

Package ‘XSLmodels’

October 18, 2026

Title Run and Optimize Cross-Situational Word Learning Models

Version 0.4.0

Maintainer George Kachergis <george.kachergis@gmail.com>

Description A collection of cognitive models of cross-situational word learning, together with methods and datasets for fitting and cross-validating them against experimental and corpus data, and for plotting the results.

URL <https://www.kachergis.com/XSLmodels/>,
<https://github.com/kachergis/XSLmodels>

BugReports <https://github.com/kachergis/XSLmodels/issues>

Encoding UTF-8

Depends R (>= 4.1.0)

Imports DEoptim, dplyr, ggplot2, purrr, rlang, tibble, tidyr

Suggests ganimate, knitr, rmarkdown, testthat (>= 3.0.0), viridis

VignetteBuilder knitr

LazyData true

Config/testthat/edition 3

License GPL (>= 3)

Config/roxygen2/version 8.0.0

NeedsCompilation no

Author George Kachergis [aut, cre] (ORCID:
<<https://orcid.org/0000-0003-4153-4167>>),
Mika Braginsky [aut] (ORCID: <<https://orcid.org/0000-0001-9039-3220>>),
Kevin Liang [aut]

Repository CRAN

Date/Publication 2026-10-18 14:50:07 UTC

Contents

baseline	3
bayesian_decay	3
benitez2020	4
create_cooc_matrix	5
cross_validated_group_fits	5
decay	7
fazly	7
fgt2009	8
fgt2009_rsa	9
fgt2009_sweep_alpha	10
fm_corpus	12
gangwani2011_category	13
get_crossvalidated_model_fit	14
get_example_ambiguous_condition	15
get_example_unambiguous_condition	15
get_fscore	16
get_group_model_fit	17
get_perf	18
get_roc	19
get_tp	20
get_train_test_split	20
guess_and_test	21
kachergis2012_highlighting	22
kachergis_initial_accuracy	22
kalman_filter	23
mafc_test	24
minerva2	25
multi_sampling	26
new_xslFit	26
plot_training_trials	27
predict_referent	28
propose_but_verify	29
pursuit	30
rem	31
rescorla_wagner	32
rollins_corpus	32
rsa_listener	33
show_datasets	34
show_models	35
softmax_rl	35
tilles	36
todam	37
uncfam	38
uncfam_attention	39
uncfam_elimination	39
uncfam_gamma	40

baseline

3

uncfam_predictive

42

uncfam_sampling

42

update_params

43

vlach_debrock2017

44

xslControl-class

44

xslData-class

45

xslFit-class

46

xsl_datasets

47

xsl_fit

47

xsl_holdout_datasets

48

xsl_run

49

Index

50

baseline	<i>Baseline model</i>
----------	-----------------------

Description

Baseline model

Usage

baseline()

Value

An object of class xslMod

Examples

```
mod <- baseline()  
xsl_run(mod, get_example_ambiguous_condition())
```

bayesian_decay	<i>Bayesian decay model</i>
----------------	-----------------------------

Description

Bayesian decay model

Usage

bayesian_decay(alpha, delta, ch_dec)

Arguments

alpha	Decay for word/object non-co-occurrence (0.1, 0.5, 0.9)
delta	Multiplier for word/object co-occurrence (1 for no increase)
ch_dec	Decision parameter (e.g., 1)

Value

An object of class `xslMod`

Examples

```
mod <- bayesian_decay(alpha = 0.5, delta = 1, ch_dec = 1)
xsl_run(mod, get_example_ambiguous_condition())
```

benitez2020

Benitez et al. (2020) temporal-structure-of-naming experiment data

Description

A named list of 7 `xslData` objects from Benitez, V. L., et al. (2020), "The temporal structure of naming events differentially affects children's and adults' cross-situational word learning" (data/materials: <https://osf.io/2hmxr/>): one per condition (Interleaved, Massed, Unstructured) x age-group (kids, adults) combination, except Unstructured's first training order, which only has child data. Unstructured has two training orders; its second and third were confirmed identical by the original extraction and are pooled as "order 2/3".

Usage

```
benitez2020
```

Format

An object of class `list` of length 7.

Details

8 symmetric word-object pairs (word *i* pairs with object *i*), 2 pairs per training trial, 12 trials, followed by an 8-item 2AFC test. `accuracy` and `response_matrix` are real per-word 2AFC accuracy and raw response counts for that condition/age-group combination (not an estimate) – structurally this could join `xsl_datasets`, but is kept standalone deliberately as held-out generalization-test data; see `xsl_holdout_datasets`. `data-raw/Benitez2020-rep/Benitez2020_XSLdata_extraction.R` extracts the raw OSF data into training orders, 2AFC test trials, and response matrices; `data-raw/add_benitez2020.R` builds the `xslData` objects from those.

Examples

```
benitez2020[["Interleaved, kids"]][extract_itex]accuracy
m <- xsl_run(baseline(), benitez2020[["Massed, adults"]][extract_itex]fits[[1]][extract_itex]matrix
mafc_test(m, benitez2020[["Massed, adults"]][extract_itex]test)
```

create_cooc_matrix	<i>Creates co-occurrence matrix from training trials</i>
--------------------	--

Description

Given a training order (list of words and objects per trial), this function returns a matrix tallying word-object co-occurrences across the trials. This matrix can be used to analyze the frequency with which certain words and objects appear together during the training phase.

Usage

```
create_cooc_matrix(train)
```

Arguments

train	A list representing the training data, where each element is a trial that includes both words and objects. The structure is expected to have sub-elements words and objects for each trial.
-------	---

Value

A matrix where each element represents the count of co-occurrences between a word (rows) and an object (columns). The row and column names correspond to the unique words and objects found in the training data, respectively.

Examples

```
create_cooc_matrix(xsl_datasets[[1]]$train)
```

cross_validated_group_fits	<i>Perform Cross-Validation on Group Fits of a Model</i>
----------------------------	--

Description

This function runs k-fold cross-validation for group fits of a specified model using a set of conditions. On each fold, the model is fit (via `xsl_fit()`) to the training conditions, and the fitted model is then evaluated (via `xsl_run()`) on the held-out test conditions to obtain an out-of-sample SSE.

Usage

```
cross_validated_group_fits(
  model,
  combined_data,
  lower,
  upper,
  n_folds = 5,
  control = xslControl(),
  deoptim_control = DEoptim::DEoptim.control(reltol = 0.001, NP = 100, itermax = 100),
  seed = NULL
)
```

Arguments

<code>model</code>	An object of class <code>xslMod</code> giving the model to cross-validate (its starting parameter values are only used to determine how many parameters <code>xsl_fit()</code> should optimize).
<code>combined_data</code>	A list of <code>xslData</code> conditions used for cross-validation.
<code>lower</code>	Numeric vector of lower bounds for the model's parameters.
<code>upper</code>	Numeric vector of upper bounds for the model's parameters.
<code>n_folds</code>	Number of folds (default: 5).
<code>control</code>	Control arguments passed to <code>xsl_fit()/xsl_run()</code> .
<code>deoptim_control</code>	Control parameters passed to <code>DEoptim()</code> (via <code>xsl_fit()</code>).
<code>seed</code>	Optional random seed, for reproducible fold assignment.

Value

A list containing cross-validation results:

- `folds`: a list (one per fold) with the fitted params, `train_sse` (in-sample), and `test_sse` (out-of-sample).
- `params`: a matrix of fitted parameters, one row per fold.
- `train_sse / test_sse`: numeric vectors of per-fold SSE.

Examples

```
# A reduced DEoptim search keeps this example fast; drop
# deoptim_control for a real, thorough fit.
cross_validated_group_fits(decay(C = 0.98), xsl_datasets[1:6],
  lower = 0.8, upper = 1.0, n_folds = 2,
  deoptim_control = DEoptim::DEoptim.control(NP = 10, itermax = 5))
```

decay

Decay model

Description

Decay model

Usage

```
decay(C)
```

Arguments

C decay

Value

An object of class xslMod

Examples

```
mod <- decay(C = 0.5)
xsl_run(mod, get_example_ambiguous_condition())
```

fazly

Fazly et al. 2010 probabilistic associative model

Description

Fazly et al. 2010 probabilistic associative model

Usage

```
fazly(lambda, beta, alpha = 10, epsilon = 0.001, theta = NULL)
```

Arguments

lambda	Small smoothing factor (e.g. 1e-5)
beta	Upper bound on number of symbol types to expect (e.g. 8500)
alpha	Fixed smoothing parameter (e.g. 10)
epsilon	Fixed smoothing parameter (e.g. 0.001)
theta	Threshold to determine when associations enter the lexicon (defaults to NULL, i.e. model variant without thresholding).

