarded (non-constant model). Use `set.seed()` for reproducibility.

data_out — If TRUE (constant model only), return the posterior probability of persistence at the end of each period from the one after the last sighting to test_year, assuming no further detections.

Details

Constant model (model = "constant", section "Model specification - constant population"). Each period, the population goes extinct with probability phi and, if extant, is detected with probability lambda. The priors are phi ~ Beta(a_phi, b_phi) and lambda ~ Beta(a_lambda, b_lambda). The paper's likelihood is [y | z] = prod_{t < z} lambda^y[t] (1 - lambda)^(1 - y[t]), and [z] = (1 - phi)^(z - 1) phi, or (1 - phi)^T when z > T. The paper samples the posterior with a Gibbs sampler (Code S1). Here it is computed exactly instead: given z, the integrals over phi and lambda are Beta functions, so the posterior of z is a finite set of weights, and the posteriors of phi and lambda are mixtures of Beta distributions.

The published numbers come from Code S1. Its lambda update, Beta(a_lambda + S, b_lambda + z - S), counts the extinction period itself as a non-detection, one period more than the likelihood above. The z update has no such term. The sampler is still a valid Gibbs sampler: the extra factor (1 - lambda) is the same for every z, so it targets the model above with the prior lambda ~ Beta(a_lambda, b_lambda + 1). formulation = "published" (the default) computes that posterior and reproduces the paper's results. formulation = "equations" uses the likelihood as printed in the paper.

Non-constant model (model = "nonconstant", Code S2). The population size is N[t] = exp(r (t - 1)), with N[1] = 1 fixed because N[1] and delta are not separately identifiable. Detection is lambda[t] = 1 - exp(-delta N[t]) (equation 3) and extinction is logit(phi[t]) = eps0 - eps1 log(N[t]) (equation 4). The priors are uniform on prior_growth, prior_delta, prior_eps0 and prior_eps1. Each iteration of the Metropolis-within-Gibbs sampler of Code S2 is ported as published:

1. impute z from its conditional distribution;
2. take a log-normal random-walk step for delta (sd 0.5);
3. take a joint step for eps0 (normal, sd 1.5) and eps1 (log-normal, sd 0.15);
4. take a normal random-walk step for r (sd 0.05).

The chain starts from r = 0, delta = log(2), eps0 = 0 and eps1 = 1. As in Code S1, the lambda likelihood used in the delta and r updates includes the extinction period as a non-detection, while the z update does not. Because lambda[t] now varies with t, this factor differs between values of z. **The full conditionals of Code S2 are therefore not those of a single joint model**, and the sampler is reproduced only for comparability with the paper. Code S2 uses r ~ U(-2.3, 0.69), whereas the paper's text gives U(-1.6, 0.69); the code's value is the default here. Code S2 also does not report an implied prior probability of persistence that can be checked, so prior and bayes_factor are NA for this model.

Input. Periods are the integer time units from the earliest time in sd to test_year. A period is a detection if its count is positive; periods not listed in sd are non-detections. Unlike in solow1993(), the first period is an observation period even if it has a sighting. To start observation before the first sighting, add a zero-count row (e.g. the 1999 start of the paper's fox example).

Value

A list of class ede_bayes with components:

- `posterior`: posterior probability that the population is extant at the end of `test_year`.
- `prior`: implied prior probability of that event, `B(a_phi, b_phi + T) / B(a_phi, b_phi)`, for the constant model; `NA` for the non-constant model.
- `bayes_factor`: posterior odds over prior odds of persistence; `NA` for the non-constant model.
- `test_year` and `method`.
- `extinction_time`: a data frame with the posterior probability that extinction occurred in each period after the last sighting.
- Posterior medians and 95% equal-tailed intervals: `phi` and `lambda` (constant model), or `growth`, `delta`, `eps0` and `eps1` (non-constant model).
- For the non-constant model, also the chain after burn-in (`samples`) and the acceptance rates (`acceptance`).

If `data_out = TRUE`, a data frame with columns `time`, `bayes_factor` and `posterior` is returned instead.

References

Caley, P., & Barry, S. C. (2014). Quantifying extinction probabilities from sighting records: inference and uncertainties. PLoS ONE, 9(4), e95857.

Examples

```
# Tasmanian red fox carcasses (Caley & Barry, 2014, Table 1), 2001 start
fox <- sighting_data(data.frame(year = c(2001, 2003, 2005, 2006),
                                count = 1))
caley_barry2014(fox, test_year = 2012)
# 1999 start
fox99 <- sighting_data(data.frame(year = c(1999, 2001, 2003, 2005, 2006),
                                  count = c(0, 1, 1, 1, 1)))
caley_barry2014(fox99, test_year = 2012)$posterior
```

---

dodo *Sighting record of the dodo*

Description

Reliable sightings of the dodo (*Raphus cucullatus*) in Mauritius, used by Roberts & Solow (2003) to estimate its extinction date with `ole()`. Two sightings were made in 1628. A later report from 1674 was excluded by the authors as unreliable.

Usage

```
dodo
```

Format

A data frame with 9 rows: `year` and `count` (number of sightings).

Source

Roberts, D. L., & Solow, A. R. (2003). Flightless birds: When did the dodo become extinct? Nature, 426(6964), 245.

Examples

```
ole(sighting_data(dodo))
```

`ede_compare` *Apply every applicable method to one sighting record*

Description

Runs the estimators and persistence tests that need only a sighting record and returns their results side by side. Methods that need additional input (such as sighting reliabilities) are not included. A method that cannot be applied to the record, or that warns, is kept in the table with the message in `note`.

Usage

```
ede_compare(sd, test_year, alpha = 0.05)
```

Arguments

`sd`	A sighting_data object.
`test_year`	Latest time to test for the persistence tests. Must be later than the last sighting.
`alpha`	Significance level, in (0, 1).

Details

The methods rest on different models, so their results answer different questions and should not be averaged; see the vignette for guidance on choosing a method.

Value

A data frame with one row per method and columns `method`, `function_name`, `estimate`, `lower`, `upper`, `interval_type` and `note`. Persistence tests report the one-sided upper bound in `upper`.

Examples

```
ede_compare(sighting_data(monk_seal), test_year = 2050)
```

ede_estimate *Extinction date estimate*

Description

Common S3 result class returned by every estimator in this package.

Format

A list with components:

estimate Point estimate, or NA if not defined for this method.

lower, upper Interval bounds, or NA if not defined. An upper bound that does not exist is returned as Inf.

interval_type Either "two-sided" or "one-sided".

method Character string identifying the estimator.

alpha Significance level used to compute the estimate.

Bayesian estimators (strauss1989_bayes(), weiss_marshall1999()) return posterior probability (credible) intervals and set interval_kind to "credible". Some estimators append method-specific components, documented on their help pages.

Details

Persistence tests (solow1993(), solow2005(), mcinerny2006(), burgman1995(), solow_roberts2003() and jaric2010()) do not define a point estimate. Inverting the test gives a one-sided 1 - alpha upper confidence bound for the extinction time, which is returned in upper with estimate = NA. solow1993b() also returns the minimum variance unbiased estimate of its model in estimate.

extinct_birds *Sighting records of four possibly extinct birds, by evidence quality*

Description

Sighting years of the Eskimo Curlew, the Ivory-billed Woodpecker (U.S.), the Nukupu'u and the O'ahu 'Alauahio, classified by Roberts et al. (2010) into physical evidence, independent expert opinion and controversial sightings, as tabulated by Jaric & Roberts (2014). Jaric & Roberts assign reliabilities of 0.8-0.9, 0.6-0.8 and 0.1-0.4 to the three classes. A year can appear under more than one class.

Usage

```
extinct_birds
```

Format

A data frame with columns `species`, `scientific_name`, `year` and `evidence` (`"physical"`, `"expert"` or `"controversial"`).

Source

Jaric, I., & Roberts, D. L. (2014). Accounting for observation reliability when inferring extinction based on sighting records. Biodiversity and Conservation, 23(11), 2801-2815, Table 2.

Roberts, D. L., Elphick, C. S., & Reed, J. M. (2010). Identifying anomalous reports of putatively extinct species and why it matters. Conservation Biology, 24(1), 189-196.

Examples

```
curlew <- extinct_birds[extinct_birds$species == "Eskimo Curlew", ]
# keep the most reliable class for each year
curlew <- curlew[order(curlew$year, match(curlew$evidence,
  c("physical", "expert", "controversial"))), ]
curlew <- curlew[!duplicated(curlew$year), ]
reliability <- c(physical = 0.85, expert = 0.7, controversial = 0.25)
jaric_roberts2014(
  sighting_data(data.frame(year = curlew$year, count = 1)),
  reliability = reliability[curlew$evidence]
)
```

---

ferret *Sighting record of the black-footed ferret*

Description

Monthly sightings of the black-footed ferret (*Mustela nigripes*) in South Dakota from 1972 to 1984, one row per month with a sighting, used by Solow (1993b) and Jaric & Ebenhard (2010).

Usage

```
ferret
```

Format

A data frame with 28 rows: `year`, `month` (1-12) and `count`.

Details

The two papers index months differently. Solow (1993b) counts months from January 1972 (= 1), starts observation at 0 and tests to `T = 229`; Jaric & Ebenhard (2010) count from June 1972 (= 1) and test to December 1990 (`T = 223`).

Source

Solow, A. R. (1993b). Inferring extinction in a declining population. Journal of Mathematical Biology, 32(1), 79-82.

Jaric, I., & Ebenhard, T. (2010). A method for inferring extinction based on sighting records that change in frequency over time. Wildlife Biology, 16(3), 267-275.

Examples

```
# Solow (1993b): months from January 1972, observation starting at 0
months <- (ferret$year - 1972) * 12 + ferret$month
sd <- sighting_data(data.frame(time = c(0, months), count = c(0, ferret$count)))
tail(solow1993b(sd, test_year = 229, data_out = TRUE), 1)
```

---

jaric2010 *Jaric & Ebenhard (2010) sighting-trend index*

Description

Index of persistence from the average interval `I` between distinct sighting times, optionally adjusted by the trend coefficient `c`. For candidate time `T`, the index is `I / (I + T - t[n])` (equations 3-4 of Jaric & Ebenhard, 2010) or, with `trend = TRUE`, `(I + c) / (I + c + T - t[n])` (equations 5-6). The paper calls this the probability that the species is still present; it is an index rather than a p-value from a fitted probability model.

