

# Package: LSX (via r-universe)

August 22, 2024

**Type** Package

**Title** Semi-Supervised Algorithm for Document Scaling

**Version** 1.4.1

**Description** A word embeddings-based semi-supervised model for document scaling Watanabe (2020) <[doi:10.1080/19312458.2020.1832976](https://doi.org/10.1080/19312458.2020.1832976)>. LSS allows users to analyze large and complex corpora on arbitrary dimensions with seed words exploiting efficiency of word embeddings (SVD, Glove). It can generate word vectors on a users-provided corpus or incorporate a pre-trained word vectors.

**License** GPL-3

**LazyData** TRUE

**Encoding** UTF-8

**Depends** methods, R (>= 3.5.0)

**Imports** quanteda (>= 2.0), quanteda.textstats, stringi, digest, Matrix, RSpectra, irlba, rsvd, rsparse, proxyC, stats, ggplot2, ggrepel, reshape2, locfit

**Suggests** spelling, knitr, rmarkdown, testthat

**RoxygenNote** 7.3.1

**Roxygen** list(markdown = TRUE)

**BugReports** <https://github.com/koheiw/LSX/issues>

**URL** <https://koheiw.github.io/LSX/>

**Language** en-US

**Repository** <https://koheiw.r-universe.dev>

**RemoteUrl** <https://github.com/koheiw/lxs>

**RemoteRef** HEAD

**RemoteSha** 1a06b33dd840ea23c5bc25f4fbb1ac5bd37c6f2e

## Contents

as.seedwords . . . . .	2
bootstrap_lss . . . . .	3
coef.textmodel_lss . . . . .	4
data_dictionary_ideology . . . . .	4
data_dictionary_sentiment . . . . .	4
data_textmodel_lss_russianprotests . . . . .	5
optimize_lss . . . . .	5
predict.textmodel_lss . . . . .	6
seedwords . . . . .	7
smooth_lss . . . . .	8
textmodel_lss . . . . .	9
textplot_simil . . . . .	11
textplot_terms . . . . .	11
textstat_context . . . . .	12
<b>Index</b>	<b>14</b>

---

as.seedwords	<i>Convert a list or a dictionary to seed words</i>
--------------	---

---

### Description

Convert a list or a dictionary to seed words

### Usage

```
as.seedwords(x, upper = 1, lower = 2, concatenator = "_")
```

### Arguments

x	a list of characters vectors or a <a href="#">dictionary</a> object.
upper	numeric index or key for seed words for higher scores.
lower	numeric index or key for seed words for lower scores.
concatenator	character to replace separators of multi-word seed words.

### Value

named numeric vector for seed words with polarity scores

---

bootstrap_lss	<i>[experimental] Compute polarity scores with different hyper-parameters</i>
---------------	---

---

### Description

A function to compute polarity scores of words and documents by resampling hyper-parameters from a fitted LSS model.

### Usage

```
bootstrap_lss(  
  x,  
  what = c("seeds", "k"),  
  mode = c("terms", "coef", "predict"),  
  remove = FALSE,  
  from = 100,  
  to = NULL,  
  by = 50,  
  verbose = FALSE,  
  ...  
)
```

### Arguments

x	a fitted <code>textmodel_lss</code> object.
what	choose the hyper-parameter to resample in bootstrapping.
mode	choose the type of the result of bootstrapping. If <code>coef</code> , returns the polarity scores of words; if <code>terms</code> , returns words sorted by the polarity scores in descending order; if <code>predict</code> , returns the polarity scores of documents.
remove	if TRUE, remove each seed word when <code>what = "seeds"</code> .
from, to, by	passed to <code>seq()</code> to generate values for k; only used when <code>what = "k"</code> .
verbose	show messages if TRUE.
...	additional arguments passed to <code>as.textmodel_lss()</code> and <code>predict()</code> .

### Details

`bootstrap_lss()` creates LSS fitted `textmodel_lss` objects internally by resampling hyper-parameters and computes polarity of words or documents. The resulting matrix can be used to assess the validity and the reliability of seeds or k.

Note that the objects created by `as.textmodel_lss()` does not contain data, users must pass newdata via ... when `mode = "predict"`.