Value

An object of class `xslMod`

Examples

```
mod <- fazly(lambda = 1e-5, beta = 8500)
xsl_run(mod, get_example_ambiguous_condition())

mod_t <- fazly(lambda = 1e-5, beta = 8500, theta = 0.7)
xsl_run(mod_t, get_example_ambiguous_condition())
```

fgt2009

Frank, Goodman & Tenenbaum (2009) intentional Bayesian word-learning model

Description

A joint Bayesian model that infers a *lexicon* – a set of word-object edges – from the whole corpus at once, marginalizing over the speaker’s latent referential intention on each trial. Words are generated either referentially (probability γ , naming an object in the intention) or non-referentially ($1 - \gamma$, with in-lexicon words down-weighted by κ); a geometric prior $P(L)$ proportional to $\exp(-\alpha |L|)$ favors small lexicons. Inference is MCMC (blocked per-object Gibbs plus MI-biased single-edge Metropolis moves), and `xslFit$matrix` holds the resulting posterior edge marginals $P(\text{word names object} \mid \text{corpus})$.

Usage

```
fgt2009(
  alpha,
  gamma = 1,
  kappa = 0.5,
  n_chains = 4L,
  n_warmup = 120L,
  n_samples = 300L,
  gibbs_per = 1L,
  edge_per = 1L,
  top_k = 6L,
  seed = 1L
)
```

Arguments

<code>alpha</code>	Geometric lexicon-size prior: $P(L)$ proportional to $\exp(-\alpha L)$. Must be tuned to corpus size.
<code>gamma</code>	Probability that a word is used referentially. 1 (the default) is appropriate when every word heard is a label, as in controlled XSL experiments.
<code>kappa</code>	Down-weight applied to an in-lexicon word used non-referentially (only relevant when $\gamma < 1$).

n_chains, n_warmup, n_samples, gibbs_per, edge_per, top_k, seed

Sampler controls. Defaults are tuned for the small, dense corpora of XSL experiments; raise n_chains/n_samples to reduce Monte Carlo noise.

Details

This is a **batch** model, unlike the incremental associative models in this package: it has no trial-by-trial learning trajectory, and `control$reps` is ignored. Ported from the Python package `wordlearn` (github.com/mcfrank/wurwur); see also [fgt2009_rsa\(\)](#) for the pragmatic extension.

`alpha` scales with corpus size (larger corpora want a larger `alpha`) and the fit-vs-`alpha` curve is single-peaked, so `alpha` is best chosen with a sweep (see [fgt2009_sweep_alpha\(\)](#)) rather than a general-purpose optimizer. `gamma` and `kappa` default to the values used for the analogous cross-situational simulation in the source package (`gamma = 1`: every word in a controlled XSL experiment is a referential label, so there is no non-referential channel; `kappa = 0.5`). The naturalistic-speech values are `gamma = 0.1`, `kappa = 0.05`.

Value

An object of class `xslMod`

Examples

```
mod <- fgt2009(alpha = 1)
xsl_run(mod, get_example_ambiguous_condition())
```

fgt2009_rsa

Frank, Goodman & Tenenbaum word-learning model with pragmatic (RSA) reasoning

Description

[fgt2009\(\)](#) with the literal speaker replaced by a Rational Speech Act pragmatic speaker: a speaker who chooses words to be informative to a listener who reasons about that speaker (Frank & Goodman, 2012; Smith, Goodman & Frank, 2013). The non-referential channel, the `kappa` penalty, and the marginalization over intentions are unchanged, so the literal [fgt2009\(\)](#) model is the `rsa_depth = 0` special case. The recursion is what produces *strong* mutual exclusivity – the speaker could have used a different word – which the literal model underproduces.

Usage

```
fgt2009_rsa(
  alpha,
  gamma = 1,
  kappa = 0.5,
  rsa_alpha = 3,
  rsa_depth = 1L,
  n_chains = 4L,
  n_warmup = 120L,
```

```

    n_samples = 300L,
    gibbs_per = 1L,
    edge_per = 1L,
    top_k = 6L,
    seed = 1L
  )

```

Arguments

alpha	Geometric lexicon-size prior: $P(L)$ proportional to $\exp(-\alpha L)$. Must be tuned to corpus size.
gamma	Probability that a word is used referentially. 1 (the default) is appropriate when every word heard is a label, as in controlled XSL experiments.
kappa	Down-weight applied to an in-lexicon word used non-referentially (only relevant when $\gamma < 1$).
rsa_alpha	Speaker rationality (softmax temperature) in the RSA recursion.
rsa_depth	Number of pragmatic recursions; 1 gives the usual one-step pragmatic listener/speaker.
n_chains, n_warmup, n_samples, gibbs_per, edge_per, top_k, seed	Sampler controls. Defaults are tuned for the small, dense corpora of XSL experiments; raise n_chains/n_samples to reduce Monte Carlo noise.

Value

An object of class `xslMod`

Examples

```

mod <- fgt2009_rsa(alpha = 1, rsa_alpha = 3)
xsl_run(mod, get_example_ambiguous_condition())

```

fgt2009_sweep_alpha *Sweep the lexicon-size prior alpha for an FGT model*

Description

`fgt2009()`'s alpha trades lexicon size against fit and must scale with corpus size, and the fit-vs-alpha curve is single-peaked – so a coarse sweep is the right way to choose it (and to evaluate the model fairly), rather than a general-purpose optimizer. This runs `xsl_run()` for each alpha and returns the per-value SSE against the dataset's human accuracy.

Usage

```
fgt2009_sweep_alpha(
  data,
  alphas,
  rsa = FALSE,
  gold = NULL,
  ...,
  control = xslControl()
)
```

Arguments

<code>data</code>	An <code>xslData</code> object (or list of them).
<code>alphas</code>	Numeric vector of alpha values to try.
<code>rsa</code>	Logical; if TRUE, sweep fgt2009_rsa() instead of fgt2009() .
<code>gold</code>	Optional gold-standard lexicon as <code>list(words, objects)</code> . When given, the result gains an <code>f_max</code> column (best-threshold F-score against gold). <code>data</code> must be a single <code>xslData</code> in this case.
<code>...</code>	Further arguments to the model constructor (e.g. <code>gamma</code> , <code>kappa</code> , sampler controls).
<code>control</code>	Control arguments passed to <code>xsl_run()</code> .

Details

For a naturalistic corpus (e.g. [rollins_corpus](#), [fm_corpus](#)) there is no human accuracy vector; pass `gold` (a `list(words, objects)` lexicon) to also get the best-threshold F-score of the learned matrix against it (via [get_roc_max\(\)](#)) at each alpha. Remember such corpora also want the naturalistic-speech `gamma = 0.1`, `kappa = 0.05` rather than the defaults (see [fgt2009\(\)](#)).

Value

A data frame with columns `alpha` and `sse` (plus `unweighted_sse` when `data` is a list, or `f_max` when `gold` is given).

Examples

```
fgt2009_sweep_alpha(get_example_ambiguous_condition(),
  alphas = c(1, 2, 4, 8))
```

fm_corpus

*Frank, Tenenbaum & Fernald naturalistic word-learning corpus***Description**

4763 caregiver utterances (24 mother-infant sessions), each paired with the objects present in the scene. 1122 word types, 30 object types present in scenes; roughly half the utterances are non-referential. Unlike [rollins_corpus](#), the speaker's referential intention is hand-coded per utterance.

Usage

```
fm_corpus
```

Format

An object of class `list` of length 5.

Details

A list with:

`data` an [xslData](#) object with the 4763 training utterances (character-vector words/objects per trial). No accuracy/test.

`intents` a length-4763 list; `intents[[t]]` is the object(s) the speaker referred to on utterance `t` (character vector, empty for non-referential utterances). A per-utterance gold signal.

`gold` a hand-curated gold lexicon as `list(words, objects)` (41 pairs), for [get_fscore\(\)](#) / [get_roc\(\)](#) / [get_tp\(\)](#).

`gold_variants` `list(strict, permissive)` – two auto-derived alternatives (39 and 116 pairs) from looser vs. tighter co-occurrence thresholds on the coded intentions.

`reference` the citation string.

Not part of [xsl_datasets](#). Gaze/hand attentional cues present in the source CSVs are not imported.

Imported from the wurwur package (<https://github.com/mcfrank/wurwur>); see `data-raw/add_wurwur_corpora.R`.

Source

Frank, M. C., Tenenbaum, J. B., & Fernald, A. (2013). Social and discourse contributions to the determination of reference in cross-situational word learning. *Language Learning and Development*, 9(1), 1-24.

Examples

```
length(fm_corpus$data$train$words)
fm_corpus$intents[[1]]
fm_corpus$gold_variants$strict$words
```

gangwani2011_category *Gangwani, Kachergis & Yu simultaneous category + object name learning*

Description

Experiment 1 of Gangwani, Kachergis & Yu, "Simultaneous Cross-situational Learning of Category and Object Names" (an undergraduate-journal version is Gangwani & Kachergis, *Indiana Undergraduate Journal of Cognitive Science* 5, 2010). A **hierarchical** cross-situational design: on every training trial the learner sees 2 objects and hears 3 words – the two objects' own 1-to-1 names *plus* one 1-to-many "category" label that names those two objects and two further objects seen on other trials. Every object is thus part of a 2-to-1 word-referent mapping (a basic name and a category label), and the paper's finding is that learners violate mutual exclusivity: they learn *both* labels for the same object.