Usage

```
jaric2010(sd, alpha = 0.05, test_year, data_out = FALSE, trend = TRUE)
```

Arguments

`sd`	A sighting_data object.
`alpha`	Significance level, in (0, 1). The returned upper bound is the earliest candidate time at which persistence is rejected at this level.
`test_year`	Latest time to test. Must be later than the last sighting. A bound beyond `test_year` is reported as `NA` with a warning.
`data_out`	If `TRUE`, return the p-value curve at integer steps after the last sighting instead of the upper bound.
`trend`	Logical. If `TRUE` (the default), adjust the average interval by the trend coefficient (equations 5-6 and 8).

Details

Several sightings in one time unit count as one sighting, as in the paper. `I` is the mean interval between the first and last sighting times, so the observation origin is not used. `c` is the mean change in consecutive interval lengths, which reduces to the difference between the last and first intervals divided by `n - 2`; it is positive when sightings become less frequent.

The returned upper bound is `t[n] + (I + c) (1 - alpha) / alpha`, the upper bound of the confidence interval in equations 7-8 of the paper (`c = 0` when `trend = FALSE`).

Value

An ede_estimate object with `estimate = NA`, the one-sided upper bound in `upper`, and additional components `average_interval` and `trend_coefficient`; or (if `data_out = TRUE`) a data frame with columns `time` and `chance` (index values) carrying the same two quantities as attributes.

References

Jaric, I., & Ebenhard, T. (2010). A method for inferring extinction based on sighting records that change in frequency over time. Wildlife Biology, 16(3), 267-275.

`jaric_roberts2014` *Reliability-adjusted extinction inference*

Description

Modifies the constant-rate Solow (1993a) model by assigning a reliability `p[i]` in `[0, 1]` to each sighting. The number of sightings is replaced by the effective number `r = sum(p[i])` (equation 4 of Jaric & Roberts, 2014) and the last sighting time by its expectation `t[r]` over which record is the last valid one (equations 5-6). The point estimate is `(r + 1) / r * t[r]` (equation 9) and the upper confidence bound is `t[r] / alpha^(1 / r)` (equation 10). For a candidate time `T`, the modified Solow p-value is `(t[r] / T)^r` (equations 7-8), with times measured from the observation origin.

Usage

```
jaric_roberts2014(
  sd,
  reliability,
  alpha = 0.05,
  test_year = NULL,
  data_out = FALSE
)
```

Arguments

`sd`	A sighting_data object. Counts must be zero or one. This is a restriction of the package, not of the paper: reliability belongs to an individual record, and the paper does not say how to combine several records of different reliability in one time unit.
`reliability`	Numeric reliabilities in `[0, 1]`, supplied either for every row of `sd` or for the positive-count rows only.
`alpha`	Significance level, in (0, 1), for the upper confidence bound.
`test_year`	Optional candidate time at which to compute the p-value of equations 7-8. Must be later than the last sighting. Required when `data_out = TRUE`.
`data_out`	If `TRUE`, return the p-value curve at integer steps after the last sighting up to `test_year` instead of the estimate.

Details

The earliest supplied time defines the observation origin. A sighting at that time fixes the origin and does not enter `r` or `t[r]`, so its reliability is not used; prepend a zero-count row to include it. Reliabilities given for zero-count rows are ignored.

Value

An ede_estimate object with additional components `effective_sightings` (`r`), `effective_endpoint` (`t[r]`, on the original time scale) and, if `test_year` is given, `p_value`; or (if `data_out = TRUE`) a data frame with columns `time` and `chance` (p-values).

References

Jaric, I., & Roberts, D. L. (2014). Accounting for observation reliability when inferring extinction based on sighting records. Biodiversity and Conservation, 23(11), 2801-2815.

`kodikara2020` *Kodikara et al. (2020) Bayesian extinction year from yearly sighting records*

Description

Posterior distribution of the extinction year from a record of yearly certain and (optionally) uncertain sightings, following Kodikara, Demirhan, Wang, Solow & Stone (2020). Each year `t = 1, ..., T` of the observation period is a trial. `tau_E` is the first year following extinction: the species is extant in years `1, ..., tau_E - 1`, so `tau_E` is at least one year after the last certain sighting `c_n`.

Usage

```
kodikara2020(certain, uncertain = NULL, test_year, model = NULL, level = 0.95)
```

Arguments

`certain`	A sighting_data object with whole-number times (years) and the certain sightings. Its earliest time is the first year of observation. At least one certain sighting is required.
`uncertain`	A sighting_data object with the uncertain sightings, or `NULL`. Ignored by Model 1.
`test_year`	Last year `T` of the observation period. Must be later than the last certain sighting and not earlier than the last uncertain sighting.
`model`	`1`, `2` or `3` (see Description). Defaults to `1` without uncertain sightings and `2` otherwise.
`level`	Probability content of the HDIs.

Details

- **Model 1** (certain sightings only, section 2.1). An extant species is sighted (with certainty) in a year with probability pc. With nc certain sighting years, the likelihood is equation 3: `0` for `tau_E <= c_n`, `pc^nc (1 - pc)^(tau_E - 1 - nc)` for `c_n < tau_E <= T`, and `pc^nc (1 - pc)^(T - nc)` for `tau_E > T`.
- **Model 2** (certain and uncertain sightings, section 2.2). Before extinction, valid and invalid uncertain sightings occur independently with probabilities puv and pui, so an uncertain sighting occurs with probability `pu = puv (1 - pui) + pui (1 - puv) + puv pui`, and is recorded only in years without a certain sighting (the "censoring" process). After extinction, only invalid uncertain sightings occur, with probability pui. With `nu(tau_E)` uncertain sightings before `tau_E` and `Nu` in total, the likelihood is equation 9: `pc^nc ((1 - pc) pu)^nu(tau_E) ((1 - pc)(1 - pu))^(tau_E pui^(Nu - nu(tau_E)) (1 - pui)^(T - (tau_E - 1) - (Nu - nu(tau_E)))` for `c_n < tau_E <= T`, and `pc^nc ((1 - pc) pu)^Nu ((1 - pc)(1 - pu))^(T - nc - Nu)` for `tau_E > T`.
- **Model 3** (section 4.2, Table 4): Model 2 with `puv = pc`.

The prior of `tau_E` is geometric with parameter `theta` (equation 4), `p(tau_E | theta) = (1 - theta)^(tau_E - 1) theta`, and `theta`, pc, puv and pui are standard uniform (equations 5-6); the joint posterior is equation 7.

The paper samples this posterior with JAGS. Because `tau_E` is discrete and all priors are uniform, the posterior of `tau_E` is computed here exactly instead (a package construct): integrating `theta` gives the marginal prior `p(tau_E) = 1 / (tau_E (tau_E + 1))`, whose tail beyond `T` has mass `1 / (T + 1)`; the integrals over pc, puv and pui are finite sums of Beta functions, obtained by writing `1 - (1 - puv)(1 - pui) = puv + (1 - puv) pui` and expanding its power binomially (all terms positive). Since the likelihood is constant for `tau_E > T`, the posterior of `tau_E` beyond `T` is proportional to `1 / (tau_E (tau_E + 1))`, so its quantiles there are analytic. The posterior median and highest posterior density interval (HDI, the shortest run of consecutive years with posterior probability at least `level`) of `tau_E` are reported as years, as in Tables 2-4 of the paper.

Input. Periods are the integer years from the earliest time in `certain` to `test_year` (`T`); include a zero-count row in `certain` when observation began before the first certain sighting (as for the O'ahu 'Alauahio record, which starts in 1837). A year is a sighting year if its count is positive (several sightings in a year count once); unlisted years have no sighting. Uncertain sightings in a year that also has a certain sighting are ignored, as the censoring process of Model 2 assumes they cannot be recorded.

Value

A list of class `ede_bayes` with components:

- `posterior`: posterior probability `p(tau_E > T | S)` that the species was extant in every year of the record (one minus the paper's `p(tau_E <= T | S)`).
- `prior`: the implied prior of that event, `1 / (T + 1)`.
- `bayes_factor`: posterior odds over prior odds of persistence.
- `test_year`, `method` and `model`.
- `extinction_year`: posterior median and `level` HDI of `tau_E`, as years (`c(median, lower, upper)`); `level`.
- `extinction_time`: data frame with the posterior probability of each value of `tau_E` (`time`, as a year) from the year after the last certain sighting to `test_year`.
- `p_certain`: posterior median and `level` HDI of pc (Models 1 and 2; `NA` for Model 3).

References

Kodikara, S., Demirhan, H., Wang, Y., Solow, A. & Stone, L. (2020). Inferring extinction year using a Bayesian approach. Methods in Ecology and Evolution, 11(8), 964-973. doi:10.1111/2041-210X.13408

Examples

```
# Ivory-billed Woodpecker, physical-evidence sightings (Kodikara et al.,
# 2020, Table 2)
ibw_c <- sighting_data(data.frame(
  year = c(1898:1902, 1904:1910, 1913, 1914, 1917, 1924, 1925, 1932,
           1935, 1938, 1939),
  count = 1))
fit <- kodikara2020(ibw_c, test_year = 2010)
fit$extinction_year
```

`kodikara2021` *Kodikara et al. (2021) extinction time with non-homogeneous Poisson sightings*

Description

Posterior distribution of the extinction time `tau_E` when certain and uncertain sightings follow independent non-homogeneous Poisson processes with Weibull (power-law) rates, and `tau_E` is a change-point for the uncertain sightings (Kodikara, Demirhan, Wang & Stone, 2021). Times are measured from the observation origin, which is the earliest time in `certain`. The observation period is `(0, T)`.

Usage

```
kodikara2021(
  certain,
  uncertain = NULL,
  test_year,
  formulation = c("published", "equations"),
  homogeneous = FALSE,
  rate = 0.005,
  sigma_max = 1000,
  alpha_max = 1000,
  level = 0.95,
  n_grid = 200
)
```

Arguments

certain	A sighting_data object with the certain sightings. Its earliest time is the origin; at least one certain sighting must be later than it.
uncertain	A sighting_data object with the uncertain sightings, or NULL (certain sightings only, equation 10). Zero-count rows are ignored.
test_year	End T of the observation period. Must be later than the last certain sighting and not earlier than any uncertain sighting.
formulation	"published" (the authors' code; reproduces the paper) or "equations" (equation 9 as printed).
homogeneous	If TRUE, all shape parameters are fixed at 1. Only available without uncertain sightings: the paper's value for the homogeneous model with uncertain sightings (Figure 7) is not reproduced.
rate	Rate of the exponential prior of tau_E, per time unit.
sigma_max, alpha_max	Upper bounds of the uniform priors of the scale and shape parameters.
level	Probability content of the HDIs.
n_grid	Approximate number of grid points for tau_E in (last certain sighting, T).