---

`coef.textmodel_lass`      *Extract model coefficients from a fitted `textmodel_lass` object*

---

### Description

`coef()` extract model coefficients from a fitted `textmodel_lass` object. `coefficients()` is an alias.

### Usage

```
## S3 method for class 'textmodel_lass'
coef(object, ...)

coefficients.textmodel_lass(object, ...)
```

### Arguments

`object`            a fitted `textmodel_lass` object.  
`...`              not used.

---

`data_dictionary_ideology`  
*Seed words for analysis of left-right political ideology*

---

### Description

Seed words for analysis of left-right political ideology

### Examples

```
as.seedwords(data_dictionary_ideology)
```

---

`data_dictionary_sentiment`  
*Seed words for analysis of positive-negative sentiment*

---

### Description

Seed words for analysis of positive-negative sentiment

### References

Turney, P. D., & Littman, M. L. (2003). Measuring Praise and Criticism: Inference of Semantic Orientation from Association. *ACM Trans. Inf. Syst.*, 21(4), 315–346. doi:10.1145/944012.944013

**Examples**

```
as.seedwords(data_dictionary_sentiment)
```

---

```
data_textmodel_lss_russianprotests
```

*A fitted LSS model on street protest in Russia*

---

**Description**

This model was trained on a Russian media corpus (newspapers, TV transcripts and newswires) to analyze framing of street protests. The scale is protests as "freedom of expression" (high) vs "social disorder" (low). Although some slots are missing in this object (because the model was imported from the original Python implementation), it allows you to scale texts using `predict`.

**References**

Lankina, Tomila, and Kohei Watanabe. "‘Russian Spring’ or ‘Spring Betrayal’? The Media as a Mirror of Putin’s Evolving Strategy in Ukraine." *Europe-Asia Studies* 69, no. 10 (2017): 1526–56. doi:10.1080/09668136.2017.1397603.

---

```
optimize_lss
```

*[experimental] Compute variance ratios with different hyper-parameters*

---

**Description**

[experimental] Compute variance ratios with different hyper-parameters

**Usage**

```
optimize_lss(x, ...)
```

**Arguments**

`x` a fitted `textmodel_lss` object.  
`...` additional arguments passed to `bootstrap_lss`.

**Details**

`optimize_lss()` computes variance ratios with different values of hyper-parameters using `bootstrap_lss`. The variance ration  $v$  is defined as

$$v = \sigma_{documents}^2 / \sigma_{words}^2.$$

It maximizes when the model best distinguishes between the documents on the latent scale.

**Examples**

```
## Not run:
# the unit of analysis is not sentences
dfmt_grp <- dfm_group(dfmt)

# choose best k
v1 <- optimize_lss(lss, what = "k", from = 50,
                  newdata = dfmt_grp, verbose = TRUE)
plot(names(v1), v1)

# find bad seed words
v2 <- optimize_lss(lss, what = "seeds", remove = TRUE,
                  newdata = dfmt_grp, verbose = TRUE)
barplot(v2, las = 2)

## End(Not run)
```

---

predict.textmodel\_lss *Prediction method for textmodel\_lss*

---

**Description**

Prediction method for textmodel\_lss

**Usage**

```
## S3 method for class 'textmodel_lss'
predict(
  object,
  newdata = NULL,
  se_fit = FALSE,
  density = FALSE,
  rescale = TRUE,
  cut = NULL,
  min_n = 0L,
  ...
)
```

**Arguments**

object	a fitted LSS textmodel.
newdata	a dfm on which prediction should be made.
se_fit	if TRUE, returns standard error of document scores.
density	if TRUE, returns frequency of polarity words in documents.
rescale	if TRUE, normalizes polarity scores using scale().

<code>cut</code>	a vector of one or two percentile values to dichotomized polarity scores of words. When two values are given, words between them receive zero polarity.
<code>min_n</code>	set the minimum number of polarity words in documents.
<code>...</code>	not used

### Details

Polarity scores of documents are the means of polarity scores of words weighted by their frequency. When `se_fit = TRUE`, this function returns the weighted means, their standard errors, and the number of polarity words in the documents. When `rescale = TRUE`, it converts the raw polarity scores to z scores for easier interpretation. When `rescale = FALSE` and `cut` is used, polarity scores of documents are bounded by `[-1.0, 1.0]`.

