

Fast binary embeddings with Gaussian circulant matrices

In many applications, ranging from signal processing to machine learning, the encountered data is typically massive and often encoded in the form of high-dimensional vectors. Dimensionality reduction techniques help to reduce storage demand and decrease computation complexity by transforming high-dimensional vectors into low-dimensional ones while simultaneously retaining certain geometric information between vectors from the data set. In this talk, I will discuss efficient methods to further reduce the storage consumption of data sets, which are based on binarizing a Johnson-Lindenstrauss transform involving a Gaussian circulant matrix.