Details

- Certain sightings occur before tau_E with rate lambda_c(t) = (alpha_c / sigma_c) (t / sigma_c)^(alpha_c - 1) and mean function m_c(t) = (t / sigma_c)^alpha_c (equations 3 and 5).
- Uncertain sightings occur with rate lambda_u1 (parameters alpha_u1, sigma_u1) before tau_E and lambda_u2 (alpha_u2, sigma_u2) after it (equations 4 and 6).
- The likelihood is equation 9: with Kc certain sightings at times y_ci, and Ku(tau_E) of the Ku uncertain sightings before tau_E, prod_i lambda_c(y_ci) exp(-m_c(tau_E)) x prod_(j <= Ku(tau_E)) lambda_u x prod_(j > Ku(tau_E)) lambda_u2(y_uj) exp(-(m_u2(T) - m_u2(tau_E))); without uncertain sightings only the first factor is used (equation 10). The likelihood is zero for tau_E before the last certain sighting, and is evaluated at tau_E = T for tau_E >= T.
- Priors: Uniform(0, alpha_max) for the shape parameters, Uniform(0, sigma_max) for the scale parameters (both 1000 in the paper) and an exponential prior with rate rate (0.005 in the paper) for tau_E (equation 11). homogeneous = TRUE fixes all shape parameters at 1 (constant rates; section 3.2 and Figure 7).

Formulation. The paper samples the posterior with JAGS. In the authors' code (R_script_Ferret.R, R_script_IBW.R), the likelihood is a product over sightings in which each sighting's factor contains the exponential term of its block. The term exp(-m_c(tau_E)) is therefore raised to the power Kc, exp(-m_u1(tau_E)) to the power Ku(tau_E) and exp(-(m_u2(T) - m_u2(tau_E))) to the power Ku - Ku(tau_E). The code also uses m_u1(tau_E) without truncating tau_E at T. formulation = "published" (the default) uses this likelihood and reproduces the published results. formulation = "equations" uses equation 9 as printed. Its results differ from the published ones (for the Ivory-billed Woodpecker, the posterior median of tau_E is 1946 instead of 1950), so it is not verified against the paper.

Computation (package construct). This function does not use MCMC. Given tau_E, the likelihood factorises into the certain, pre- and post-extinction uncertain blocks. In each block, the integral

over the scale sigma is an upper incomplete gamma function, int_0^sigma_max sigma^(-alpha K) exp(-c sigma^(-alph
The shape alpha is integrated numerically on a log-spaced grid. The posterior density of tau_E is evaluated on a grid between the discontinuities at the uncertain sighting times, and integrated with the trapezoidal rule. The density beyond T is integrated exactly for the equations (constant likelihood), and on a grid for the published formulation. The median and the HDI (the shortest interval with posterior probability level) are as reported in Tables 1 and 2 of the paper. The marginal posteriors of the shape parameters are summarised in the same way.

Input. Times are in the units of certain and uncertain. The priors depend on these units: the paper uses months for the black-footed ferret and years for the Ivory-billed Woodpecker. A count greater than one is that number of sightings at the same time. Every sighting must be later than the origin. Add a zero-count row in certain at the start of observation.

Value

A list of class ede_bayes with components:

- posterior: posterior probability P(tau_E > T) of persistence beyond the observation period.
- prior: the prior probability of that event, exp(-rate (T - origin)).
- bayes_factor: posterior odds over prior odds of persistence.
- test_year, method, formulation and level.
- extinction_time: posterior median and HDI of tau_E, as times (c(median, lower, upper)).
- shape: posterior medians and HDIs of alpha_c, alpha_u1 and alpha_u2 (a matrix with one row for each estimated shape).
- density: data frame with the posterior density (density) and distribution function (cdf) of tau_E on the grid (time). Use it for P(tau_E <= t).

References

Kodikara, S., Demirhan, H., Wang, Y. & Stone, L. (2021). Inferring extinction date of a species using non-homogeneous Poisson processes with a change-point. Methods in Ecology and Evolution, 12(3), 530-538. doi:10.1111/2041-210X.13542

Examples

```
# Black-footed ferret, months since January 1972 (Kodikara et al., 2021,
# Table 1; dates from Solow, 1993)
months <- (ferret$year - 1972) * 12 + ferret$month
sd <- sighting_data(data.frame(month = c(0, months),
                               count = c(0, ferret$count)))
fit <- kodikara2021(sd, test_year = 229, n_grid = 200)
1972 + fit$extinction_time / 12
fit$shape
```

lee2014 *Lee et al. (2014) Bayesian extinction probability with certain and uncertain sightings*

Description

Posterior probability that a species is extant at the end of `test_year`, from a discrete record of certain sightings and, optionally, a second record of uncertain sightings that may be false positives (Lee, McCarthy, Wintle, Bode, Roberts & Burgman, 2014). The model is the OpenBUGS model of the paper's Appendix S1, computed exactly instead of by MCMC.

Usage

```
lee2014(certain, uncertain = NULL, test_year, prior = 0.5, rate_max = 100)
```

Arguments

`certain`	A sighting_data object with the certain sightings, with whole-number times. It may contain only zero-count rows (no certain sightings).
`uncertain`	A sighting_data object with the uncertain sightings, or `NULL` (certain sightings only).
`test_year`	Last period of observation. Must be a whole number later than every sighting.
`prior`	Prior probability that the species is extant at the end of `test_year`, in (0, 1).
`rate_max`	Upper bound `M` of the uniform priors on the sighting rates (Appendix S1: 100).

Details

Model (Appendix S1; equations 5-12). The record is a series of periods `i = 1, ..., T` (e.g. years). In each period at most one sighting of each kind is recorded (equation 1). The extinction time is `te ~ U(0, T)` given extinction (equation 12) and the species is present in period `i` if it is extant at the end of the record or if `i <= te` (`step(te - i)` in Appendix S1). A certain sighting occurs in a period with probability `1 - exp(-mc)` while the species is present and `0` after extinction (equation 5); an uncertain sighting occurs with probability `1 - exp(-(mu + f))` while it is present and `1 - exp(-f)` after extinction (equation 6). The two records are independent given `te` (equations 7-8). The rates have independent uniform priors on `(0, rate_max)` (equation 11; Appendix S1 uses `rate_max = 100` for `mc`, `mu` and `f`), and the prior probability that the species is extant is `prior` (`Extant ~ dbern(0.5)` in Appendix S1, `pi` in the paper). The posterior is equation 2 with the Bayes factor of equation 3.

Exact computation. Given extinction, the likelihood depends on `te` only through the number of periods `k = floor(te)` with the species present, which is uniform on `0, ..., T - 1`. For each `k`, the certain likelihood is `(1/M) int_0^M (1 - e^-m)^S e^(-m (k - S)) dm` (equations 9-10 in discrete form), an incomplete Beta function, and is zero if a certain sighting occurs after period `k`. The uncertain likelihood is a double integral over `mu` and `f`: the integral over `mu` is again an incomplete Beta function and the integral over `f` is computed with `stats::integrate()`. The likelihood under extinction (equation 10) is the average over `k`; the likelihood under persistence is the case `k = T`.

Input. Periods are the integer time units from the observation origin, the earliest time in `certain` or `uncertain` (zero-count rows included), to `test_year`, both included. A period is a sighting if its count is positive (several sightings in one period count once). As in Appendix S1 (whose example record starts with a sighting), a sighting in the first period counts as a sighting; add a zero-count row to start observation before the first sighting. Uncertain sightings may occur anywhere in the record, including before the first certain sighting.

Published conventions. With these conventions the bars of Fig. 4 for the paper's own model (Birdlife records, observation from `t0` in Table 3 to 2011) are reproduced. The bars of Fig. 3 and the 0.697 quoted for the greater stick-nest rat are reproduced when the sighting at the origin is not counted (the origin given as a zero-count row, as in Table 2) and the record has `2011 - t0` periods, i.e. `test_year = 2010`.

Package constructs: the exact computation (the paper uses MCMC), the `extinction_time` output, and the requirement that `test_year` be later than every sighting. The paper's eqn 14 (continuous extinction time) and its prior for the false-positive rate built from the probability that a sighting is true (eqn 13, Table 1; used for the light grey bars of Fig. 4) are not implemented: the paper does not give enough detail to reproduce those results. Appendix S1 has a single uncertain record; so has this function (the uncertain sightings are grouped together, as in the dark grey bars of Fig. 4).

Value

A list of class `ede_bayes` with components `bayes_factor` (`l(s | H0) / l(s | H1)`, equation 3), `posterior` (probability that the species is extant, equation 2), `prior`, `test_year`, `method`, `periods` (`T`), `rate_max`, and `extinction_time`: a data frame with, for each period after the last certain sighting, the posterior probability that it is the first period in which the species is absent (these probabilities sum to `1 - posterior`).

References

Lee, T. E., McCarthy, M. A., Wintle, B. A., Bode, M., Roberts, D. L. & Burgman, M. A. (2014). Inferring extinctions from sighting records of variable reliability. Journal of Applied Ecology, 51(1), 251-258.

Examples

```
# Greater stick-nest rat (Lee et al. 2014, Table 2), observation from 1835
yrs <- c(1844, 1857, 1879, 1896, 1897, 1907, 1920, 1921, 1930, 1933,
         1938, 1986)
rat <- sighting_data(data.frame(year = c(1835, yrs),
                                count = c(0, rep(1, 12))))
lee2014(rat, test_year = 2010)$posterior  # Fig. 3: 0.697
# the 1986 sighting treated as uncertain
rat_c <- sighting_data(data.frame(year = c(1835, yrs[-12]),
                                  count = c(0, rep(1, 11))))
rat_u <- sighting_data(data.frame(year = 1986, count = 1))
lee2014(rat_c, rat_u, test_year = 2010)
```

`marshall1994` *Marshall (1994) distribution-free confidence intervals on a range end-point*

Description

Confidence interval on the end-point of a stratigraphic (or sighting) range that does not assume randomly placed finds, only that gap sizes come from a continuous distribution with no trend along the section. The `C` confidence interval equals the `C`th percentile of the gap-size distribution (Marshall, 1994, p. 460), which is bracketed by order statistics of the `N` observed gaps. With `F(x)` the binomial `Binomial(N, C)` distribution function and `gamma = (1 - conf_prob) / 2`: the lower limit on the interval length is the `(x + 1)`th smallest gap, `x` the largest integer `< N` with `F(x) <= gamma` (equation 1); the upper limit is the `(x + 1)`th smallest gap, `x` the smallest integer with `F(x) >= 1 - gamma` (equation 2). With probability `conf_prob` (the paper's confidence probability `1 - 2 gamma`) the length of the `C` interval lies between the two. If no `x < N` satisfies equation 1 there is no lower limit; if only `x = N` satisfies equation 2 there is no upper limit and the interval is unbounded (e.g. a 95% interval needs at least 72 gaps at `conf_prob = 0.95`).