Documents tend to receive extreme polarity scores when they have only few polarity words. This is problematic when LSS is applied to short documents (e.g. social media posts) or individual sentences, but users can alleviate this problem by adding zero polarity words to short documents using `min_n`. This setting does not affect empty documents.

---

seedwords

*Seed words for Latent Semantic Analysis*

---

### Description

Seed words for Latent Semantic Analysis

### Usage

```
seedwords(type)
```

### Arguments

<code>type</code>	type of seed words currently only for sentiment ( <code>sentiment</code> ) or political ideology ( <code>ideology</code> ).
-------------------	---

### References

Turney, P. D., & Littman, M. L. (2003). Measuring Praise and Criticism: Inference of Semantic Orientation from Association. *ACM Trans. Inf. Syst.*, 21(4), 315–346. doi:10.1145/944012.944013

### Examples

```
seedwords('sentiment')
```

---

`smooth_lss`*Smooth predicted polarity scores by local polynomial regression*

---

## Description

Smooth predicted polarity scores by local polynomial regression

## Usage

```
smooth_lss(  
  x,  
  lss_var = "fit",  
  date_var = "date",  
  span = 0.1,  
  group = NULL,  
  from = NULL,  
  to = NULL,  
  by = "day",  
  engine = c("loess", "locfit"),  
  ...  
)
```

## Arguments

<code>x</code>	a <a href="#">data.frame</a> containing polarity scores and dates.
<code>lss_var</code>	the name of the column in <code>x</code> for polarity scores.
<code>date_var</code>	the name of the column in <code>x</code> for dates.
<code>span</code>	the level of smoothing.
<code>group</code>	the name of the column in <code>x</code> to smooth the data by group.
<code>from, to, by</code>	the the range and the internal of the smoothed scores; passed to <a href="#">seq.Date</a> .
<code>engine</code>	specifies the function to be used for smoothing.
<code>...</code>	additional arguments passed to the smoothing function.

## Details

Smoothing is performed using [stats::loess\(\)](#) or [locfit::locfit\(\)](#). When the `x` has more than 10000 rows, it is usually better to choose the latter by setting `engine = "locfit"`. In this case, `span` is passed to `locfit::lp(nn = span)`.



---

textmodel_1ss	<i>Fit a Latent Semantic Scaling model</i>
---------------	--

---

### Description

Latent Semantic Scaling (LSS) is a word embedding-based semisupervised algorithm for document scaling.

### Usage

```
textmodel_1ss(x, ...)

## S3 method for class 'dfm'
textmodel_1ss(
  x,
  seeds,
  terms = NULL,
  k = 300,
  slice = NULL,
  weight = "count",
  cache = FALSE,
  simil_method = "cosine",
  engine = c("RSpectra", "irlba", "rsvd"),
  prop_slice = NULL,
  auto_weight = FALSE,
  include_data = FALSE,
  group_data = FALSE,
  verbose = FALSE,
  ...
)

## S3 method for class 'fcm'
textmodel_1ss(
  x,
  seeds,
  terms = NULL,
  w = 50,
  max_count = 10,
  weight = "count",
  cache = FALSE,
  simil_method = "cosine",
  engine = c("rsparse"),
  auto_weight = FALSE,
  verbose = FALSE,
  ...
)
```

