

Package ‘seededlda’

May 31, 2023

Type Package

Title Seeded Sequential LDA for Topic Modeling

Version 1.0.0

Description

Seeded Sequential LDA can classify sentences of texts into pre-define topics with a small number of seed words (Watanabe & Baturo, 2023) <[doi:10.1177/08944393231178605](https://doi.org/10.1177/08944393231178605)>.

Implements Seeded LDA (Lu et al., 2010) <[doi:10.1109/ICDMW.2011.125](https://doi.org/10.1109/ICDMW.2011.125)> and Sequential LDA (Du et al., 2012) <[doi:10.1007/s10115-011-0425-1](https://doi.org/10.1007/s10115-011-0425-1)> with the distributed LDA algorithm (Newman, et al., 2009) for parallel computing.

License GPL-3

LazyData TRUE

URL <https://github.com/koheiw/seededlda>

BugReports <https://github.com/koheiw/seededlda/issues>

Encoding UTF-8

Depends R (>= 3.5.0), quanteda (>= 3.3.0), methods, proxyC (>= 0.3.1)

Imports Matrix

LinkingTo Rcpp, RcppParallel, RcppArmadillo (>= 0.7.600.1.0),
quanteda, testthat

Suggests testthat, topicmodels, keyATM

RoxygenNote 7.2.3

NeedsCompilation yes

Author Kohei Watanabe [aut, cre, cph],
Phan Xuan-Hieu [aut, cph] (GibbsLDA++)

Maintainer Kohei Watanabe <watanabe.kohei@gmail.com>

Repository CRAN

Date/Publication 2023-05-31 04:20:02 UTC

R topics documented:

data_corpus_moviereviews	2
divergence	2
sizes	3
terms	4
textmodel_lda	4
textmodel_seededlda	6
textmodel_seqlda	8
topics	9

Index	10
--------------	-----------

data_corpus_moviereviews	<i>Movie reviews from Pang and Lee (2004)</i>
--------------------------	---

Description

A corpus object containing 2,000 movie reviews.

Source

<https://www.cs.cornell.edu/people/pabo/movie-review-data/>

References

Pang, B., Lee, L. (2004) "A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts.", Proceedings of the ACL.

divergence	<i>Optimize the number of topics</i>
------------	--------------------------------------

Description

divergence() computes the regularized topic divergence to find the optimal number of topics for LDA.

Usage

```
divergence(x, min_size = 0.01, select = NULL, regularize = TRUE)
```

Arguments

x	a LDA model fitted by <code>textmodel_seededlda()</code> or <code>textmodel_lda()</code> .
min_size	the minimum size of topics for regularized topic divergence. Ignored when <code>regularize = FALSE</code> .
select	names of topics for which the divergence is computed.
regularize	if TRUE, returns the regularized divergence.

Details

`divergence()` computes the average Jensen-Shannon divergence between all the pairs of topic vectors in `x$phi`. The divergence score maximizes when the chosen number of topic `k` is optimal (Deveaud et al., 2014). The regularized divergence penalizes topics smaller than `min_size` to avoid fragmentation (Watanabe & Baturo, forthcoming).

References

- Deveaud, Romain et al. (2014). "Accurate and Effective Latent Concept Modeling for Ad Hoc Information Retrieval". doi:10.3166/DN.17.1.61-84. *Document Numérique*.
- Watanabe, Kohei & Baturo, Alexander. (2023). "Seeded Sequential LDA: A Semi-supervised Algorithm for Topic-specific Analysis of Sentences". doi:10.1177/08944393231178605. *Social Science Computer Review*.

See Also

[sizes](#)

sizes	<i>Compute the sizes of topics</i>
-------	------------------------------------

Description

Compute the sizes of topics as the proportions of topic words in the corpus.

Usage

```
sizes(x)
```

Arguments

x	a LDA model fitted by <code>textmodel_seededlda()</code> or <code>textmodel_lda()</code>
---	--

terms *Extract most likely terms*

Description

terms() returns the most likely terms, or words, for topics based on the phi parameter.

Usage

```
terms(x, n = 10)
```

Arguments

x a LDA model fitted by `textmodel_seededlda()` or `textmodel_lda()`
n number of terms to be extracted

Details

Users can access the original matrix `x$phi` for likelihood scores.

textmodel_lda *Unsupervised Latent Dirichlet allocation*

Description

Implements unsupervised Latent Dirichlet allocation (LDA). Users can run Seeded LDA by setting `gamma > 0`.