Usage

```
marshall1994(
  sd,
  alpha = 0.05,
  conf_prob = 0.95,
  horizons = c("distinct", "count")
)
```

Arguments

`sd`	A sighting_data object.
`alpha`	One minus the confidence level `C` of the interval, in (0, 1).
`conf_prob`	Confidence probability `1 - 2 gamma` with which the length of the `C` interval is bracketed, in (0, 1). The paper uses 0.95 or 0.99.
`horizons`	`"distinct"` (each distinct positive-count time is one horizon) or `"count"` (a count of `k` is `k` horizons at that time).

Details

The method gives a range of possible interval lengths rather than one interval. The returned `upper` is the conservative end, last find plus the upper-limit gap (`Inf` if there is none); `upper_min` is the last find plus the lower-limit gap, so the `C` bound lies in `[upper_min, upper]` with probability `conf_prob`. This mapping onto ede_estimate is a package construct.

Gaps are the distances between successive find horizons. By default (`horizons = "distinct"`) each distinct time with a positive count is one horizon, as in `strauss1989()`. With `horizons = "count"` a row with count `k` stands for `k` horizons at the same (binned) time, i.e. `k - 1` gaps of size zero, as in the paper's Metrarabdotos ages binned to 0.05 Myr. Zero-count rows are not

used. The end point is the one at the largest `time`: for a first appearance, or for ages in Ma, code `time = -age` (or `-height`). The method assumes continuous sampling; it is not meant for discretely sampled cores (use `weiss_marshall1999()` for those). Check for a trend in gap size with `marshall1994_trend()` first.

Value

An ede_estimate object with `estimate = NA`, `lower = NA`, `upper` (last find plus upper-limit gap, `Inf` if not available) and `interval_type = "one-sided"`, plus components `upper_min` (last find plus lower-limit gap, `NA` if not available), `ci_length` (named `lower`/`upper` limits on the interval length), `gap_ranks` (ranks `l` and `u` of the gaps used, `NA` if not available), `n_gaps`, `conf_prob` and `last_sighting`.

References

Marshall, C. R. (1994). Confidence intervals on stratigraphic ranges: partial relaxation of the assumption of randomly distributed fossil horizons. Paleobiology, 20(4), 459-469.

`marshall1994_trend` *Marshall (1994) transposition test for a trend in gap size*

Description

Tests the assumption of `marshall1994()` (and of the classical intervals) that gap size is not correlated with stratigraphic position. Following Marshall (1994, p. 462-463), the gaps are listed in stratigraphic order and the number of transpositions `T` (pairs in which a larger gap precedes a smaller one) is counted; each group of `m` equal gaps adds `m (m - 1) / 4`. `T` near 0 indicates gaps increasing in size, `T` near `T_max = N (N - 1) / 2` gaps decreasing in size.

Usage

```
marshall1994_trend(sd, horizons = c("distinct", "count"))
```

Arguments

`sd`	A sighting_data object.
`horizons`	`"distinct"` (each distinct positive-count time is one horizon) or `"count"` (a count of `k` is `k` horizons at that time).

Details

The paper compares `min(T, T_max - T)` with tabulated critical values (its Table 3, from Kreyszig 1979). Here a p-value is computed instead from the exact null distribution of the number of transpositions in a random ordering of `N` distinct gaps: `p = min(1, 2 P(I <= min(T, T_max - T)))`, with a non-integer `T` (from ties) rounded down. This two-sided p-value is a package construct. Table 3's critical values are two-sided (each tail `p / 2`): the exact distribution reproduces 67 of its 72 entries, and the other five differ by one at borderline probabilities.

Value

A list of class `marshall1994_trend` with components `transpositions` (`T`, including the tie correction), `t_max`, `n_gaps`, `direction` (`"increasing"` if `T < T_max / 2`, `"decreasing"` if `T > T_max / 2`, otherwise `"none"`), `p_value` and `gaps` (in stratigraphic order).

References

Marshall, C. R. (1994). Confidence intervals on stratigraphic ranges: partial relaxation of the assumption of randomly distributed fossil horizons. Paleobiology, 20(4), 459-469.

`marshall1997` *Marshall (1997) generalized confidence interval for a range end-point*

Description

Confidence interval on the end-point of a stratigraphic (or sighting) range when the probability of recovering a find varies along the axis according to a user-supplied fossil recovery potential `f(t)`. With `H` distinct find horizons, the observed range from the first (FAD) to the last (LAD) find carries recovery mass `R* = integral of f from FAD to LAD` (equation 10 of Marshall, 2010). The `C = 1 - alpha` interval extends beyond the LAD by the distance `r` whose recovery mass `r* = integral of f from LAD to LAD + r` (equation 11) equals `R* [(1 - C)^(-1 / (H - 1)) - 1]` (equation 9, the generalization of the classical interval, equation 4). If even the whole recovery mass beyond the LAD (up to `max_time`) is smaller than that, the interval is unbounded and `upper = Inf`, as in the short half-life case of Marshall (2010, Case 3). With constant `f` the method reduces to the classical interval of `strauss1989()`.

Usage

```
marshall1997(
  sd,
  recovery = function(t) rep(1, length(t)),
  alpha = 0.05,
  test_year = NULL,
  max_time = Inf
)
```

Arguments

`sd`	A sighting_data object.
`recovery`	A vectorized function of time giving the (relative) fossil recovery potential, non-negative. Only ratios of its integrals matter, so it need not be normalized. The default (constant) gives the classical interval.
`alpha`	Significance level, in (0, 1); the confidence level is `1 - alpha`.
`test_year`	Optional time after the last find at which to compute the confidence that the end-point had already been reached.
`max_time`	Upper limit of the axis on which `recovery` is defined (e.g. the top of the section). Defaults to `Inf`.

Details

The confidence that the end-point lies before a candidate time `T`, `C = 1 - [R* / (R* + r*)]^(H - 1)` with `r*` the recovery mass from the LAD to `T` (equations 3 and 6 with `R*`, `r*` in place of the range and its extension), is returned as `confidence` when `test_year` is given.

Each distinct time with a positive count is one find horizon; counts are ignored and zero-count rows are not used, as in `strauss1989()`. The end point estimated is always the one at the largest `time`: to place an interval on a last appearance dated in Ma, code `time = -age`; for a first appearance, code `time = -height` (or `time = age`). `recovery` must be written on the same coded axis.

The formulas are from Marshall (2010, Box 3), who attributes the method to Marshall (1997); they have been checked against the worked example of Marshall (2010, section 4.1), not against Marshall (1997) itself. The integrals are computed numerically with `stats::integrate()`, so a discontinuous (e.g. stepwise) recovery potential should be supplied with breakpoints smoothed or with care.

Value

An ede_estimate object with `estimate = NA`, `lower = NA`, `upper =` the end of the one-sided `1 - alpha` interval (possibly `Inf`) and `interval_type = "one-sided"`, plus components `last_sighting` (the LAD), `recovery_range` (`R*`), `recovery_extension` (the `r*` required) and, if `test_year` is given, `confidence`.

References

Marshall, C. R. (1997). Confidence intervals on stratigraphic ranges with non-random distributions of fossil horizons. Paleobiology, 23(2), 165-173.

Marshall, C. R. (2010). Using confidence intervals to quantify the uncertainty in the end-points of stratigraphic ranges. In J. Alroy & G. Hunt (Eds.), Quantitative Methods in Paleobiology (pp. 291-316). The Paleontological Society Papers, 16.

`mcinerny2006` *McInerny, Roberts, Davy & Cribb (2006) sighting-rate persistence test*

Description

A modification of the Solow (1993a) persistence test that conditions on the sighting rate observed up to the last sighting, `s = n / t[n]`, instead of on the length of the whole observation window. For candidate time `T`, the p-value is `(1 - s)^(T - t[n])` (equation 2 of the source paper). This makes the test comparable across records with very different total observation periods.

Usage

```
mcinerny2006(sd, alpha = 0.05, test_year, data_out = FALSE)
```

Arguments

sd	A sighting_data object.
alpha	Significance level, in (0, 1).
test_year	Latest time to test. Must be later than the last sighting. A bound beyond test_year is reported as NA with a warning.
data_out	If TRUE, return the p-value curve instead of the upper bound.

Details

The record is treated as binary: n counts time units with at least one sighting, as in the source paper. The earliest supplied time defines the observation origin. If that time has a sighting, it is used as the start date and not counted, so n is reduced by one, as in the paper. If the record starts with a zero-count row, that row is the known start date and every sighting time is counted.

The model is in whole time units, so the returned upper bound is the first time unit after the last sighting at which persistence is rejected, which is how the paper itself reports the test (Fig. 1).

Value

An ede_estimate object with estimate = NA and the one-sided upper bound in upper, or (if data_out = TRUE) a data frame with columns time and chance (p-values).

References

McInerny, G. J., Roberts, D. L., Davy, A. J., & Cribb, P. J. (2006). Significance of sighting rate in inferring extinction and threat. Conservation Biology, 20(2), 562-567.

---

monk_seal *Sighting record of the Caribbean monk seal*

Description

Reliable sightings of the Caribbean monk seal (*Monachus tropicalis*) used as the example in Solow (1993a), Burgman et al. (1995) and Solow (2005). There is no natural start of the observation period, so the first sighting defines the origin.

Usage

```
monk_seal
```

Format

A data frame with 5 rows: year and count.

Source

Solow, A. R. (1993a). Inferring extinction from sighting data. Ecology, 74(3), 962-964.

Examples

```
solow1993(sighting_data(monk_seal), test_year = 2000)
solow1993_bayes(sighting_data(monk_seal), test_year = 1992)
```

`ole` *Optimal Linear Estimation of extinction date*

Description

Estimates the extinction date by optimal linear estimation (Cooke, 1980), as applied by Roberts & Solow (2003). The joint distribution of the `k` most recent sighting times (the extreme order statistics) is assumed to have the Weibull extreme-value form; the estimate is the weighted sum of those times given by equations 12-15 of Solow (2005).

Usage

```
ole(sd, alpha = 0.05, k = NULL)
```

Arguments

`sd`	A sighting_data object.
`alpha`	Significance level for the confidence interval, in (0, 1).
`k`	Number of most recent sighting events to use. Counts greater than one are treated as independent events at the recorded time; Solow (2005) recommends combining sightings that are obviously related. By default, all sighting events are used. Must be at least 3; Solow (2005) suggests at least 5, and notes that the extreme-value approximation may fail if `k` is too large.

