

t-Distributed Stochastic Neighbourhood Embedding

Summer Research Internship Proposal

Supervised by: Dr. Jason Wyse

Group size: One person

Introduction

t-SNE is a statistical dimensionality-reduction technique introduced in 2008 by van der Maaten and Hinton [1]. The method enables users to take datasets with large dimensions (hundreds or thousands) and find a faithful representation of all the datapoints in a lower dimension. In other words, the method attempts to preserve the clustering of datapoints in high dimensions and project this information onto a lower dimension, typically the 2D or 3D plane, which then allows the users to visualize the structure of the dataset. This method holds quite a few advantages over traditional dimensionality-reduction techniques such as Principal Component Analysis and Multidimensional Scaling due to its non-linearity and ability to handle outliers.

Projected Goals

Begin the project by acquiring a general idea of the method. Accomplish this by understanding how the various concepts and methodologies involved are related to one another [1 Week]. The different components will then be studied individually in greater detail. More specifically, we will examine how one can choose a suitable complexity parameter [0.5 Weeks], inspect the method of calculating the similarity between datapoints [0.5 Weeks] and investigate the method of minimizing the Kullback-Leibler divergence via gradient descent in order to obtain the projection map [1.5 Weeks].

After we've assimilated a complete understanding, we will compare the performance of t-SNE to other dimensionality-reduction techniques by applying these methods (using statistical software) to various datasets. The performance will be measured using a suitable metric to be determined by first running some pilot tests [1.5 Weeks].

With the remaining time, we shall attempt to generalize, or at least to modify the method so as to better suit specific scenarios which the model may have struggled with in the preceding stages [1 Week].

Results

A manuscript containing a detailed description of the t-SNE method along with any other significant observations or results will be produced and published on my webpage. If any noteworthy novel results are produced, a manuscript detailing these will probably be published on arXiv and/or submitted to an appropriate journal.

References

- [1] Laurens Van der Maaten and Geoffrey Hinton. "Visualizing data using t-SNE." *Journal of machine learning research* 9.11 (2008).
- [2] Martin Wattenberg, Fernanda Viégas, and Ian Johnson. "How to use t-SNE effectively". *Distill* 1.10 (2016), e2.

Figure: Preservation of clusters in 2D