Usage

```
textmodel_lda(  
  x,  
  k = 10,  
  max_iter = 2000,  
  auto_iter = FALSE,  
  alpha = 0.5,  
  beta = 0.1,  
  gamma = 0,  
  model = NULL,  
  batch_size = 1,  
  verbose = quanteda_options("verbose")  
)
```

Arguments

x	the dfm on which the model will be fit.
k	the number of topics.
max_iter	the maximum number of iteration in Gibbs sampling.
auto_iter	if TRUE, stops Gibbs sampling on convergence before reaching max_iter. See details.
alpha	the value to smooth topic-document distribution.
beta	the value to smooth topic-word distribution.
gamma	a parameter to determine change of topics between sentences or paragraphs. When $\gamma > 0$, Gibbs sampling of topics for the current document is affected by the previous document's topics.
model	a fitted LDA model; if provided, textmodel_lda() inherits parameters from an existing model. See details.
batch_size	split the corpus into the smaller batches (specified in proportion) for distributed computing; it is disabled when a batch include all the documents batch_size = 1.0. See details.
verbose	logical; if TRUE print diagnostic information during fitting.

Details

If auto_iter = TRUE, the iteration stops even before max_iter when $\delta \leq 0$. δ is computed to measure the changes in the number of words whose topics are updated by the Gibbs sampler in every 100 iteration as shown in the verbose message.

If batch_size < 1.0, the corpus is partitioned into sub-corpora of $\text{ndoc}(x) * \text{batch_size}$ documents for Gibbs sampling in sub-processes with synchronization of parameters in every 10 iteration. Parallel processing is more efficient when batch_size is small (e.g. 0.01). The algorithm is the Approximate Distributed LDA proposed by Newman et al. (2009). User can changed the number of sub-processes used for the parallel computing via options(seededlda_threads).

To predict topics of new documents (i.e. out-of-sample), first, create a new LDA model from a existing LDA model passed to model in textmodel_lda(); second, apply `topics()` to the new model. The model argument takes objects created either by textmodel_lda() or textmodel_seededlda().

Value

textmodel_seededlda() and textmodel_lda() returns a list of model parameters. theta is the distribution of topics over documents; phi is the distribution of words over topics. alpha and beta are the small constant added to the frequency of words to estimate theta and phi, respectively, in Gibbs sampling. Other elements in the list subject to change.

References

Newman, D., Asuncion, A., Smyth, P., & Welling, M. (2009). Distributed Algorithms for Topic Models. *The Journal of Machine Learning Research*, 10, 1801–1828.

See Also[LDA weightedLDA](#)**Examples**

```
require(seededlda)
require(quanteda)

corp <- head(data_corpus_moviereviews, 500)
toks <- tokens(corp, remove_punct = TRUE, remove_symbols = TRUE, remove_number = TRUE)
dfmt <- dfm(toks) %>%
  dfm_remove(stopwords('en'), min_nchar = 2) %>%
  dfm_trim(min_termfreq = 0.90, termfreq_type = "quantile",
           max_docfreq = 0.1, docfreq_type = "prop")

lda <- textmodel_lda(head(dfmt, 450), 6)
terms(lda)
topics(lda)
lda2 <- textmodel_lda(tail(dfmt, 50), model = lda) # new documents
topics(lda2)
```

textmodel_seededlda *Semisupervised Latent Dirichlet allocation*

Description

Implements semisupervised Latent Dirichlet allocation (Seeded LDA). `textmodel_seededlda()` allows users to specify topics using a seed word dictionary. Users can run Seeded Sequential LDA by setting `gamma > 0`.

Usage

```
textmodel_seededlda(
  x,
  dictionary,
  valuetype = c("glob", "regex", "fixed"),
  case_insensitive = TRUE,
  residual = 0,
  weight = 0.01,
  uniform = TRUE,
  max_iter = 2000,
  auto_iter = FALSE,
  alpha = 0.5,
  beta = 0.1,
  gamma = 0,
  batch_size = 1,
```

```

    ...,
    verbose = quanteda_options("verbose")
)

```

Arguments

x	the dfm on which the model will be fit.
dictionary	a <code>quanteda::dictionary()</code> with seed words that define topics.
valuetype	see <code>quanteda::valuetype</code>
case_insensitive	see <code>quanteda::valuetype</code>
residual	the number of undefined topics. They are named "other" by default, but it can be changed via <code>base::options(seededlda_residual_name)</code> .
weight	determines the size of pseudo counts given to matched seed words.
uniform	if FALSE, adjusts the weights of seed words to make their total amount equal across topics.
max_iter	the maximum number of iteration in Gibbs sampling.
auto_iter	if TRUE, stops Gibbs sampling on convergence before reaching <code>max_iter</code> . See details.
alpha	the value to smooth topic-document distribution.
beta	the value to smooth topic-word distribution.
gamma	a parameter to determine change of topics between sentences or paragraphs. When $\gamma > 0$, Gibbs sampling of topics for the current document is affected by the previous document's topics.
batch_size	split the corpus into the smaller batches (specified in proportion) for distributed computing; it is disabled when a batch include all the documents <code>batch_size = 1.0</code> . See details.
...	passed to <code>quanteda::dfm_trim</code> to restrict seed words based on their term or document frequency. This is useful when glob patterns in the dictionary match too many words.
verbose	logical; if TRUE print diagnostic information during fitting.