Details

The confidence interval is the two-sided `1 - alpha` interval of Roberts & Solow (2003). Its upper limit does not exist when `-log(alpha / 2) >= k`; it is then returned as `Inf` with a warning. `solow2005()` inverts the corresponding test one-sidedly (Solow, 2005, equation 17), so its bound at level `alpha` equals the upper limit here at level `2 * alpha`.

Value

An ede_estimate object.

References

Cooke, P. (1980). Optimal linear estimation of bounds of random variables. Biometrika, 67(1), 257-258.

Roberts, D. L., & Solow, A. R. (2003). Flightless birds: When did the dodo become extinct? Nature, 426(6964), 245.

Solow, A. R. (2005). Inferring extinction from a sighting record. Mathematical Biosciences, 195(1), 47-55.

orchid *Collection record of a Mauritian orchid*

Description

Collections of *Angraecum* aff. *rutenbergianum*, one of the nine Mauritian orchids analysed by Solow & Roberts (2003). With 1996 as the end of the observation period, their nonparametric test gives a p-value of 0.10.

Usage

```
orchid
```

Format

A data frame with 11 rows: `year` and `count`.

Source

Solow, A. R., & Roberts, D. L. (2003). A nonparametric test for extinction based on a sighting record. Ecology, 84(5), 1329-1332.

Examples

```
solow_roberts2003(sighting_data(orchid), test_year = 2030)
```

robson1964 *Robson & Whitlock (1964) truncation point estimator*

Description

Estimates the extinction date by a Quenouille (jackknife-type) bias reduction of the most recent sighting. If the two most recent distinct sighting times are `t[n-1]` and `t[n]`, the estimate is `t[n] + (t[n] - t[n-1])` and the one-sided upper confidence bound is `t[n] + (t[n] - t[n-1]) (1 - alpha) / alpha`.

Usage

```
robson1964(sd, alpha = 0.05)
```

Arguments

`sd`	A sighting_data object.
`alpha`	Significance level, in (0, 1), for the approximate one-sided confidence interval.

Details

The estimator is mean-unbiased to order `n^-2` in general and exactly unbiased for a uniform distribution. The bound is exact for a uniform distribution and asymptotically valid when the density is continuous and positive at the endpoint; its coverage is below `1 - alpha` when the density decreases towards the endpoint, as with a declining sighting rate.

Only distinct sighting times are used; counts are ignored, so the paper's order statistics are applied to distinct times.

Value

An ede_estimate object.

References

Robson, D. S., & Whitlock, J. H. (1964). Estimation of a truncation point. Biometrika, 51(1/2), 33-39.

---

rout2009 *Rout et al. (2009) Bayes factor with a declining sighting rate*

Description

Bayesian test of persistence (e.g. of an invasive species after an eradication programme) in which the pre-extinction sighting rate declines as `lambda(s) = m * s^(-a)` (equation 6 of Rout, Salomon & McCarthy, 2009), with `m` a constant and `0 <= a < 1`. Integrating over `m` (with the prior of Solow, 1993a) gives the Bayes factor in favour of persistence to `S`: `B = (n(1 - a) - 1) / ((S / s[n])^(n(1 - a) - 1) - 1)` (Rout et al., 2009, unnumbered equation after equation 6; derivation in their Appendix S1). With `a = 0` it is the Bayes factor of Solow (1993a) (Rout et al., equation 2), and this function then equals `solow1993_bayes()`. When `n(1 - a) = 1` the limit `1 / log(S / s[n])` is used. The posterior probability of persistence is `1 / (1 + (1 - prior) / (prior * B))` (equation 1).

Usage

```
rout2009(sd, test_year, a = 0, prior = 0.5, data_out = FALSE)
```

Arguments

`sd`	A sighting_data object.
`test_year`	Candidate time `S` (e.g. the current survey number). Must be later than the last sighting.
`a`	Rate of decline of the sighting rate, in `[0, 1)`. Rout et al. (2009) estimated `a = 0.1` for *Helenium amarum*.
`prior`	Prior probability that the species persisted to `test_year`, in (0, 1).
`data_out`	If `TRUE`, return the Bayes factor and posterior probability at integer steps after the last sighting up to `test_year`.

Details

Rout et al. (2009) measure time in surveys: `S` surveys have been made and the species was seen in `n` of them, the last being survey `s[n]`. Here times are measured from the observation origin, the earliest time in `sd`, and sighting events are counted as in `solow1993()`: a sighting at the origin defines it and is not counted. To follow the paper, supply one row per survey with binary counts and a leading zero-count row for survey 0.

The paper's model has the rate declining with the number of surveys, so the result depends on where the origin is placed; this is a property of the model, not of the implementation.

Value

A list of class `ede_bayes` with components `bayes_factor`, `posterior`, `prior`, `test_year`, `method` and `a`; or (if `data_out = TRUE`) a data frame with columns `time`, `bayes_factor` and `posterior`.

References

Rout, T. M., Salomon, Y., & McCarthy, M. A. (2009). Using sighting records to declare eradication of an invasive species. Journal of Applied Ecology, 46(1), 110-117.

Solow, A. R. (1993a). Inferring extinction from sighting data. Ecology, 74(3), 962-964.

See Also

`rout2009_absences()` for the number of absent surveys after which eradication should be declared.

`rout2009_absences` *Optimal number of absent surveys before declaring eradication*

Description

Decision rule of Rout, Salomon & McCarthy (2009): after `d` consecutive surveys without a sighting, the net expected cost of stopping, in units of the cost of escape and damage `C[e]`, is `NEC(d) / C[e] = (d - 1) R + p(extant` (equation 3, the "rule of thumb"), where `R = C[s] / C[e]` is the ratio of the cost of one survey to the cost of escape, and `p(extant | d)` is the posterior probability of persistence of `rout2009()` with `S = s[n] + d`. The optimal number of absent surveys `d*` minimises `NEC(d)` over `d = 0, 1, ..., max_absent`.

Usage

```
rout2009_absences(sd, cost_ratio, a = 0, prior = 0.5, max_absent = 1000)
```

Arguments

`sd`	A sighting_data object.
`cost_ratio`	The cost ratio `R = C[s] / C[e]`, positive.
`a`	Rate of decline of the sighting rate, in `[0, 1)`. Rout et al. (2009) estimated `a = 0.1` for *Helenium amarum*.
`prior`	Prior probability that the species persisted to `test_year`, in (0, 1).
`max_absent`	Largest number of absent surveys considered.

Details

For the constant-rate model (a = 0) the closed-form approximation d* ~ s[n] ((n - 1)^2 prior / (s[n] (1 - prior) R))^
(equation 5) is also returned. The paper gives no approximation for a > 0, so it is NA then. The paper's stochastic dynamic programme is not implemented.

Value

A list with components optimal_absent (d*), stop_time (last sighting time plus d*), approximation (equation 5, or NA), nec (a data frame with columns absent, time, posterior and nec, the latter in units of C[e]), n, sn, a, prior and cost_ratio.

References

Rout, T. M., Salomon, Y., & McCarthy, M. A. (2009). Using sighting records to declare eradication of an invasive species. Journal of Applied Ecology, 46(1), 110-117.

---

sighting_data *Construct a validated sighting record*

Description

Builds the common input object used by every estimator in the package: a time-ordered table of sighting counts, checked for the conditions each estimator in the sighting-record literature assumes (numeric, non-negative, no duplicated times).

Usage

```
sighting_data(data, time_col = 1L, count_col = 2L)
```

Arguments

data	A data frame or matrix. By default the first column is read as time (e.g. year) and the second as the number of sightings recorded at that time. The earliest supplied time defines the observation origin; include an initial zero-count row when observation began before the first sighting.
time_col, count_col	Column name or position for time and sighting count.

Value

An object of class sighting_data: a data frame with columns time and count, sorted by time.

`solow1993` *Solow (1993a) constant-rate persistence test*

Description

Parametric test of the null hypothesis that a species was still extant at a candidate test year, under a stationary Poisson sighting process. For candidate time `T`, the p-value is `(t[n] / T)^n` (equation 2 of Solow, 1993a), with times measured from the observation origin.

Usage

```
solow1993(sd, alpha = 0.05, test_year, data_out = FALSE)
```

Arguments

`sd`	A sighting_data object.
`alpha`	Significance level, in (0, 1). The returned upper bound is the earliest candidate time at which persistence is rejected at this level.
`test_year`	Latest time to test. Must be later than the last sighting. A bound beyond `test_year` is reported as `NA` with a warning.
`data_out`	If `TRUE`, return the p-value curve at integer steps after the last sighting instead of the upper bound.

Details

The earliest supplied time defines the observation origin. If that time has a sighting, one sighting defines the origin and is omitted, as in the Caribbean monk seal example of the source paper; any further sightings in the same time unit are counted. Every sighting event after that is counted, so a time unit with several sightings contributes several events to `n`.

Solow (1993a) defines only the test. Inverting it gives the one-sided upper confidence bound `t[n] / alpha^(1/n)` for the extinction time (Solow, 2005, equation 5), which is returned as `upper`.

Value

An ede_estimate object with `estimate = NA` and the one-sided upper bound in `upper`, or (if `data_out = TRUE`) a data frame with columns `time` and `chance`. `chance` contains p-values; the column name is retained for compatibility with EDE 0.1.0.

References

Solow, A. R. (1993a). Inferring extinction from sighting data. Ecology, 74(3), 962-964.

Solow, A. R. (2005). Inferring extinction from a sighting record. Mathematical Biosciences, 195(1), 47-55.

solow1993b *Solow (1993b) declining-population persistence test*

Description

Conditional test of the null hypothesis that a declining species was still extant at a candidate test year, under a non-stationary Poisson process with an exponentially declining rate. The test conditions on `n` and on the sum `S` of the sighting times, so the rate parameters are not estimated; the p-value is `F(t[n]) / F(T)`, where `F` is Fisher's distribution of the largest gap (equations 3-4 of the source paper).

Usage

```
solow1993b(sd, alpha = 0.05, test_year, data_out = FALSE)
```

Arguments

`sd`	A sighting_data object.
`alpha`	Significance level, in (0, 1). The returned upper bound is the earliest candidate time at which persistence is rejected at this level.
`test_year`	Latest time to test. Must be later than the last sighting. A bound beyond `test_year` is reported as `NA` with a warning.
`data_out`	If `TRUE`, return the p-value curve at integer steps after the last sighting instead of the upper bound.

Details

Sighting events are counted as in `solow1993()`: one sighting at the observation origin defines it and is omitted, and every other sighting event is counted, including repeated sightings in one time unit. Include a zero-count first row when observation began before the first sighting, as in the black-footed ferret example of the source paper.