Usage

```
gangwani2011_category
```

Format

An object of class `list` of length 3.

Details

A named list of three [xslData](#) objects, one per training block – "block 1 (ad hoc)", "block 2 (natural)", "block 3 (ad hoc)" – run in that fixed order. All three share the same 36-trial co-occurrence schedule (only the stimuli differ); each carries its own block-specific human data.

Per object: `train$words[[t]]` is `c(name_i, name_j, category_label)` and `train$objects[[t]]` is `c(i, j)`. Words 1-12 are object names (name *i* pairs with object *i*); words 13-15 are the category labels (13/14/15 name objects 1-4 / 5-8 / 9-12). `accuracy` is per-word human $P(\text{correct referent} | \text{word})$, length 15: positions 1-12 are 12-AFC name accuracy, positions 13-15 are 3-AFC category-membership accuracy ($n = 34$ canonical-block-order participants; the paper reports 33, dropping one more low performer). `response_matrix` (15 x 12, row-normalised) is the full human choice distribution including the category labels. `test` holds only the 12 name trials (12-AFC).

Not part of [xsl_datasets](#), and not scorable by the built-in `xsl_run()` SSE: the association matrix is 15 words x 12 objects, so `get_perf()` / `mafc_test()` (which score via the diagonal `m[w, w]`) cover only words 1-12. Score the category labels – and the ME-violation measure that is the paper's point – with `score_gangwani2011_category()` in `tests/bakeoff_comparison/gangwani2011_category_fit.R`. The natural-category feature-similarity effect (block 2) and generalization to novel objects are out of scope for the package's amodal, index-based models and are not represented. See `data-raw/add_gangwani2011_category` for construction details.

Examples

```

block1 <- gangwani2011_category[["block 1 (ad hoc)"]]
# 15 x 12 matrix; suppressWarnings() because get_perf()'s diagonal rule
# warns on the non-square matrix (the built-in SSE is not meaningful here).
m <- suppressWarnings(xsl_run(baseline(), block1))$fits[[1]]$matrix
mafc_test(m, block1$test) # names (words 1-12) only

```

```
get_crossvalidated_model_fit
```

Get cross-validated model fit

Description

Performs k-fold cross-validation on a model using the available datasets: the datasets are split into `n_folds` groups, and for each fold the model is fit (via `xsl_fit()`) to the other folds and evaluated (via `xsl_run()`) on the held-out fold.

Usage

```

get_crossvalidated_model_fit(
  model_name,
  n_folds = 5,
  datasets = NULL,
  control = xslControl(),
  deoptim_control = DEoptim::DEoptim.control(reltol = 0.001, NP = 100, itermax = 100)
)

```

Arguments

<code>model_name</code>	Name of the model to cross-validate (see <code>show_models()</code>)
<code>n_folds</code>	Number of folds for cross-validation (default: 5)
<code>datasets</code>	Optional list of datasets to use (defaults to all available)
<code>control</code>	Control arguments passed to <code>xsl_run()/xsl_fit()</code> ; notably <code>n_sim</code> , the number of simulations averaged per evaluation of a stochastic model (default 500 – lowering this can speed up cross-validation of stochastic models substantially).
<code>deoptim_control</code>	Control parameters passed to <code>DEoptim()</code> (via <code>cross_validated_group_fits()</code>).

Value

A list containing cross-validation results (see `cross_validated_group_fits()`)

Examples

```
# Get cross-validated fit for the decay model. A reduced DEoptim search
# keeps this example fast; drop deoptim_control for a real, thorough fit.
cv_fit <- get_crossvalidated_model_fit(
  "decay", n_folds = 2, datasets = xsl_datasets[1:6],
  deoptim_control = DEoptim::DEoptim.control(NP = 10, itermax = 5))
```

```
get_example_ambiguous_condition
```

Get example ambiguous condition

Description

This function creates an example condition with two trials, and two word-object pairs per trial. Each word-object pair appears only once, and thus each word has a 50% chance of being associated with the correct referent. The function returns a list of the trials with nested words and objects per trial, as well as a vector of the conditional probability of correctly selecting the correct referent given each word.

Usage

```
get_example_ambiguous_condition()
```

Value

An object of class `xslData` containing a list of training trials (with nested words and objects per trial) and accuracy ($P(\text{correct referent} \mid \text{word})$).

Examples

```
get_example_ambiguous_condition()
```

```
get_example_unambiguous_condition
```

Get example unambiguous condition

Description

This function creates an example condition with three trials and two word-object pairs per trial. Each word-object pair appears twice, and so if the most frequently co-occurring referent for each word is selected then all three words would be learned.

Usage

```
get_example_unambiguous_condition()
```

Value

An object of class `xslData` containing a list of training trials (with nested words and objects per trial) and accuracy ($P(\text{correct referent} \mid \text{word})$).

Examples

```
get_example_unambiguous_condition()
```

<code>get_fscore</code>	<i>Calculate F-score, precision, recall, and specificity for a knowledge matrix at a given threshold</i>
-------------------------	--

Description

This function calculates the F-score, precision, recall, and specificity for a given knowledge matrix at a specified threshold. It uses the concept of true positives, false positives, and false negatives, determined from the knowledge matrix and an optional gold lexicon. The function is useful for evaluating the performance of a model in terms of its ability to correctly identify associations between words and referents.

Usage

```
get_fscore(m, threshold, gold_lexicon = NULL)
```

Arguments

<code>m</code>	A matrix representing the knowledge matrix with words as rows and referents as columns.
<code>threshold</code>	A numeric value representing the threshold for considering an association between a word and a referent as positive.
<code>gold_lexicon</code>	Optional; a data frame or list where each row/element represents a word-object pair in the gold lexicon. If provided, it is used to calculate true positives, false positives, and false negatives.

Value

A tibble with columns for threshold, precision, recall, specificity, and F-score.

Examples

```
dat <- xsl_datasets[[1]]
x <- xsl_run(baseline(), dat)
mat <- x$fits[[1]]$matrix
lex <- list(words = rep(1:18), objects = rep(1:18))
get_fscore(mat, 0.5, lex)
```

get_group_model_fit	<i>Get group model fit for a specific model</i>
---------------------	---

Description

Fits a model to all available datasets and returns the group-level fit. This function provides a convenient way to get pre-computed model fits for common models.

Usage

```
get_group_model_fit(
  model_name,
  datasets = NULL,
  control = xslControl(),
  deoptim_control = DEoptim::DEoptim.control(reltol = 0.001, NP = 100, itermax = 100)
)
```

Arguments

model_name	Name of the model to fit (see show_models())
datasets	Optional list of datasets to fit to (defaults to all available)
control	Control arguments passed to xsl_run() (via xsl_fit()); notably n_sim, the number of simulations averaged per evaluation of a stochastic model (default 500 – lowering this can speed up fitting of stochastic models substantially).
deoptim_control	Control parameters passed to DEoptim() (via xsl_fit()).

Value

A list containing the fitted model results

Examples

```
# Get group fit for the uncjam model. A small dataset subset and a
# reduced DEoptim search keep this example fast; drop those arguments
# for a real, thorough fit.
group_fit <- get_group_model_fit(
  "uncjam", datasets = xsl_datasets[1:3],
  deoptim_control = DEoptim::DEoptim.control(NP = 10, itermax = 5))
```

get_perf	<i>Calculate Luce choice (proportion correct) for each item in a model knowledge matrix</i>
----------	---

Description

This function computes the Luce choice, or the proportion of correct selections, for each item in a given model knowledge matrix. It assesses the probability of correctly identifying each referent based on the knowledge matrix, providing a measure of model performance per item.

Usage

```
get_perf(m, d = NULL)
```

Arguments

m	A square matrix representing the model's knowledge, where rows correspond to words and columns correspond to referents. The diagonal elements represent correct associations, and off-diagonal elements represent incorrect associations.
d	Exponent for exponentiated choice rule

Details

Negative associations (possible in error-driven models such as [rescorla_wagner\(\)](#) and [kalman_filter\(\)](#)) are treated as 0 before the choice rule is applied – Luce choice is undefined for negative values, and would otherwise yield "probabilities" outside [0, 1] – matching [predict_referent\(\)](#). A word whose associations are then all 0 is answered at chance ($1 / \text{ncol}(m)$).

Value

A named numeric vector where each element corresponds to an item in the matrix. The value of each element represents the proportion of correct selections (Luce choice) for that item, calculated as the ratio of the correct association (diagonal element) to the total associations for that item.