References

- Lu, Bin et al. (2011). "Multi-aspect Sentiment Analysis with Topic Models". doi:10.5555/2117693.2119585. *Proceedings of the 2011 IEEE 11th International Conference on Data Mining Workshops*.
- Watanabe, Kohei & Zhou, Yuan (2020). "Theory-Driven Analysis of Large Corpora: Semisupervised Topic Classification of the UN Speeches". doi:10.1177/0894439320907027. *Social Science Computer Review*.
- Watanabe, Kohei & Baturo, Alexander. (2023). "Seeded Sequential LDA: A Semi-supervised Algorithm for Topic-specific Analysis of Sentences". doi:10.1177/08944393231178605. *Social Science Computer Review*.

See Also[keyATM](#)**Examples**

```

require(seededlda)
require(quanteda)

corp <- head(data_corpus_moviereviews, 500)
toks <- tokens(corp, remove_punct = TRUE, remove_symbols = TRUE, remove_number = TRUE)
dfmt <- dfm(toks) %>%
  dfm_remove(stopwords('en'), min_nchar = 2) %>%
  dfm_trim(min_termfreq = 0.90, termfreq_type = "quantile",
           max_docfreq = 0.1, docfreq_type = "prop")

dict <- dictionary(list(people = c("family", "couple", "kids"),
                        space = c("alien", "planet", "space"),
                        moster = c("monster*", "ghost*", "zombie*"),
                        war = c("war", "soldier*", "tanks"),
                        crime = c("crime*", "murder", "killer")))
sllda <- textmodel_seededlda(dfmt, dict, residual = TRUE, min_termfreq = 10)
terms(sllda)
topics(sllda)

```

textmodel_seqlda

Sequential Latent Dirichlet allocation

Description

Implements Sequential Latent Dirichlet allocation (Sequential LDA). `textmodel_seqlda()` allows the users to classify sentences of texts. It considers the topics of previous document in inferring the topics of currency document. `textmodel_seqlda()` is a shortcut equivalent to `textmodel_lda(gamma = 0.5)`. Seeded Sequential LDA is `textmodel_seededlda(gamma = 0.5)`.

Usage

```

textmodel_seqlda(
  x,
  k = 10,
  max_iter = 2000,
  auto_iter = FALSE,
  alpha = 0.5,
  beta = 0.1,
  batch_size = 1,
  model = NULL,
  verbose = quanteda_options("verbose")
)

```


Arguments

x	the dfm on which the model will be fit.
k	the number of topics.
max_iter	the maximum number of iteration in Gibbs sampling.
auto_iter	if TRUE, stops Gibbs sampling on convergence before reaching max_iter. See details.
alpha	the value to smooth topic-document distribution.
beta	the value to smooth topic-word distribution.
batch_size	split the corpus into the smaller batches (specified in proportion) for distributed computing; it is disabled when a batch include all the documents batch_size = 1.0. See details.
model	a fitted LDA model; if provided, textmodel_lda() inherits parameters from an existing model. See details.
verbose	logical; if TRUE print diagnostic information during fitting.

References

- Du, Lan et al. (2012). "Sequential Latent Dirichlet Allocation". doi.org/10.1007/s10115-011-0425-1. *Knowledge and Information Systems*.
- Watanabe, Kohei & Baturo, Alexander. (2023). "Seeded Sequential LDA: A Semi-supervised Algorithm for Topic-specific Analysis of Sentences". doi:10.1177/08944393231178605. *Social Science Computer Review*.

topics	<i>Extract most likely topics</i>
--------	-----------------------------------

Description

topics() returns the most likely topics for documents based on the theta parameter.

Usage

```
topics(x, min_prob = 0, select = NULL)
```

Arguments

x	a LDA model fitted by <code>textmodel_seededlda()</code> or <code>textmodel_lda()</code>
min_prob	ignores topics if their probability is lower than this value.
select	returns the selected topic with the highest probability; specify by the names of columns in x\$theta.

Details

Users can access the original matrix x\$theta for likelihood scores; run `max.col(x$theta)` to obtain the same result as `topics(x)`.

Index

* **textmodel**

textmodel_lda, 4

data_corpus_moviereviews, 2
divergence, 2

keyATM, 8

LDA, 6

quanteda::dfm_trim, 7
quanteda::dictionary(), 7
quanteda::valuetype, 7

sizes, 3, 3

terms, 4

textmodel_lda, 4

textmodel_lda(), 3, 4, 9

textmodel_seededlda, 6

textmodel_seededlda(), 3, 4, 9

textmodel_seqlda, 8

topics, 9

topics(), 5

weightedLDA, 6