Solow (1993b) defines only the test. Under the assumption that the species is extinct, the minimum variance unbiased estimate of the extinction time under the same truncated exponential model (Beg, 1982) is equation 8 of Solow (2005), also listed in Table 1 of McCrea et al. (2024): `T_E = t[n] + sum_{i=0}^{[s/t[n]]} (-1)^i C(n, i) (s - i t[n])^(n-1) / (n (n-1) sum_{i=0}^{[s/t[n]]-1} (-1`
which equals `t[n] + F(t[n]) / F'(t[n])`. It is returned as `estimate` (in the original time units, i.e. shifted by the origin). Both sums are evaluated as Fisher's largest-gap probabilities through the stable B-spline (Irwin-Hall) recursion rather than the alternating series, which loses all precision when the sighting times are concentrated near `t[n]`. The estimate is `NA` when it is undefined (all sightings other than the last at the origin with `n > 2`, or all at `t[n]`).

Inverting the test numerically gives a one-sided upper confidence bound for the extinction time, returned as `upper` (a package construct; Solow, 2005, Section 3, describes the numerical inversion). Under this model the bound may not exist: the p-value cannot fall below `F(t[n])`, however late the candidate time (Solow, 2005, Section 3).

Value

An ede_estimate object with the minimum variance unbiased estimate in `estimate` and the one-sided upper bound in `upper` (`NA`, with a warning, if it does not exist before `test_year`), or (if `data_out = TRUE`) a data frame with columns `time` and `chance` (p-values).

References

Solow, A. R. (1993b). Inferring extinction in a declining population. Journal of Mathematical Biology, 32(1), 79-82.

Solow, A. R. (2005). Inferring extinction from a sighting record. Mathematical Biosciences, 195(1), 47-55.

Beg, M. A. (1982). Optimal tests and estimators for truncated exponential families. Metrika, 29, 103-113.

Solow, A. R., & Helser, T. (2000). Detecting extinction in sighting data. In S. Ferson & M. Burgman (Eds.), Quantitative Methods for Conservation Biology (pp. 1-6). Springer, New York.

McCrea, R. S., Cheale, T., Campillo-Funollet, E., & Roberts, D. L. (2024). Inferring species extinction from sighting data. Cambridge Prisms: Extinction, 2, e19.

`solow1993_bayes` *Solow (1993a) Bayes factor for persistence*

Description

Bayesian counterpart of `solow1993()` under the same stationary Poisson model. With a uniform prior on the extinction time over the observation period and the noninformative prior `dm / m` on the sighting rate, the Bayes factor in favour of persistence to `T` is `B = (n - 1) / ((T / t[n])^(n - 1) - 1)` (equation 3 of Solow, 1993a), which tends to `1 / log(T / t[n])` when `n = 1`. Given the prior probability `prior` of persistence, the posterior probability is `1 / (1 + (1 - prior) / (prior * B))`.

Usage

```
solow1993_bayes(sd, test_year, prior = 0.5, data_out = FALSE)
```

Arguments

`sd`	A sighting_data object.
`test_year`	Candidate time `T`. Must be later than the last sighting.
`prior`	Prior probability that the species persisted to `test_year`, in (0, 1).
`data_out`	If `TRUE`, return the Bayes factor and posterior probability at integer steps after the last sighting up to `test_year`.

Details

Times are measured from the observation origin and sighting events are counted as in `solow1993()`.

Value

A list of class ede_bayes with components bayes_factor, posterior, prior, test_year and method; or (if data_out = TRUE) a data frame with columns time, bayes_factor and posterior.

References

Solow, A. R. (1993a). Inferring extinction from sighting data. Ecology, 74(3), 962-964.

---

solow2005 *Solow (2005) Weibull extreme-value persistence test*

Description

Tests persistence using the Weibull extreme-value model for the k most recent sighting events, the model behind ole(). For a candidate time T, the approximate p-value is equation 16 of Solow (2005), with the shape parameter estimated by equation 15.

Usage

```
solow2005(sd, alpha = 0.05, test_year, data_out = FALSE, k = NULL)
```

Arguments

sd	A sighting_data object.
alpha	Significance level, in (0, 1). The returned upper bound is the earliest candidate time at which persistence is rejected at this level.
test_year	Latest time to test. Must be later than the last sighting. A bound beyond test_year is reported as NA with a warning.
data_out	If TRUE, return the p-value curve at integer steps after the last sighting instead of the upper bound.
k	Number of most recent sighting events to use. Counts greater than one are treated as independent events at the recorded time; Solow (2005) recommends combining sightings that are obviously related. By default, all sighting events are used. Must be at least 3; Solow (2005) suggests at least 5.

Details

Inverting the test gives the one-sided upper confidence bound of equation 17, (t[1] - c t[k]) / (1 - c) with c = (-log(alpha) / k)^m, returned as upper. The bound does not exist when -log(alpha) >= k. Note that ole() reports the two-sided interval of Roberts & Solow (2003), so its upper limit at level alpha equals this bound at level alpha / 2.

Value

An ede_estimate object with estimate = NA and the one-sided upper bound in upper, or (if data_out = TRUE) a data frame with columns time and chance (p-values).

References

Solow, A. R. (2005). Inferring extinction from a sighting record. Mathematical Biosciences, 195(1), 47-55.

`solow_roberts2003` *Solow & Roberts (2003) nonparametric persistence test*

Description

Tests the null hypothesis that a species persisted to a candidate time using only the two most recent distinct sighting times. For candidate time `T`, the p-value is `(t[n] - t[n-1]) / (T - t[n-1])` (equation 4 of the source paper). It is exact for a constant sighting rate and asymptotically correct otherwise.

Usage

```
solow_roberts2003(sd, alpha = 0.05, test_year, data_out = FALSE)
```

Arguments

`sd`	A sighting_data object.
`alpha`	Significance level, in (0, 1). The returned upper bound is the earliest candidate time at which persistence is rejected at this level.
`test_year`	Latest time to test. Must be later than the last sighting. A bound beyond `test_year` is reported as `NA` with a warning.
`data_out`	If `TRUE`, return the p-value curve at integer steps after the last sighting instead of the upper bound.

Details

The paper derives the test from the Robson & Whitlock (1964) interval; inverting it gives the one-sided upper confidence bound `t[n] + (t[n] - t[n-1]) (1 - alpha) / alpha`, returned as `upper`. It is the same bound as `robson1964()`.

Value

An ede_estimate object with `estimate = NA` and the one-sided upper bound in `upper`, or (if `data_out = TRUE`) a data frame with columns `time` and `chance` (p-values).

References

Solow, A. R., & Roberts, D. L. (2003). A nonparametric test for extinction based on a sighting record. Ecology, 84(5), 1329-1332.

strauss1989 *Strauss & Sadler (1989) confidence interval for the end of a range*

Description

Unbiased point estimator and classical one-sided confidence interval for the true endpoint of a temporal range, assuming finds are uniformly distributed between the true endpoints. With `n` finds spanning the range `y` to `z`, the estimate is `z + (z - y) / (n - 1)` (equation 8 of the source paper) and the interval is `(z, z + c (z - y))` with `c = alpha^(-1 / (n - 1)) - 1` (equations 13, 16 and 20).

Usage

```
strauss1989(sd, alpha = 0.05)
```

Arguments

`sd`	A sighting_data object.
`alpha`	Significance level, in (0, 1).

Details

Each distinct sighting time is one find, as each fossil horizon is in the paper; counts are ignored. The paper's model is continuous, so with dense yearly records the unbiasedness and the coverage hold only approximately. Both endpoints are estimated from the data, so the observation origin (a leading zero-count row) is not used.

Value

An ede_estimate object containing the unbiased point estimate and a one-sided confidence interval.

References

Strauss, D., & Sadler, P. M. (1989). Classical confidence intervals and Bayesian probability estimates for ends of local taxon ranges. Mathematical Geology, 21(4), 411-427.

strauss1989_bayes *Strauss & Sadler (1989) Bayesian estimate of the end of a range*

Description

Bayesian counterpart of strauss1989() (Section 4.3 of the source paper). Finds are independent and uniform between the true endpoints theta1 < theta2 (equation 23), and the prior on the endpoints is uniform over a < theta1 < theta2 < b (equation 22, there scaled to (0, 1)), where (a, b) = prior_range is the interval that the endpoints are known to lie in. With n finds spanning y to z, the posterior density of the extinction time theta2 is h(theta2) = (n - 2) [(theta2 - y)^(1-n) - (theta2 - a)^(1-n)] / u_n on z < theta2 < b (equation 26), with u_n = (z-y)^(2-n) - (b-y)^(2-n) - (z-a)^(2-n) + (b-a)^(2-n) (equation 25). The estimate is the posterior mean (equation 27), the Bayes point estimate under quadratic loss. Since the density decreases from its maximum at z, the shortest 1 - alpha posterior probability interval is (z, q), with q the 1 - alpha posterior quantile (as shaded in Fig. 3 of the paper); it is returned in lower and upper.

Usage

```
strauss1989_bayes(sd, alpha = 0.05, prior_range, data_out = FALSE)
```

Arguments

sd	A sighting_data object.
alpha	Posterior probability outside the interval, in (0, 1).
prior_range	Numeric vector c(a, b): the prior is uniform for the endpoints over a < theta1 < theta2 < b. a must be finite and earlier than the first find; b must be later than the last find and may be Inf.
data_out	If TRUE, return the posterior on a grid of 200 times from z to the 0.999 posterior quantile instead of the estimate.

Details

The paper writes the equations for (a, b) = (0, 1); they are applied here after the corresponding linear change of scale, which leaves the posterior unchanged. The paper suggests (0, Inf) for taxa that may be extant; b = Inf is allowed (the posterior is proper for n >= 3, but its mean is infinite for n = 3). The paper's example takes (0, 1500) m for the ammonite section of Seymour Island.

As in strauss1989(), each distinct sighting time is one find and counts are ignored; the observation origin (a leading zero-count row) is not used, but it may be a natural choice for a. The paper requires n > 2.

Value

An ede_estimate object with the posterior mean in estimate and the one-sided posterior probability interval (z, q) in lower and upper, plus components density and cdf (functions of time giving the posterior density and distribution function of the extinction time) and prior_range. If data_out = TRUE, a data frame with columns time, density and chance, where chance is the posterior probability that the extinction time is at or before time.

References

Strauss, D., & Sadler, P. M. (1989). Classical confidence intervals and Bayesian probability estimates for ends of local taxon ranges. Mathematical Geology, 21(4), 411-427.

`strauss1989_curve` *Confidence curve for the Strauss & Sadler (1989) upper bound*

Description

Evaluates the upper confidence bound of `strauss1989()`, `z + c (z - y)` with `c = alpha^(-1 / (n - 1)) - 1`, for alpha from 0.01 to 1 in steps of 0.01, for plotting instead of a single bound.