**Arguments**

x	a dfm or fcm created by <code>quanteda::dfm()</code> or <code>quanteda::fcm()</code>
...	additional arguments passed to the underlying engine.
seeds	a character vector or named numeric vector that contains seed words. If seed words contain "*", they are interpreted as glob patterns. See <code>quanteda::valuetype</code> .
terms	a character vector or named numeric vector that specify words for which polarity scores will be computed; if a numeric vector, words' polarity scores will be weighted accordingly; if NULL, all the features of <code>quanteda::dfm()</code> or <code>quanteda::fcm()</code> will be used.
k	the number of singular values requested to the SVD engine. Only used when x is a dfm.
slice	a number or indices of the components of word vectors used to compute similarity; <code>slice &lt; k</code> to further truncate word vectors; useful for diagnosis and simulation.
weight	weighting scheme passed to <code>quanteda::dfm_weight()</code> . Ignored when engine is "rsparse".
cache	if TRUE, save result of SVD for next execution with identical x and settings. Use the <code>base::options(lss_cache_dir)</code> to change the location cache files to be save.
simil_method	specifies method to compute similarity between features. The value is passed to <code>quanteda.textstats::textstat_simil()</code> , "cosine" is used otherwise.
engine	select the engine to factorize x to generate word vectors. Choose from <code>RSpectra::svds()</code> , <code>irlba::irlba()</code> , <code>rsvd::rsvd()</code> , and <code>rsparse::GloVe()</code> .
prop_slice	[experimental] specify the number of compontes to use by proportion.
auto_weight	automatically determine weights to approximate the polarity of terms to seed words. See details.
include_data	if TRUE, fitted model includes the dfm supplied as x.
group_data	if TRUE, apply <code>dfm_group(x)</code> before saving the dfm.
verbose	show messages if TRUE.
w	the size of word vectors. Used only when x is a fcm.
max_count	passed to <code>x_max</code> in <code>rsparse::GloVe\$new()</code> where cooccurrence counts are ceiled to this threshold. It should be changed according to the size of the corpus. Used only when x is a fcm.

**Details**

Latent Semantic Scaling (LSS) is a semisupervised document scaling method. `textmodel_lss()` constructs word vectors from use-provided documents (x) and weights words (terms) based on their semantic proximity to seed words (seeds). Seed words are any known polarity words (e.g. sentiment words) that users should manually choose. The required number of seed words are usually 5 to 10 for each end of the scale.

If seeds is a named numeric vector with positive and negative values, a bipolar LSS model is construct; if seeds is a character vector, a unipolar LSS model. Usually bipolar models perform better in document scaling because both ends of the scale are defined by the user.

A seed word's polarity score computed by `textmodel_1ss()` tends to diverge from its original score given by the user because it's score is affected not only by its original score but also by the original scores of all other seed words. If `auto_weight = TRUE`, the original scores are weighted automatically using `stats::optim()` to minimize the squared difference between seed words' computed and original scores. Weighted scores are saved in `seed_weighted` in the object.

Please visit the [package website](#) for examples.

## References

Watanabe, Kohei. 2020. "Latent Semantic Scaling: A Semisupervised Text Analysis Technique for New Domains and Languages", *Communication Methods and Measures*. doi:10.1080/19312458.2020.1832976.

Watanabe, Kohei. 2017. "Measuring News Bias: Russia's Official News Agency ITAR-TASS' Coverage of the Ukraine Crisis" *European Journal of Communication*. doi:10.1177/0267323117695735.

---

textplot_simil	<i>Plot similarity between seed words</i>
----------------	---

---

## Description

Plot similarity between seed words

## Usage

```
textplot_simil(x)
```

## Arguments

x                    fitted textmodel\_1ss object.

---

textplot_terms	<i>Plot polarity scores of words</i>
----------------	--------------------------------------

---

## Description

Plot polarity scores of words

## Usage

```
textplot_terms(  
  x,  
  highlighted = NULL,  
  max_highlighted = 50,  
  max_words = 1000,  
  ...  
)
```

**Arguments**

x	a fitted textmodel_1ss object.
highlighted	<code>quanteda::pattern</code> to select words to highlight. If a <code>quanteda::dictionary</code> is passed, words in the top-level categories are highlighted in different colors.
max_highlighted	the maximum number of words to highlight. When <code>highlighted = NULL</code> , words to highlight are randomly selected proportionally to $\text{polarity}^2 * \log(\text{frequency})$ .
max_words	the maximum number of words to plot. Words are randomly sampled to keep the number below the limit.
...	passed to underlying functions. See the Details.

**Details**

Users can customize the plots through ..., which is passed to `ggplot2::geom_text()` and `ggrepel::geom_text_repel()`. The colors are specified internally but users can override the settings by appending `ggplot2::scale_colour_manual()` or `ggplot2::scale_colour_brewer()`. The legend title can also be modified using `ggplot2::labs()`.

