

Package ‘tsne’

July 2, 2014

Type Package

Title T-distributed Stochastic Neighbor Embedding for R (t-SNE)

Version 0.1-2

Date 2010-02-19

Author Justin Donaldson <jdonaldson@gmail.com>

Maintainer Justin Donaldson <jdonaldson@gmail.com>

Description A “pure R” implementation of the t-SNE algorithm.

License GPL

LazyLoad yes

Repository CRAN

Date/Publication 2012-05-02 06:34:28

NeedsCompilation no

R topics documented:

tsne-package	1
tsne	2

Index	4
--------------	----------

tsne-package	<i>The tsne-package for multidimensional scaling</i>
--------------	--

Description

This package contains one function called `tsne` which contains all the functionality.

Details

Package: tsne
 Type: Package
 Version: 0.1
 Date: 2010-02-19
 License: GPL
 LazyLoad: yes

Author(s)

Justin Donaldson <http://www.scwn.net> Maintainer: Justin Donaldson (jdonaldson@gmail.com)

References

L.J.P. van der Maaten and G.E. Hinton. Visualizing High-Dimensional Data Using t-SNE. *Journal of Machine Learning Research* 9 (Nov) : 2579-2605, 2008.

L.J.P. van der Maaten. Learning a Parametric Embedding by Preserving Local Structure. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR W&CP 5:384-391, 2009.

tsne

The t-SNE method for dimensionality reduction

Description

Provides a simple function interface for specifying t-SNE dimensionality reduction on R matrices or "dist" objects.

Usage

```
tsne(X, initial_config = NULL, k = 2, initial_dims = 30, perplexity = 30, max_iter = 1000, min_cost = 0)
```

Arguments

X	The R matrix or "dist" object
initial_config	an argument providing a matrix specifying the initial embedding for X. See Details.
k	the dimension of the resulting embedding.
initial_dims	The number of dimensions to use in reduction method.
perplexity	Perplexity parameter. (optimal number of neighbors)
max_iter	Maximum number of iterations to perform.
min_cost	The minimum cost value (error) to halt iteration.

epoch_callback	A callback function used after each epoch (an epoch here means a set number of iterations)
whiten	A boolean value indicating whether the matrix data should be whitened.
epoch	The number of iterations in between update messages.

Details

When the `initial_config` argument is specified, the algorithm will automatically enter the *final momentum* stage. This stage has less large scale adjustment to the embedding, and is intended for small scale tweaking of positioning. This can greatly speed up the generation of embeddings for various similar X datasets, while also preserving overall embedding orientation.

Value

An R object containing a *ydata* embedding matrix, as well as a the matrix of probabilities *P*

Author(s)

Justin Donaldson (jdonaldson@gmail.com)

References

L.J.P. van der Maaten and G.E. Hinton. Visualizing High-Dimensional Data Using t-SNE. *Journal of Machine Learning Research* 9 (Nov) : 2579-2605, 2008.

L.J.P. van der Maaten. Learning a Parametric Embedding by Preserving Local Structure. In *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics (AISTATS)*, JMLR W&CP 5:384-391, 2009.

See Also

[dist](#)

Examples

```
## Not run:
colors = rainbow(length(unique(iris$Species)))
names(colors) = unique(iris$Species)
ecb = function(x,y){ plot(x,t='n'); text(x,labels=iris$Species, col=colors[iris$Species]) }
tsne_iris = tsne(iris[,1:4], epoch_callback = ecb, perplexity=50)

# compare to PCA
dev.new()
pca_iris = princomp(iris[,1:4])$scores[,1:2]
plot(pca_iris, t='n')
text(pca_iris, labels=iris$Species,col=colors[iris$Species])

## End(Not run)
```

Index

*Topic **package**
tsne-package, 1

dist, 3

tsne, 1, 2

tsne-package, 1