Examples

```
x <- xsl_run(baseline(), get_example_ambiguous_condition())
mat <- x$fits[[1]]$matrix
get_perf(mat)
```

get_roc	<i>Calculate receiver operating characteristic (ROC) scores for a model association Matrix</i>
---------	--

Description

This function computes receiver operating characteristic (ROC) scores for a given model association matrix. It evaluates the performance of the model at various thresholds, providing metrics like f-scores, precision, and recall. The function can operate with a range of thresholds and optionally consider a gold lexicon for calculating true and false positives and negatives. The result is a comprehensive assessment of model performance over a continuum of classification thresholds.

This function computes the maximum F-score from the receiver operating characteristic (ROC) scores of a model. It leverages the `get_roc()` function to calculate the ROC scores and then extracts the highest F-score, providing a concise metric for the best classification performance of the model.

Usage

```
get_roc(m, thresholds = seq(0, 1, 0.01), gold_lexicon = NULL)

plot_roc(m, thresholds = seq(0, 1, 0.01), gold_lexicon = NULL)

get_roc_max(m, thresholds = seq(0, 1, 0.01), gold_lexicon = NULL)
```

Arguments

<code>m</code>	A matrix representing the knowledge matrix with words as rows and referents as columns.
<code>thresholds</code>	Vector of thresholds to use.
<code>gold_lexicon</code>	Optional; a data frame or list where each row/element represents a word-object pair in the gold lexicon. If provided, it is used to calculate true positives, false positives, and false negatives.

Value

A tibble with columns for the threshold, f-score, precision, and recall.

A single numeric value representing the maximum F-score obtained from the ROC scores of the model.

Examples

```
dat <- xsl_datasets[[1]]
x <- xsl_run(baseline(), dat)
mat <- x$fits[[1]]$matrix
lex <- list(words = rep(1:18), objects = rep(1:18))
get_roc(mat, gold_lexicon = lex)
plot_roc(mat, gold_lexicon = lex)
get_roc_max(mat, gold_lexicon = lex)
```

get_tp	<i>Get true positives (TP), given a knowledge matrix and a gold-standard lexicon</i>
--------	--

Description

This function iterates over words in a given gold lexicon and accumulates the associative strength (can be integral e.g. 1, or real-valued) in the knowledge matrix for the intended referents (present in the gold lexicon). Returns the number of expected true positives (TP) for this gold lexicon and knowledge matrix.

Usage

```
get_tp(m, gold_lexicon)
```

Arguments

m	A matrix representing the knowledge matrix with words as rows and referents as columns.
gold_lexicon	A data frame or list where each row/element represents a word-object pair in the gold lexicon.

Value

A single value with the expected number of true positives.

Examples

```
dat <- xsl_datasets[[10]]
x <- xsl_run(baseline(), dat)
mat <- x$fits[[1]]$matrix
lex <- list(words = rep(1:18), objects = rep(1:18))
get_tp(mat, lex)
```

get_train_test_split	<i>Split Data into Training and Testing Sets Based on Indices</i>
----------------------	---

Description

This function divides a given set of conditions into training and testing sets. The division is based on a provided vector of indices, which specify the conditions to be used for testing.

Usage

```
get_train_test_split(test_inds, conds)
```

Arguments

test_inds	A numeric vector of indices indicating which elements in conds should be used for the test set.
conds	A list of conditions (e.g. xslData objects), where each condition is represented as a separate element.

Value

A list with two elements: train and test. Each element is a list of conditions, where train includes the conditions not indexed by test_inds, and test includes the conditions indexed by test_inds.

Examples

```
split <- get_train_test_split(c(1, 3), xsl_datasets[1:5])
length(split$train)
length(split$test)
```

guess_and_test

Guess and test model

Description

Trueswell & Gleitman 2011 guess-and-test model

Usage

```
guess_and_test(f, sa)
```

Arguments

f	forget at retrieval
sa	prob of storage (slow learning down)

Value

An object of class xslMod

Examples

```
mod <- guess_and_test(f = .1, sa = .5)
xsl_run(mod, get_example_ambiguous_condition())
```

kachergis2012_highlighting

Kachergis (2012) highlighting experiment data

Description

A list of two `xslData` objects ("words as cues" and "objects as cues") from the highlighting experiment in Kachergis (2012, CogSci), "Learning Nouns with Domain-General Associative Learning Mechanisms". Training trials (genuine simultaneous 2-cue trials, matching the paper's Figure 2: 2 words + 1 object, or 2 objects + 1 word, per trial) are read directly from the real per-trial ordering file for the real N=67 dataset behind the paper's published statistics – confirmed by reproducing its exact reported response proportions and chi-square values. See `data-raw/add_kachergis2012_highlighting.R` for the full construction rationale, including index relabeling needed to work around a `uncfam_model()` indexing limitation.

Usage

`kachergis2012_highlighting`

Format

An object of class `list` of length 2.

Details

Each condition replicates the classic 3-role highlighting structure (PE, PL, I) twice. Words-as-cues' PL items have their real target at a different index than their own word index, and its I items have two legitimate targets – neither fits `xslData`'s one-correct-object-per-word accuracy convention, so only its two PE items get a real accuracy value (the rest are NA). Objects-as-cues has no such issue (every tested word's own index is among its trained objects after relabeling), so all 4 of its items have real accuracy. Because of the partial-NA accuracy vector, this dataset is kept standalone rather than appended to `xsl_datasets`: an NA would silently poison every other dataset's aggregate fit in `get_group_model_fit()/get_crossvalidated_model_fit()`, which sum SSE across all of `xsl_datasets`. See `tests/bakeoff_comparison/kachergis2012_highlighting_fit.R` for a custom scorer that fits a model against all four of the paper's reported proportions (PE-E, PL-L, I-E, I-L) directly off the association matrix, the way the paper's own (non-package) fitting procedure did.

kachergis_initial_accuracy

Kachergis, Grimmick, & Gureckis initial accuracy experiment data

Description

A list of two `xslData` objects ("High Initial Accuracy" and "Low Initial Accuracy") from an unpublished MTurk cross-situational word learning experiment manipulating how many of 18 words are first shown ("familiarized") with a different object than the one they are then cross-situationally studied with. The study pairing is the test-correct one, and is on the diagonal ($m[w, w]$); an initially inaccurate word's familiarization object is off-diagonal. The first 18 training trials are the familiarization phase. Not included in `xsl_datasets` (see `data-raw/add_kachergis_initial_accuracy.R` for construction details and `tests/bakeoff_comparison/kachergis_initial_accuracy_fit.R` for an example model comparison).

Usage

`kachergis_initial_accuracy`

Format

An object of class `list` of length 2.

<code>kalman_filter</code>	<i>Kalman filter associative model</i>
----------------------------	--

Description

A Kalman-filter generalization of `rescorla_wagner()`'s error-driven update rule (Dayan & Kakade, 2000; Kruschke, 2008; Gershman, 2015): each word-object association is tracked as a Gaussian belief (mean μ and variance σ^2), and each trial's update is scaled by that pair's own Kalman gain, $\sigma^2 / (\sigma^2 + \sigma^2_{\text{obs}})$. This makes the effective learning rate adaptive rather than fixed – large (fast learning) while an association is still uncertain, and automatically shrinking as confidence accumulates – and it emerges from three interpretable parameters rather than being hand-set. Between observations, every association's variance grows by τ^2 (the model assumes true associations drift slowly over time), which keeps the model able to revise a belief rather than converging to a fixed point.

Usage

`kalman_filter(tau2, sigma2_obs, sigma2_0)`

Arguments

<code>tau2</code>	Process (diffusion) noise: how much uncertainty about every association grows per trial, whether or not it was observed
<code>sigma2_obs</code>	Observation noise: how uninformative a single co-occurrence observation is (larger = slower learning per trial)
<code>sigma2_0</code>	Initial (prior) uncertainty about every association, before any training

Details

Note that the three parameters share a scale redundancy: multiplying `tau2`, `sigma2_obs` and `sigma2_0` together by any constant leaves every Kalman gain – and hence the learned matrix – unchanged. Only two of the three are identifiable; when fitting, hold one fixed (conventionally `sigma2_obs`).

Value

An object of class `xslMod`

Examples

```
mod <- kalman_filter(tau2 = .01, sigma2_obs = .5, sigma2_0 = 1)
xsl_run(mod, get_example_ambiguous_condition())
```

mafc_test

Evaluate m-alternative forced choice test

Description

Scores a set of test trials against a model's word-object matrix, returning the probability of choosing the correct object on each trial. Each trial presents one word and a set of candidate referents; the correct object is the one whose id matches the word's id (the package's diagonal convention). A thin wrapper over [predict_referent\(\)](#).

Usage

```
mafc_test(m, test, ...)
```

Arguments

<code>m</code>	A matrix representing model memory with words as rows and referents as columns.
<code>test</code>	A list representing the test trials, each containing a word and its associated referents.
<code>...</code>	Further arguments passed to predict_referent() (e.g. <code>pragmatic</code> , <code>threshold</code>).

Value

A vector with the probability of choosing the correct object, given each word.