---

textstat_context	<i>Identify context words using user-provided patterns</i>
------------------	--

---

**Description**

Identify context words using user-provided patterns

**Usage**

```
textstat_context(
  x,
  pattern,
  valuetype = c("glob", "regex", "fixed"),
  case_insensitive = TRUE,
  window = 10,
  min_count = 10,
  remove_pattern = TRUE,
  n = 1,
  skip = 0,
  ...
)

char_context(
  x,
  pattern,
  valuetype = c("glob", "regex", "fixed"),
  case_insensitive = TRUE,
  window = 10,
```

```

    min_count = 10,
    remove_pattern = TRUE,
    p = 0.001,
    n = 1,
    skip = 0
  )

```

### Arguments

x	a tokens object created by <a href="#">quanteda::tokens()</a> .
pattern	<a href="#">quanteda::pattern()</a> to specify target words.
valuetype	the type of pattern matching: "glob" for "glob"-style wildcard expressions; "regex" for regular expressions; or "fixed" for exact matching. See <a href="#">quanteda::valuetype()</a> for details.
case_insensitive	if TRUE, ignore case when matching.
window	size of window for collocation analysis.
min_count	minimum frequency of words within the window to be considered as collocations.
remove_pattern	if TRUE, keywords do not contain target words.
n	integer vector specifying the number of elements to be concatenated in each n-gram. Each element of this vector will define a <i>n</i> in the <i>n</i> -gram(s) that are produced.
skip	integer vector specifying the adjacency skip size for tokens forming the n-grams, default is 0 for only immediately neighbouring words. For skipgrams, skip can be a vector of integers, as the "classic" approach to forming skip-grams is to set skip = <i>k</i> where <i>k</i> is the distance for which <i>k</i> or fewer skips are used to construct the <i>n</i> -gram. Thus a "4-skip-n-gram" defined as skip = 0:4 produces results that include 4 skips, 3 skips, 2 skips, 1 skip, and 0 skips (where 0 skips are typical n-grams formed from adjacent words). See Guthrie et al (2006).
...	additional arguments passed to <a href="#">textstat_keyness()</a> .
p	threshold for statistical significance of collocations.

### See Also

[tokens\\_select\(\)](#) and [textstat\\_keyness\(\)](#)

# Index

- \* **data**
  - data\_textmodel\_lass\_russianprotests,  
5
- as.seedwords, 2
- as.textmodel\_lass(), 3
- bootstrap\_lass, 3, 5
- char\_context(textstat\_context), 12
- coef.textmodel\_lass, 4
- coefficients.textmodel\_lass  
(coef.textmodel\_lass), 4
- data.frame, 8
- data\_dictionary\_ideology, 4
- data\_dictionary\_sentiment, 4
- data\_textmodel\_lass\_russianprotests, 5
- dictionary, 2
- ggplot2::geom\_text(), 12
- ggplot2::labs(), 12
- ggplot2::scale\_colour\_brewer(), 12
- ggplot2::scale\_colour\_manual(), 12
- ggrepel::geom\_text\_repel(), 12
- irlba::irlba(), 10
- locfit::locfit(), 8
- optimize\_lass, 5
- predict(), 3
- predict.textmodel\_lass, 6
- quanteda.textstats::textstat\_simil(),  
10
- quanteda::dfm(), 10
- quanteda::dfm\_weight(), 10
- quanteda::dictionary, 12
- quanteda::fcm(), 10
- quanteda::pattern, 12
- quanteda::pattern(), 13
- quanteda::tokens(), 13
- quanteda::valuetype, 10
- quanteda::valuetype(), 13
- rsparse::GloVe(), 10
- RSpectra::svds(), 10
- rsvd::rsvd(), 10
- seedwords, 7
- seq.Date, 8
- smooth\_lass, 8
- stats::loess(), 8
- stats::optim(), 11
- textmodel\_lass, 4, 9
- textplot\_simil, 11
- textplot\_terms, 11
- textstat\_context, 12
- textstat\_keyness(), 13
- tokens\_select(), 13