
RUSLAN SALAKHUTDINOV, University of Toronto
Non-linear Dimensionality Reduction using Neural Networks

Scientists, working with large amounts of high-dimensional data are constantly facing the problem of dimensionality reduction: how to discover low-dimensional structure from high-dimensional observations. The compact representation can be used for exploratory data analysis, preprocessing, data visualization, and information retrieval.

One way to discover low-dimensional structure is to convert high-dimensional data to low-dimensional codes by training a multi-layer neural network with a small central layer to reconstruct high-dimensional input vectors. Gradient descent can be used for fine-tuning the weights in such “autoencoder” networks, but this only works well if the initial weights are close to a good solution. In this talk we will describe an effective way of initializing the weights which allows deep autoencoder networks to learn low-dimensional codes that work much better than widely used Principal Components Analysis.

When trained on large document corpus, autoencoders are capable of extracting low-dimensional “semantic” codes that allow for much more accurate and faster retrieval than Latent Semantic Analysis, a well-known document retrieval method based on Singular Value Decomposition.