Examples

```
dat <- xsl_datasets[[10]]
x <- xsl_run(baseline(), dat)
mat <- x$fits[[1]]$matrix
mafc_test(mat, dat$test)
```

minerva2

*MINERVA2 episodic memory model***Description**

Hintzman's (1984, 1986) MINERVA2 adapted to cross-situational word learning: each trial is stored as a single episodic trace formed by concatenating that trial's word-vector sum and object-vector sum (words and objects are random Gaussian vectors, one per vocabulary item). At test, a word probe (word-half only, object-half zeroed) activates every trace by its similarity cubed; the activation-weighted sum of traces (the "echo") is compared, on its object half, against every candidate object to give a graded association – exponentiated by X into a Luce-choice-ready matrix rather than MINERVA2's usual hard best-match.

Usage

```
minerva2(X, D = 100)
```

Arguments

X	Choice-rule temperature: how sharply the echo/object cosine similarity is exponentiated into a choice weight (larger = more winner-take-all)
D	Dimensionality of the random word/object feature vectors (fixed hyperparameter, not a free cognitive parameter)

Details

Performance note: memory is exact (no forgetting/decay), so the model keeps one $2 \times D$ -length trace per trial and a test pass costs $O(\text{voc_sz} * n_traces * D)$. D is a fixed representational hyperparameter (not usually treated as a free cognitive parameter), not included by default in this package's DEoptim registry bounds – see `xsl_model_registry()`.

Value

An object of class `xslMod`

Examples

```
mod <- minerva2(X = 5, D = 50)
xsl_run(mod, get_example_ambiguous_condition())
```

multi_sampling	<i>Multi-sampling associative model</i>
----------------	---

Description

Multi-sampling associative model

Usage

```
multi_sampling(C, X, B, K = 1)
```

Arguments

C	Decay
X	Total associative weight to distribute per trial
B	Associative weight to distribute each sample
K	Number of associations to update per word (constraint: $B \leq X$)

Value

An object of class xslMod

Examples

```
mod <- multi_sampling(C = 1, X = .2, B = .3, K = 1)
xsl_run(mod, get_example_ambiguous_condition())
```

new_xslFit	<i>Constructor for xsl_model S3 class</i>
------------	---

Description

Constructor for xsl_model S3 class

xslMod S3 class

Constructor for xslMod S3 class

Usage

```

new_xslFit(x = list())

xslMod(
  name = character(),
  description = character(),
  model,
  params = numeric(),
  stochastic = logical(),
  supports_start_matrix = FALSE
)

new_xslMod(x = list())

```

Arguments

x	List with elements perf, matrix, traj, sse
name	Name
description	Description
model	Model fitting function
params	List of parameters
stochastic	Logical indicating whether model is stochastic
supports_start_matrix	Logical: does model initialize from control\$start_matrix? xsl_run() errors if a start matrix is supplied to a model that doesn't, rather than letting it be silently ignored.

Value

An object of class xslMod

plot_training_trials *Create Animated Plot of Co-Occurrence Matrix from Training Trials*

Description

This function takes a sequence of training trials, each consisting of words and objects, and creates an animated plot of the word-object co-occurrence matrix. The plot evolves over each trial, visually representing the accumulation of co-occurrences. The animation is saved to a file if a filename is provided.

Usage

```
plot_training_trials(train, filename = NULL)
```

Arguments

train	A list representing the training data, where each element is a trial containing lists of words (words) and objects (objects).
filename	A string specifying the filename where the animation should be saved. If NULL, no file will be saved.

Value

A long-format data frame where each row corresponds to a trial and includes columns for trial number, word, object, and the count of co-occurrences up to that trial.

Examples

```
plot_training_trials(xsl_datasets[[1]]$train)
```

predict_referent	<i>Predict referent selection for a single trial from an association matrix</i>
------------------	---

Description

Given a model's word-object association matrix, returns the probability distribution over a trial's candidate objects for one heard word – the quantity to compare against a participant's choice on a referent-selection (m-alternative forced choice) trial. This is model-agnostic: it works on the matrix of any xslFit (or any word-by-object matrix), reading each entry as an association strength / meaning weight.

Usage

```
predict_referent(
  m,
  word,
  objects,
  prior = NULL,
  pragmatic = FALSE,
  threshold = NULL,
  rsa_alpha = 3,
  depth = 1
)
```

Arguments

m	A word-by-object association matrix (e.g. xsl_run(mod, data)\$fits[[1]]\$matrix). Assumed non-negative; negative entries are clamped to 0. dimnames are used to resolve character word/objects.
word	The heard word: a row name of m, or a positive integer row index. A value not present in m is treated as a novel word.

objects	The candidate objects present on the trial: column names of <code>m</code> , or positive integer column indices. An entry not present in <code>m</code> is treated as an unnamed (novel) object.
prior	Optional prior over objects (length <code>length(objects)</code> , non-negative). Defaults to uniform. Set an entry to 0 to remove that object from consideration entirely (e.g. an object the speaker has no epistemic access to).
pragmatic	If TRUE, use the RSA pragmatic listener instead of the literal one.
threshold	Optional; binarize <code>m</code> at this value (<code>m >= threshold</code>) before predicting. The RSA layer is sharpest on a binary lexicon.
rsa_alpha, depth	Speaker rationality and recursion depth for the RSA listener (only used when <code>pragmatic = TRUE</code>).

Details

The literal rule is Bayes' rule on the word's row, $P(\text{object} \mid \text{word})$ proportional to $m[\text{word}, \text{object}] * \text{prior}(\text{object})$ normalized over the objects present. With `pragmatic = TRUE` the objects are instead resolved by a one-step Rational Speech Act pragmatic listener (`rsa_listener()`), which additionally reasons that the speaker could have used a different word – this is what yields *strong* mutual exclusivity when a heard word is lexically ambiguous and a competitor word names one of the candidates. A truly novel (unseen) word carries no lexical evidence and returns the prior in either mode.

Value

A numeric vector of probabilities over objects, in the given order, summing to 1.

Examples

```
m <- xsl_run(fgt2009(alpha = 1), get_example_unambiguous_condition())$fits[[1]]$matrix
predict_referent(m, 1, c(1, 2, 3))
predict_referent(m, 1, c(1, 2, 3), prior = c(1, 0, 1)) # object 2 unavailable
```

propose_but_verify	<i>Trueswell et al. 2013 propose-but-verify model</i>
--------------------	---

Description

Trueswell et al. 2013 propose-but-verify model

Usage

```
propose_but_verify(alpha, alpha_increase)
```

Arguments

`alpha` initial association strength
`alpha_increase` learning rate

Value

An object of class `xslMod`

Examples

```
mod <- propose_but_verify(alpha = .1, alpha_increase = .5)
xsl_run(mod, get_example_ambiguous_condition())
```

pursuit

Stevens et al. 2014 pursuit model

Description

Stevens et al. 2014 pursuit model

Usage

```
pursuit(gamma, threshold, lambda)
```

Arguments

<code>gamma</code>	learning rate
<code>threshold</code>	threshold (probability) to move an association to the known lexicon
<code>lambda</code>	smoothing probability

Value

An object of class `xslMod`

Examples

```
mod <- pursuit(gamma = .2, threshold = .3, lambda = .05)
xsl_run(mod, get_example_ambiguous_condition())
```

rem

*REM (Retrieving Effectively from Memory) model***Description**

Shiffrin & Steyvers' (1997) REM adapted to cross-situational word learning. Word/object identity is a length- w feature vector drawn from a geometric base-rate distribution (g); each trial stores one episodic trace per candidate word-object pairing present, built by imperfectly copying the pair's true features (each feature independently stored with probability u , and if stored, correct with probability c , else a fresh random draw). At test, a candidate pairing's strength is the summed likelihood ratio of every trace given the true (fully-known) probe features, REM's odds-based alternative to `minerva2()`'s dot-product activation.

Usage

```
rem(g, u, c, w = 12)
```

Arguments

g	Base-rate feature-value distribution parameter (larger = more variable/distinctive feature values, in (0, 1))
u	Per-feature storage probability at study, in (0, 1)
c	Probability a stored feature is copied correctly (vs. replaced by a random, possibly-confusable value), in (0, 1)
w	Number of features per item (fixed hyperparameter, not a free cognitive parameter)

Details

Performance note: unlike `minerva2()` (one trace per trial) this model stores one trace per candidate (word, object) pair per trial – on a highly ambiguous corpus (many words/objects per utterance) memory size grows combinatorially rather than linearly with trial count, and both encoding and the $\text{voc_sz} * \text{ref_sz}$ test pass scale with that trace count, making `rem()` the most expensive of this package's memory models on cluttered/naturalistic data. w is a fixed representational hyperparameter (not usually treated as a free cognitive parameter), not included by default in this package's DEoptim registry bounds – see `xsl_model_registry()`.