Usage

```
strauss1989_curve(sd)
```

Arguments

`sd`	A sighting_data object.

Value

A data frame with columns `time` and `chance`, where `chance` is the confidence level `1 - alpha` of the bound `time`, i.e. the probability that the true endpoint is at or before `time`.

`thompson2013` *Thompson et al. (2013) probability of persistence with uncertain sightings*

Description

Discrete-time model of Thompson, Lee, Stone, McCarthy & Burgman (2013) for a record of one or more observation types (e.g. specimens, photographs, sound recordings) seen or not seen in each time unit. The first type is definite (no false sightings); the others may be false. For type `alpha`, `p[alpha]` is the probability of not seeing that type in a time unit when the species is extant and `q[alpha]` the probability of not seeing it when the species is extinct (`q = 1` for the definite type). Time units are independent, so the probability of the record given persistence to the last time unit `T` is the product of the `p_i` (equation 5), and given extinction in time unit `j` it is `p_1 ... p_(j-1) q_j ... q_T` (equation 7), where `p_i` and `q_i` are products over types of `p` or `1 - p` (`q` or `1 - q`) according to whether the type was seen (equation 27). The probability of persistence is equation 9,

Usage

```
thompson2013(
  definite,
  uncertain = NULL,
  test_year,
  p,
  q = NULL,
  prior_extant = NULL,
  prior_extinct = NULL,
  average = c("annealed", "quenched"),
  nodes = NULL
)
```

Arguments

definite	A sighting_data object for the definite (certain) observation type.
uncertain	NULL, a sighting_data object, or a list of them, one per uncertain observation type.
test_year	Last time unit T of the record. Must be later than the last definite sighting and not earlier than any other sighting.
p	List (or numeric vector, for point values) with one element per observation type, definite first: the probability of not observing that type in a time unit given the species is extant. Each element is a point value in (0, 1) or a range c(a, b) with 0 <= a < b <= 1.
q	As p, for the uncertain types only: the probability of not observing that type given the species is extinct. Point values may be 1.
prior_extant	Prior probability P(X_T) that the species is extant at test_year. Default T_N / T (equation 19).
prior_extinct	Prior probability P_E of extinction in each time unit after the last definite sighting. Default 1 / T (equation 19).
average	"annealed" (default) or "quenched"; ignored when all parameters are point values.
nodes	Number of Gauss-Legendre nodes per random parameter for the quenched average with more than one random parameter. Default chooses about 2e5 nodes in total.

Details

P(X_T | s) = prod(p_i) P(X_T) / (prod(p_i) P(X_T) + sum_j p_1...p_(j-1) q_j...q_T P_E),

with j running over the time units after the last definite sighting T_N, P(X_T) the prior probability of persistence and P_E the prior probability of extinction in each of those time units.

Each p[alpha] and q[alpha] is either a point value or a range c(a, b) read as a uniform distribution, as in the paper. With point values the result is the deterministic estimate P_D (equation 14 when there is only the definite type). With ranges the result is averaged (Section 2.2): "annealed" averages the numerator and the denominator of equation 9 separately (a ratio of averages; equation 12), "quenched" averages the ratio itself (equation 11). Parameters of different types, and p

and q of the same type, are independent. Each parameter is held fixed across time units within a realisation.

The record is binary: a time unit counts as a sighting of a type when its count is positive. Time units are numbered from 1 at the earliest time of any record (so, unlike solow1993(), the first sighting is part of the record), and T = test_year - origin + 1. Times must be whole numbers.

Package constructs: the paper estimates the averages by simulation; here the annealed averages are computed exactly, as products of one-dimensional integrals (stats::integrate(), on the log scale), and the quenched average by stats::integrate() for a single random parameter or by tensor-product Gauss-Legendre quadrature otherwise. The Bayes factor (deterministic and annealed only) is the likelihood under persistence divided by the mean likelihood over the extinction years, which gives the paper's equation 13 in its two-year example; it is NA for the quenched average, which is not a ratio of likelihoods.

Value

A list of class ede_bayes with components bayes_factor (in favour of persistence, or NA), posterior (probability of persistence from equation 9), prior (prior_extant), test_year, method, plus prior_extinct and average ("deterministic", "annealed" or "quenched").

References

Thompson, C. J., Lee, T. E., Stone, L., McCarthy, M. A. & Burgman, M. A. (2013). Inferring extinction risks from sighting records. Journal of Theoretical Biology, 338, 16-22.

---

thompson2017 *Thompson et al. (2017) iterative probability that a taxon is extant*

Description

Year-by-year model of the probability P(X[t]) that a taxon is extant, combining records of varying reliability with unsuccessful dedicated and passive surveys (Thompson et al., 2017; the model behind the IUCN Red List extinction-probability tool). Starting from an initial probability p_initial at the beginning of start_year, each year updates the probability by equation 5 of the paper:

Usage

```
thompson2017(
  records,
  surveys,
  passive,
  p_initial,
  start_year = NULL,
  end_year = NULL,
  n_sim = 10000,
  data_out = FALSE
)
```

Arguments

records	Data frame of record years with columns year and pci_lower, pci_upper (or a single column pci), the probability of correct identification. NULL if there are no records.
surveys	Data frame of unsuccessful dedicated survey years with columns year, eps_lower, eps_upper, pi_lower, pi_upper, pr_lower, pr_upper (or single columns eps, pi, pr). NULL if there are no dedicated surveys.
passive	Passive-survey inputs applied to every gap year: a named numeric vector or list with elements eps_lower, eps_upper, pi_lower, pi_upper, pr_lower, pr_upper (or eps, pi, pr). Use c(eps = 0, pi = 0, pr = 0) for no passive surveillance.
p_initial	Probability that the taxon is extant at the beginning of start_year: one value or a c(lower, upper) range.
start_year	First year of the iteration; p_initial applies at its beginning. Defaults to the earliest year in records and surveys.
end_year	Last year of the iteration. Defaults to the latest year in records and surveys; years after the last record or survey are gap years. Records or surveys outside start_year:end_year are ignored. Both years are required when there are no records or surveys (e.g. the passive-only projection of section 8.2 and equation 14 of the paper).
n_sim	Number of Monte Carlo replicates for the percentile band; 0 skips the simulation.
data_out	If TRUE, return the year-by-year trajectory instead of the summary object.

Details

- record year: P(X[t+1]) = p(ci) + (1 - p(ci)) * P(X[t]) (equation 1), where p(ci) is the probability that the record is a correct identification of the extant taxon;
- unsuccessful survey year: P(X[t+1]) = (1 - eps * p(r) * p(i)) * P(X[t]) (equations 3-4), where eps is the proportion of the habitat surveyed, p(r) the probability that the taxon would have been recorded and p(i) the probability that it would have been reliably identified;
- gap year (neither a record nor a dedicated survey): the survey rule with the passive-survey inputs eps', p'(r) and p'(i) (section 4.2).

The model is not Bayesian (section 3 of the paper): P(X[t]) is an absolute probability updated by the rules above. Successful surveys are supplied as records.

Every input is given as a lower and upper bound, as in Table 1 of the paper, and uncertainty is propagated in the two ways of section 6:

- p_extant iterates equation 5 with the midpoint of every bound (the solid line of Figs. 1-3);
- lower and upper are the results of interval arithmetic. Equation 5 increases with P(X[t]) and p(ci) and decreases with eps, p(r) and p(i), so the exact bounds are obtained by iterating it with the corresponding extreme inputs (the light shaded region of Figs. 1-3);
- if n_sim > 0, each bound pair is also treated as the 5th and 95th percentiles of an independent normal distribution, inputs are drawn independently for every year (including each gap year)

and every replicate, and the 5th and 95th percentiles of the simulated `P(X[t])` are returned as `mc_lower` and `mc_upper` (the dark shaded region). The paper does not say how draws outside `[0, 1]` are handled; here they are truncated to `[0, 1]` (a package construct).

A point value can be supplied instead of a range by using a single column (e.g. `pci` instead of `pci_lower`/`pci_upper`).

Value

A list of class `ede_extant` with components `p_extant` (midpoint estimate of the probability that the taxon is extant at the end of `end_year`), `lower`, `upper` (interval-arithmetic bounds), `mc_lower`, `mc_upper` (Monte Carlo 5th and 95th percentiles, or `NA` when `n_sim = 0`), `p_extinct` (`1 - p_extant`), `end_year` and `method`; or (if `data_out = TRUE`) a data frame with one row per year and columns `time`, `type` (`"record"`, `"survey"` or `"passive"`), `p_extant`, `lower`, `upper`, `mc_lower` and `mc_upper`, where each row gives the probability that the taxon is extant at the end of year `time`, after that year's data (as plotted in Figs. 1-3).

References

Thompson, C. J., Koshkina, V., Burgman, M. A., Butchart, S. H. M., & Stone, L. (2017). Inferring extinctions II: a practical, iterative model based on records and surveys. Biological Conservation, 214, 328-335.

Examples

```
# Alaotra Grebe, Table 1 of Thompson et al. (2017)
rec <- data.frame(
  year = c(1929, 1960, 1963, 1969, 1970, 1971, 1972, 1982, 1985, 1986, 1988),
  pci_lower = c(0.99, 0.95, 0.75, 0.75, 0.10, 0.10, 0.60, 0.60, 0.20, 0.20, 0.20),
  pci_upper = c(1.00, 0.99, 0.94, 0.94, 0.40, 0.40, 0.80, 0.80, 0.80, 0.70, 0.50)
)
sv_years <- c(1989, 1990, 1993, 1994, 1997, 1998, 1999, 2000, 2004, 2009)
sv <- data.frame(
  year = sv_years,
  eps_lower = ifelse(sv_years == 2000, 0.70, 0.80),
  eps_upper = ifelse(sv_years == 2000, 0.90, 0.95),
  pi_lower = 0.90, pi_upper = 0.95, pr_lower = 0.70, pr_upper = 0.90
)
pas <- c(eps_lower = 0, eps_upper = 0.05, pi_lower = 0.10, pi_upper = 0.65,
         pr_lower = 0.40, pr_upper = 0.60)
thompson2017(rec, sv, pas, p_initial = c(0.95, 0.99), end_year = 2017,
             n_sim = 0)
```

`thompson2019` *Thompson et al. (2019) Bayesian updating of the probability of persistence*