Value

An object of class `xslMod`

Examples

```
mod <- rem(g = 0.4, u = 0.3, c = 0.7, w = 12)
xsl_run(mod, get_example_ambiguous_condition())
```

rescorla_wagner	<i>Rescorla-Wagner (1972) error-driven associative model</i>
-----------------	--

Description

Rescorla-Wagner (1972) error-driven associative model

Usage

```
rescorla_wagner(C, alpha, beta, lambda)
```

Arguments

C	Decay
alpha	Salience (fix at 1 unless manipulated)
beta	Learning rate – a proportion of lambda
lambda	Maximum associative value that a CS can achieve - should be larger than learning rate

Value

An object of class xslMod

Examples

```
mod <- rescorla_wagner(C = 1, alpha = 1, beta = 0.1, lambda = 3)
xsl_run(mod, get_example_ambiguous_condition())
```

rollins_corpus	<i>CHILDES/Rollins naturalistic word-learning corpus (Frank et al. 2009)</i>
----------------	--

Description

The corpus the intentional Bayesian model of Frank, Goodman & Tenenbaum (2009) was fit to: 619 mother-to-infant utterances from the Rollins corpus in CHILDES, each paired with the set of objects present in the scene (six toys rotated in groups). 416 word types, 22 object types.

Usage

```
rollins_corpus
```

Format

An object of class list of length 3.

Details

A list with:

`data` an `xslData` object with the 619 training utterances (`data$train$words[[t]]/data$train$objects[[t]]` are character vectors). No accuracy or test – there is no human referent-selection data for this corpus.

`gold` the gold-standard lexicon as `list(words, objects)` (34 word-object pairs), for scoring a learned matrix with `get_fscore()`, `get_roc()`, `get_roc_max()`, or `get_tp()`.

`reference` the citation string.

Not part of `xsl_datasets` (which is scored by SSE against a human accuracy vector these corpora don't have). Imported from the `wurwur` package (<https://github.com/mcfrank/wurwur>); see `data-raw/add_wurwur_corpora.R`.

Source

Frank, M. C., Goodman, N. D., & Tenenbaum, J. B. (2009). Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science*, 20(5), 578-585. Corpus originally from Rollins (2003), CHILDES.

Examples

```
m <- suppressWarnings(
  xsl_run(baseline(), rollins_corpus$data)$fits[[1]]$matrix)
get_roc_max(m, gold_lexicon = rollins_corpus$gold)
```

rsa_listener

Rational Speech Act pragmatic listener and speaker

Description

A small, self-contained RSA recursion over a word-by-object meaning matrix (Frank & Goodman, 2012; Smith, Goodman & Frank, 2013), used by `predict_referent()` with `pragmatic = TRUE`. `rsa_listener()` returns $P(\text{object} \mid \text{word})$ for a listener who reasons about an informative speaker; `rsa_speaker()` returns $P(\text{word} \mid \text{object})$ for that speaker. Both operate on any non-negative (word $\times$ object) matrix, reading entries as literal meaning weights.

Usage

```
rsa_listener(lex, prior, ...)
```

```
rsa_speaker(lex, prior, ...)
```

Arguments

<code>lex</code>	A (word x object) matrix of literal meaning weights.
<code>prior</code>	Numeric prior over objects (one per column of <code>lex</code>). Objects with prior 0 are dropped from the reasoning and get listener probability 0.
<code>...</code>	Further arguments to the recursion: <code>alpha</code> (speaker rationality / softmax temperature, default 3), <code>cost</code> (optional per-word production cost vector), and <code>depth</code> (pragmatic recursions beyond the literal level; 0 gives the prior-weighted literal listener, 1 the usual one-step pragmatic one).

Value

`rsa_listener()`: a (word x object) matrix of $P(\text{object} \mid \text{word})$. `rsa_speaker()`: a (word x object) matrix of $P(\text{word} \mid \text{object})$.

Examples

```
# "wug" names either object; a pragmatic listener resolves it to the second
lex <- matrix(c(1, 0, 1, 1), 2, byrow = TRUE)
rsa_listener(lex, c(0.5, 0.5), alpha = 3, depth = 1)
```

show_datasets

Show available datasets in the package

Description

Returns information about the available cross-situational word learning datasets included in the package.

Usage

```
show_datasets()
```

Value

A data frame with dataset information

Examples

```
show_datasets()
```

show_models	<i>Show available models in the package</i>
-------------	---

Description

Returns a list of all available cross-situational word learning models that can be used with the package.

Usage

```
show_models()
```

Value

A character vector of model names

Examples

```
show_models()
```

softmax_rl	<i>Softmax reinforcement learning model</i>
------------	---

Description

A Q-learning-style reinforcement learning model of cross-situational word learning. Unlike every other model in this package, which updates every word-object pair presented together on a trial, this model makes a discrete choice: for each word heard, it samples a single candidate referent from a softmax policy over its current values for *all* known objects (its current best guess at that word's meaning), then updates only that guess's value via a standard TD/delta-rule update, using whether the guess was actually present in the scene as the reward (1 = "confirmed", 0 = "disconfirmed"). This is the only feedback available in a task with no trial-by-trial correct answer, and mirrors the propose-and-verify logic of [propose_but_verify\(\)](#)/[pursuit\(\)](#), but replaces their threshold-based keep/discard rules with graded value learning and an explicit (softmax) exploration/exploitation tradeoff.

Usage

```
softmax_rl(alpha, beta)
```

Arguments

alpha	Learning rate
beta	Inverse temperature: how deterministically the model favors its highest-valued guess (0 = uniformly random guessing, large = always picks the current best guess)

Value

An object of class xslMod

Examples

```
mod <- softmax_rl(alpha = .3, beta = 3)
xsl_run(mod, get_example_ambiguous_condition())
```

tilles

Tilles and Fontanari (2013) reinforcement model

Description

Tilles and Fontanari (2013) reinforcement model

Usage

```
tilles(x, b, alpha_0)
```

Arguments

x	Reinforcement parameter for stimuli in current context
b	Inference parameter, regulating ME and prior info integration (applies either to known words that do not appear in the current context or to new words in the current context)
alpha_0	Baseline efficiency corresponding to maximal uncertainty about referent of target word

Value

An object of class xslMod

Examples

```
mod <- tilles(x = 0.6, b = 0.8, alpha_0 = 0.85)
xsl_run(mod, get_example_ambiguous_condition())
```

todam

*TODAM holographic reduced representation model***Description**

Murdock's (1982) TODAM adapted to cross-situational word learning: word and object vectors (random Gaussian, one per vocabulary item) are bound via circular convolution into a single memory vector M , attention-weighted each trial by how familiar (cosine-similar to the current M) each word/object already is. At test, a candidate object is circularly correlated back out of M and compared (cosine) against each word; similarities are shifted non-negative, row-normalized, and exponentiated by X into a graded (Luce-choice-ready) matrix rather than a hard best-match.

Usage

```
todam(X, D = 500)
```

Arguments

X	Choice-rule temperature: how sharply the retrieved-object/word cosine similarity is exponentiated into a choice weight (larger = more winner-take-all)
D	Dimensionality of the random word/object feature vectors (fixed hyperparameter, not a free cognitive parameter)

Details

Performance note: every encoding step and every test comparison requires an FFT-based circular convolution/correlation over a D -length vector, so cost scales with $D \log(D)$; unlike `minerva2()`'s per-trial trace, `todam()` holds a single fixed-size memory vector regardless of corpus length. D is a fixed representational hyperparameter (not usually treated as a free cognitive parameter), not included by default in this package's DEoptim registry bounds – see `xsl_model_registry()`.

Value

An object of class `xslMod`

Examples

```
mod <- todam(X = 5, D = 200)
xsl_run(mod, get_example_ambiguous_condition())
```

uncfam*Kachergis 2012*

Description

Kachergis et al. 2012 uncertainty- and familiarity-biased associative model

Usage

```
uncfam(X, B, C, variant = c("entropy", "novelty", "uncertainty-only"))
```

Arguments

X	Associative weight to distribute
B	Weighting of uncertainty vs. familiarity
C	Decay
variant	Which variant of the model to fit (one of "entropy", "novelty", "uncertainty-only").

Value

An object of class xslMod

Examples

```
mod <- uncfam(X = .1, C = 1, B = .98)
xsl_run(mod, get_example_ambiguous_condition())

mod <- uncfam(X = .1, C = 1, B = .98, variant = "novelty")
xsl_run(mod, get_example_ambiguous_condition())

mod <- uncfam(X = .1, C = 1, B = 0) # familiarity-only
xsl_run(mod, get_example_ambiguous_condition())

mod <- uncfam(X = .1, C = 1, B = .98, variant = "uncertainty-only")
xsl_run(mod, get_example_ambiguous_condition())
```

uncfam_attention	<i>Biased associative model with attention scaled to trial uncertainty</i>
------------------	--

Description

A variant of `uncfam()` (Kachergis et al. 2012's uncertainty- and familiarity-biased associative model) in which the associative weight distributed on a trial is additionally scaled by the ratio of that trial's mean object entropy to the mean entropy of all objects seen so far. This implements the "system-level" theory proposed by Fitneva & Christiansen (2015): that learners are more alert (i.e. allocate a higher learning rate) on trials with more uncertain items, such as in a Low Initial Accuracy condition where many items have been mis-paired.

Usage

```
uncfam_attention(X, B, C)
```

Arguments

X	Associative weight to distribute
B	Weighting of uncertainty vs. familiarity
C	Decay

Value

An object of class `xslMod`

Examples

```
mod <- uncfam_attention(X = .1, C = 1, B = .98)
xsl_run(mod, get_example_ambiguous_condition())
```

uncfam_elimination	<i>Biased associative model with learning by elimination</i>
--------------------	--

Description

A variant of `uncfam()` (Kachergis et al. 2012's uncertainty- and familiarity-biased associative model) that adds inference by elimination (mutual exclusivity) to how a trial's associative weight is allocated. If the other words and objects on a trial are already confidently paired with each other, the remaining word and object probably go together, so their pairing should draw attention even though it is still weak. In `uncfam()`, by contrast, an already-strong pair on the trial draws weight *away* from such a pairing.

Usage

```
uncfam_elimination(X, B, C, eps)
```

Arguments

X	Associative weight to distribute
B	Weighting of uncertainty vs. familiarity
C	Decay
eps	Weight of elimination inference (0 = identical to <code>uncfam()</code>)

Details

On each trial, each on-trial pair (w', o') has a mutual confidence $q(w', o') = P(o'|w')P(w'|o')$, from the row- and column-normalized associations. A pair (w, o) gets elimination support $s(w, o) = \sum_{w' \neq w, o' \neq o} q(w', o')$: how much of the rest of the trial is accounted for by other pairs. The share of the trial's weight X given to (w, o) is then

$$\frac{b(w, o) + \epsilon s(w, o)}{1 + \epsilon \sum s}$$

where $b(w, o)$ is `uncfam()`'s normalized familiarity-and-uncertainty share, so `eps = 0` reproduces `uncfam()` exactly.

Motivated by the initial-accuracy experiment (`kachergis_initial_accuracy`), where learners were more accurate on an initially mis-paired word the more of its study trials it shared with a word they already knew.

Value

An object of class `xslMod`

Examples

```
mod <- uncfam_elimination(X = .1, B = .98, C = 1, eps = 1)
xsl_run(mod, get_example_ambiguous_condition())
```

uncfam_gamma

Kachergis 2012 with a free familiarity exponent

Description

A variant of `uncfam()` (Kachergis et al. 2012's uncertainty- and familiarity-biased associative model) that decouples familiarity's contribution to a trial's attentional allocation from uncertainty's. `uncfam()`'s allocation weight for pair (w, o) is $\exp(B * \text{entropy}) * \text{assocs} - \text{entropy}$ gets a free, tunable exponent (B), but familiarity (assocs, the pair's current association strength) enters linearly, with an implicit weight of 1 and no free parameter of its own. This adds one parameter, gamma, raising familiarity to a free power instead: $\text{assocs}^{\text{gamma}} * \exp(B * \text{entropy})$. `gamma = 1` is identical to `uncfam()`; `gamma < 1` dampens the "rich get richer" effect (diminishing returns on an already-strong pairing); `gamma > 1` amplifies it (winner-take-all).

Usage

```
uncfam_gamma(X, B, C, gamma)
```

Arguments

X	Associative weight to distribute
B	Weighting of uncertainty vs. familiarity
C	Decay
gamma	Familiarity exponent (1 = identical to <code>uncfam()</code>)

Details

Cross-validated against all of [xsl_datasets](#) (SSE to human 18AFC accuracy, the package's standard model-comparison objective), `uncfam_gamma()` beats plain `uncfam()` in every fold of a 5-fold split (mean test SSE 0.48 vs. 0.56), with a best-fit gamma around 0.3 – i.e. familiarity's pull *dampens* with diminishing returns when explaining how well people learn from a passively-received training sequence. A separate analysis fitting this same free-gamma substrate to *active* cross-situational word learning (Kachergis, Yu, & Shiffrin's paradigm where learners choose which items to see named next, rather than receiving a fixed passive sequence) found a best-fit gamma around 2 instead – amplified, not dampened. Both contexts agree that familiarity should get its own free exponent rather than `uncfam()`'s hard-coded linear one, but disagree on which direction it bends: the process governing moment-to-moment choice of what to study next doesn't appear to be simply the same one governing how much is learned from a given, already-fixed sequence.

Value

An object of class `xslMod`

Examples

```
mod <- uncfam_gamma(X = .1, C = 1, B = .98, gamma = 1)
xsl_run(mod, get_example_ambiguous_condition())

# gamma = 1 reproduces uncfam() exactly
sse_gamma1 <- xsl_run(uncfam_gamma(X = .1, C = 1, B = .98, gamma = 1),
  get_example_ambiguous_condition())$sse
sse_uncfam <- xsl_run(uncfam(X = .1, C = 1, B = .98),
  get_example_ambiguous_condition())$sse
isTRUE(all.equal(sse_gamma1, sse_uncfam))

mod <- uncfam_gamma(X = .1, C = 1, B = .98, gamma = 0.3) # dampened familiarity
xsl_run(mod, get_example_ambiguous_condition())
```

uncfam_predictive	<i>Predictive biased associative model (item-level prediction error)</i>
-------------------	--

Description

A variant of `uncfam()` (Kachergis et al. 2012's uncertainty- and familiarity-biased associative model) that adds an item-level, Rescorla-Wagner-style prediction error term: on each trial, the amount learned about a word-object association is scaled by how much the word's predicted association strength (summed over the trial's objects) falls short of the maximum value, rather than normalizing to distribute a fixed amount of associative weight across the trial. This lets initially mis-paired ("surprising") items draw more learning than the un-normalized original model allows. Associations are bounded to $[\emptyset, \text{beta}]$ (with $\text{beta} = 1$): without normalization, the update can otherwise overshoot below 0 (and diverge) or far above beta on a word's first exposure.

Usage

```
uncfam_predictive(X, B, C)
```

Arguments

X	Associative weight to distribute
B	Weighting of uncertainty vs. familiarity
C	Decay

Value

An object of class `xslMod`

Examples

```
mod <- uncfam_predictive(X = .1, C = 1, B = .98)
xsl_run(mod, get_example_ambiguous_condition())
```

uncfam_sampling	<i>Kachergis 2012</i>
-----------------	-----------------------

Description

Kachergis et al. 2012 uncertainty- and familiarity-biased associative model

Usage

```
uncfam_sampling(X, B, C, K = 1)
```

Arguments

X	Associative weight to distribute
B	Weighting of uncertainty vs. familiarity
C	Decay
K	Number of associations to update per word (constraint: $B \leq X$)

Value

An object of class `xslMod`

Examples

```
mod <- uncfam_sampling(X = .1, C = 1, B = .98, K = 1)
xsl_run(mod, get_example_ambiguous_condition())
```

update_params	<i>Update parameters of xslMod</i>
---------------	------------------------------------

Description

Replaces `xsl_mod`'s leading parameter values with `params`, matched by position (not by name) – e.g. a raw numeric vector such as `DEoptim`'s `bestmem`. The result keeps `xsl_mod`'s original parameter names.

Usage

```
update_params(xsl_mod, params)
```

Arguments

<code>xsl_mod</code>	Object of class <code>xslMod</code>
<code>params</code>	Vector or list of parameter values, in the same order as the leading elements of <code>xsl_mod\$params</code>

Details

`params` may be shorter than `xsl_mod$params`, in which case only the first `length(params)` parameters are replaced and the rest keep their existing values. This is what lets `xsl_fit()` optimize a leading subset of a model's parameters (e.g. `uncfam()`'s numeric $X/B/C$ but not its categorical variant) by giving `DEoptim` lower/upper bounds shorter than the model's full parameter list.

Value

An object of class `xslMod`

vlach_debrock2017

Vlach & DeBrock (2017) spaced/massed repetition experiment data

Description

An [xslData](#) object from Vlach, H. A., & DeBrock, C. A. (2017). "Remember dax? Relations between children's cross-situational word learning, memory, and language abilities." *Journal of Memory and Language*, 93, 217-230. 12 to-be-learned word-object pairs (symmetric design: word *i* pairs with object *i*), 2 pairs per training trial, 36 trials; pairs vary in whether their 6 repetitions fall on consecutive trials (massed) or are spread out (spaced), the paper's memory/spacing manipulation. `n_subj` is 47 children (ages 2-5).

Usage

```
vlach_debrock2017
```

Format

An object of class `xslData` (inherits from `list`) of length 8.

Details

No per-word accuracy is available from this dataset, only the paper's overall mean across all pairs and participants (accuracy = .5583, sd = .1975; see `data-raw/XSL-dataset-fields.csv`), so accuracy is NA throughout. **Not included in [xsl_datasets](#)** and not scorable by `xsl_run()`'s built-in per-item SSE – usable for running/fitting models and comparing simulated performance against the one real number this paper reports. See `data-raw/add_vlach_debrock2017.R` for construction details.