Description

Bayesian counterpart of thompson2017(). Each year's posterior is the next year's prior (Thompson et al., 2019, equations 1-6), so the probability that the taxon is extant after year t is P(X[t+1]) = 1 / (1 + B[t] * (1 / P(X[1]) - 1)) (equation 7), where P(X[1]) is the initial probability at the beginning of start_year and B[t] is the cumulative Bayes factor for extinction, the product of the yearly factors b[j] (equation 8). The yearly factors (equations 9-11) are

Usage

```
thompson2019(
  records,
  surveys,
  passive,
  p_initial = 0.5,
  start_year = NULL,
  end_year = NULL,
  data_out = FALSE
)
```

Arguments

records	Data frame of record years with columns year and pci_lower, pci_upper (or a single column pci), the probability of correct identification. NULL if there are no records.
surveys	Data frame of unsuccessful dedicated survey years with columns year, eps_lower, eps_upper, pi_lower, pi_upper, pr_lower, pr_upper (or single columns eps, pi, pr). NULL if there are no dedicated surveys.
passive	Passive-survey inputs applied to every gap year: a named numeric vector or list with elements eps_lower, eps_upper, pi_lower, pi_upper, pr_lower, pr_upper (or eps, pi, pr). Use c(eps = 0, pi = 0, pr = 0) for no passive surveillance.
p_initial	Probability P(X[1]) that the taxon is extant at the beginning of start_year: one value or a c(lower, upper) range, whose midpoint is used.
start_year	First year of the iteration; p_initial applies at its beginning. Defaults to the earliest year in records and surveys.
end_year	Last year of the iteration. Defaults to the latest year in records and surveys; years after the last record or survey are gap years. Records or surveys outside start_year:end_year are ignored. Both years are required when there are no records or surveys (e.g. the passive-only projection of section 8.2 and equation 14 of the paper).
data_out	If TRUE, return the year-by-year Bayes factors and probabilities instead of the summary object.

Details

- record year: b = 1 - p(ci);
- unsuccessful dedicated survey year: b = 1 / (1 - eps * p(r) * p(i));

- gap year (passive survey): the survey factor with `eps'`, `p'(r)` and `p'(i)`.

Inputs use the same format as `thompson2017()`. Following section 2.1 and the case study of the paper, ranges are replaced by their midpoints (including a range for `p_initial`); the paper does not propagate the ranges any further.

The result is an `ede_bayes` object, as returned by `solow1993_bayes()`. Its `bayes_factor` is in favour of persistence, `1 / B[t]`, so that it has the same orientation as in `solow1993_bayes()`; the paper's cumulative Bayes factor for extinction is returned as `bayes_factor_extinction`.

Value

A list of class `ede_bayes` with components `bayes_factor` (cumulative Bayes factor in favour of persistence, `1 / B`), `bayes_factor_extinction` (`B`, equation 8), `posterior` (probability that the taxon is extant at the end of `end_year`, equation 7), `prior` (`P(X[1])`), `test_year` (`end_year`) and `method`; or (if `data_out = TRUE`) a data frame with one row per year and columns `time`, `type` (`"record"`, `"survey"` or `"passive"`), `b` (yearly Bayes factor for extinction, equation 11), `bayes_factor_extinction` (cumulative `B[t]`) and `posterior` (probability extant at the end of year `time`).

References

Thompson, C. J., Kodikara, S., Burgman, M. A., Demirhan, H., & Stone, L. (2019). Bayesian updating to estimate extinction from sequential observation data. Biological Conservation, 229, 26-29.

Examples

```
# Alaotra Grebe from 1990 (Thompson et al., 2019, Table 1)
sv <- data.frame(year = c(1990, 1993, 1994, 1997, 1998, 1999),
                 eps_lower = 0.80, eps_upper = 0.95, pi_lower = 0.90,
                 pi_upper = 0.95, pr_lower = 0.70, pr_upper = 0.90)
pas <- c(eps_lower = 0, eps_upper = 0.05, pi_lower = 0.10, pi_upper = 0.65,
         pr_lower = 0.40, pr_upper = 0.60)
thompson2019(NULL, sv, pas, p_initial = 0.5, end_year = 1994)
```

`weiss_marshall1999` *Weiss & Marshall (1999) Bayesian end point for discretely sampled sections*

Description

Posterior distribution of the true end point `S` of a stratigraphic range (or of a sighting record) when the section is sampled at discrete points `t[1] < ... < t[J]`, each recorded as presence or absence. While the taxon is extant (`t[j] <= S`) it is found in a sample with probability `p`; after `S` it is never found (equations 1-2 of Weiss & Marshall, 1999). With `a[m]` presences up to the last presence `t[m]` and `b(S)` absences up to `S`, the likelihood is `p^a[m] (1 - p)^b(S)` for `S >= t[m]` (equation 3). The prior on `S` is exponential with mean longevity `l` (equation 4) and the prior on `p` is `Beta(c, d)`. Integrating `p` out gives the step-function posterior of equation 6: on `[t[j], t[j+1])` (with `t[J+1]`

= Inf) its mass is g[j] of equation 8, B(a[m] + c, b[j] + d) (exp(-t[j] / l) - exp(-t[j+1] / l)), with b[j] = j - a[m]. Quantiles, moments and the posterior of p (a mixture of Beta(a[m] + c, b[j] + d) densities, equations 7-8) follow in closed form (Appendix).

Usage

```
weiss_marshall1999(sd, prior_p, mean_longevity, alpha = 0.05, test_year = NULL)
```

Arguments

sd	A sighting_data object, one row per sampled horizon (count 0 = absence).
prior_p	Numeric vector c(c, d) of the Beta(c, d) prior on the detection probability p; both must be positive.
mean_longevity	Mean l of the exponential prior on the time to the end point (equation 4), in the units of time; positive and finite. The paper uses 800 m (a 21 Myr mean foraminiferal longevity converted to section thickness).
alpha	Posterior tail probability, in (0, 1), for upper.
test_year	Optional time after the last presence; if given, the posterior probability P(S > test_year | Y) that the taxon was still extant at that time is returned as persistence_prob.

Details

Each row of sd is one sampled horizon: a positive count is a presence and a zero count an absence. Every sampled horizon must be supplied, including the barren samples above the last presence, which carry the information about the end point; unsampled positions must not be given as rows. The earliest row is t[1], where the paper assumes the taxon is already extant (it need not be found there). Counts above one are treated as presence; the model uses presence/absence only.

Times may be stratigraphic heights or ages, but the end point estimated is always the one at the largest time. To estimate a last appearance from ages in Ma (which decrease upwards), code time = -age; to estimate a first appearance from heights, code time = -height. mean_longevity and the results are on the coded scale.

The reported estimate is the posterior median and upper is the 1 - alpha posterior quantile, so [t[m], upper] is the minimum-length 1 - alpha posterior interval the paper reports (S.95). The paper tabulates these as distances beyond the last presence (residual lifetime); here they are on the original time scale. Choosing c and d is part of the analysis: the paper pools the records of similar taxa, see weiss_marshall1999_prior().

Value

An ede_estimate object with estimate (posterior median of S), lower (last presence t[m]), upper (1 - alpha posterior quantile) and interval_type = "one-sided", plus components posterior_mean and posterior_sd of S, p_mean and p_sd (posterior mean and SD of p), prior_p, mean_longevity and, if requested, persistence_prob.

References

Weiss, R. E., & Marshall, C. R. (1999). The uncertainty in the true end point of a fossil's stratigraphic range when stratigraphic sections are sampled discretely. Mathematical Geology, 31(4), 435-453.

Marshall, C. R. (2010). Using confidence intervals to quantify the uncertainty in the end-points of stratigraphic ranges. In J. Alroy & G. Hunt (Eds.), Quantitative Methods in Paleobiology (pp. 291-316). The Paleontological Society Papers, 16.

`weiss_marshall1999_prior`

Beta prior for p pooled from related records (Weiss & Marshall, 1999)

Description

Builds the `Beta(c, d)` prior on the detection probability used by `weiss_marshall1999()` the way Weiss & Marshall (1999, p. 442-443) do for their Eastbourne foraminifera: starting from `Beta(0, 0)`, it adds the presences (`c`) and absences (`d`) of other, comparable taxa from the first sample up to, but excluding, each taxon's last presence (the last presence is excluded because it is a presence by definition).

Usage

```
weiss_marshall1999_prior(records)
```

Arguments

`records`	A list of sighting_data objects, one per comparable taxon, each with one row per sampled horizon (count 0 = absence).

Value

A named numeric vector `c(c = , d = )` to pass as `prior_p`.

References

Weiss, R. E., & Marshall, C. R. (1999). The uncertainty in the true end point of a fossil's stratigraphic range when stratigraphic sections are sampled discretely. Mathematical Geology, 31(4), 435-453.

Index

∗ **datasets**
  dodo, 9
  extinct_birds, 11
  ferret, 12
  monk_seal, 26
  orchid, 28

brook2019, 3
burgman1995, 5
burgman1995(), *11*

caley_barry2014, 7

dodo, 9

ede_compare, 10
ede_estimate, *3, 4*, *6*, 11, *14, 15*, *22, 23*, *25–27*, *29*, *32*, *34–38*, *46*
extinct_birds, 11

ferret, 12

jaric2010, 13
jaric2010(), *11*
jaric_roberts2014, 14

kodikara2020, 15
kodikara2021, 17

lee2014, 20

marshall1994, 22
marshall1994(), *23*
marshall1994_trend, 23
marshall1994_trend(), *23*
marshall1997, 24
mcinerny2006, 25
mcinerny2006(), *11*
monk_seal, 26

ole, 27
ole(), *3*, *9*, *35*
orchid, 28

robson1964, 28
robson1964(), *3*, *36*
rout2009, 29
rout2009(), *30*
rout2009_absences, 30
rout2009_absences(), *30*

set.seed(), *7*
sighting_data, *3*, *5*, *7*, *10*, *13–15*, *18*, *20*, *22–24*, *26–30*, 31, *32–40*, *46, 47*
solow1993, 32
solow1993(), *3*, *8*, *11*, *30*, *33, 34*, *41*
solow1993_bayes, 34
solow1993_bayes(), *29*, *45*
solow1993b, 33
solow1993b(), *11*
solow2005, 35
solow2005(), *11*, *27*
solow_roberts2003, 36
solow_roberts2003(), *11*
stats::integrate(), *20*, *25*, *41*
strauss1989, 37
strauss1989(), *3*, *22*, *24, 25*, *38, 39*
strauss1989_bayes, 37
strauss1989_bayes(), *11*
strauss1989_curve, 39

thompson2013, 39
thompson2017, 41
thompson2017(), *44, 45*
thompson2019, 43

weiss_marshall1999, 45
weiss_marshall1999(), *11*, *23*, *47*
weiss_marshall1999_prior, 47
weiss_marshall1999_prior(), *46*