Examples

```
m <- xsl_run(baseline(), vlach_debrock2017)$fits[[1]]$matrix
dim(m)
```

xslControl-class

xslControl S3 class

Description

Control arguments for `xsl_run()`

Usage

```
xslControl(
  reps = 1,
  start_matrix = NULL,
  test_noise = 0,
  n_sim = 500,
  verbose = FALSE,
  keep_traj = FALSE,
  keep_sims = FALSE
)

new_xslControl(x = list())
```

Arguments

<code>reps</code>	Number of times to repeat training
<code>start_matrix</code>	Starting matrix
<code>test_noise</code>	Test noise
<code>n_sim</code>	Number of simulations for stochastic models
<code>verbose</code>	Verbosity
<code>keep_traj</code>	Whether each model run should record its per-trial association matrix (<code>xslFit\$traj</code>). Off by default: the trajectory is not used by any function in the package, and for a long corpus it is a serious memory cost (one word-by-object matrix per trial, per simulation). Set TRUE only when you want to inspect the learning trajectory yourself.
<code>keep_sims</code>	Whether <code>xsl_run()</code> should retain the full list of per-simulation <code>xslFit</code> objects (<code>fits[[i]]\$sims</code>). Off by default; <code>fits[[i]]\$responses</code> (an <code>n_sim</code> x <code>n-words</code> matrix of each simulated participant's final per-word accuracy) is always returned instead, at a fraction of the memory.
<code>x</code>	List with elements <code>train</code> , <code>test</code> , <code>accuracy</code> , <code>n_subj</code> , <code>label</code> , <code>condition</code>

Value

An object of class `xslControl`

<code>xslData-class</code>	<i>xslData S3 class</i>
----------------------------	-------------------------

Description

`xslData` S3 class

Constructor for `xslData` S3 class

Usage

```
xslData(
  train = list(),
  test = list(),
  accuracy = numeric(),
  n_subj = numeric(),
  label = "",
  condition = "",
  description = "",
  response_matrix = NULL
)

new_xslData(x = list())
```

Arguments

train	Training data
test	Test data
accuracy	Accuracy data
n_subj	Number of subjects
label	Label
condition	Condition
description	Description
response_matrix	Response matrix
x	List with elements train, test, accuracy, n_subj, label, condition, description, response_matrix

Value

An object of class xslData

xslFit-class

xslFit S3 class

Description

xslFit S3 class

Usage

```
xslFit(perf = matrix(), matrix = matrix(), traj = list(), sse = double())
```

Arguments

perf	Perf
matrix	Matrix
traj	Traj
sse	SSE

Value

An object of class xslFit

xsl_datasets	<i>Cross-situational word learning experiment data</i>
--------------	--

Description

Cross-situational word learning experiment data

Usage

xsl_datasets

Format

An object of class list of length 56.

xsl_fit	<i>Fit XSL model using differential evolution</i>
---------	---

Description

Fits a model to provided data using the Differential Evolution optimization algorithm. It optimizes the model parameters to minimize the sum of squared errors (SSE) between the model's predictions and human accuracy data.

Usage

```
xsl_fit(  
  model,  
  data,  
  lower,  
  upper,  
  by_data = FALSE,  
  control = xslControl(),  
  deoptim_control = DEoptim::DEoptim.control(reltol = 0.001, NP = 100, itermax = 100)  
)
```

Arguments

<code>model</code>	An object of class <code>xslMod</code> .
<code>data</code>	An object (or list of objects) of class <code>xslData</code> .
<code>lower</code>	Numeric vector of lower bounds for the model's parameters.
<code>upper</code>	Numeric vector of upper bounds for the model's parameters.
<code>by_data</code>	Logical indicating whether to fit to each entry in data separately.
<code>control</code>	Control parameters passed to <code>xsl_run()</code> .
<code>deoptim_control</code>	Control parameters passed to <code>DEoptim()</code> .

Value

An object of class `DEoptim` representing the fitting result, which includes the best set of parameters found and the corresponding SSE value.

`xsl_holdout_datasets` *Registry of datasets held out from `xsl_datasets`*

Description

A named list, one entry per dataset in the package that is **not** part of `xsl_datasets` and therefore never seen by `get_group_model_fit()/get_crossvalidated_model_fit()` – i.e. every dataset available for testing a fitted model's generalization to genuinely held-out data. This is a lookup table (dataset name -> how to score it), not a copy of the data itself – look up the dataset by its own name (e.g. `benitez2020`, `vlach_debrock2017`) to use it.

Usage

```
xsl_holdout_datasets
```

Format

An object of class `list` of length 7.

Details

Each entry has:

`scorable` Whether the dataset can be scored with the package's standard diagonal convention (`xsl_run()` + `mafc_test()/get_perf()`) out of the box.

`scoring` How to actually score a model's fit against it.

`note` Why it's held out, or any caveat for using it.

Most of these are standalone for a structural reason (partial-NA accuracy, a non-square association matrix, or no human accuracy data at all, e.g. the naturalistic corpora) – see each dataset's own ?help page for the full rationale. `benitez2020` and `kachergis_initial_accuracy` are the exceptions: both are directly scorable the same way `xsl_datasets` conditions are, and are kept out of the group/CV fitting pool specifically to have real held-out data to test generalization against.

Examples

```
xsl_holdout_datasets$benitez2020$scoring  
names(xsl_holdout_datasets)[sapply(xsl_holdout_datasets, `[`, "scorable")]
```

xsl_run*Run XSL model*

Description

Run XSL model

Usage

```
xsl_run(model, data, control = xslControl())
```

Arguments

model	An object of class xslMod.
data	An object (or list of objects) of class xslData.
control	Control arguments returned by xsl_control().

Value

A list with sse, unweighted_sse, and fits (one entry per dataset). Each fit has matrix (the word-by-object matrix, summed over simulations for a stochastic model), perf, sse, data, responses (an n_sim x n-words matrix of each simulated participant's final per-word accuracy), and sims (the full per-simulation xslFit list, NULL unless control had keep_sims = TRUE).

Index

* datasets

- benitez2020, 4
- fm_corpus, 12
- gangwani2011_category, 13
- kachergis2012_highlighting, 22
- kachergis_initial_accuracy, 22
- rollins_corpus, 32
- vlach_debrock2017, 44
- xsl_datasets, 47
- xsl_holdout_datasets, 48

baseline, 3

bayesian_decay, 3

benitez2020, 4

create_cooc_matrix, 5

cross_validated_group_fits, 5

decay, 7

fazly, 7

fgt2009, 8

fgt2009(), 9, 11

fgt2009_rsa, 9

fgt2009_rsa(), 9, 11

fgt2009_sweep_alpha, 10

fgt2009_sweep_alpha(), 9

fm_corpus, 11, 12

gangwani2011_category, 13

get_crossvalidated_model_fit, 14

get_crossvalidated_model_fit(), 48

get_example_ambiguous_condition, 15

get_example_unambiguous_condition, 15

get_fscore, 16

get_fscore(), 12, 33

get_group_model_fit, 17

get_group_model_fit(), 48

get_perf, 18

get_roc, 19

get_roc(), 12, 33

get_roc_max (get_roc), 19

get_roc_max(), 11, 33

get_tp, 20

get_tp(), 12, 33

get_train_test_split, 20

guess_and_test, 21

kachergis2012_highlighting, 22

kachergis_initial_accuracy, 22, 40

kalman_filter, 23

kalman_filter(), 18

mafc_test, 24

minerva2, 25

minerva2(), 31, 37

multi_sampling, 26

new_xslControl (xslControl-class), 44

new_xslData (xslData-class), 45

new_xslFit, 26

new_xslMod (new_xslFit), 26

plot_roc (get_roc), 19

plot_training_trials, 27

predict_referent, 28

predict_referent(), 18, 24, 33

propose_but_verify, 29

propose_but_verify(), 35

pursuit, 30

pursuit(), 35

rem, 31

rescorla_wagner, 32

rescorla_wagner(), 18, 23

rollins_corpus, 11, 12, 32

rsa_listener, 33

rsa_listener(), 29

rsa_speaker (rsa_listener), 33

show_datasets, 34

show_models, 35

softmax_rl, 35

tilles, 36

todam, 37

uncfam, 38

uncfam(), 39–42

uncfam_attention, 39

uncfam_elimination, 39

uncfam_gamma, 40

uncfam_predictive, 42

uncfam_sampling, 42

update_params, 43

vlach_debrock2017, 44

xsl_datasets, 4, 12, 13, 22, 23, 33, 41, 44, 47, 48

xsl_fit, 47

xsl_holdout_datasets, 48

xsl_run, 49

xslControl (xslControl-class), 44

xslControl-class, 44

xslData, 4, 12, 13, 33, 44

xslData (xslData-class), 45

xslData-class, 45

xslFit (xslFit-class), 46

xslFit-class, 46

xslMod (new_xslFit), 26

xslMod-class (new_xslFit), 26