

High-Dimensional Vector Similarity Search

mercredi 12 juin 2024 14:30 (1h 30m)

Very large amounts of high-dimensional data are now omnipresent (ranging from traditional multidimensional data to time series and deep embeddings), and the performance requirements (i.e., response-time and accuracy) of a variety of applications that need to process and analyze these data have become very stringent and demanding. In the past years, high-dimensional similarity search has been studied in its many flavors. Similarity search algorithms for exact and approximate, one-off and progressive query answering. Approximate algorithms with and without (deterministic or probabilistic) quality guarantees. Solutions for on-disk and in-memory data, static and streaming data. Approaches based on multidimensional space-partitioning and metric trees, random projections and locality-sensitive hashing (LSH), product quantization (PQ) and inverted files, k-nearest neighbor graphs and optimized linear scans. Surprisingly, the work on data-series (or time-series) similarity search has recently been shown to achieve the state-of-the-art performance for several variations of the problem, on both time-series and general high-dimensional vector data. In this talk, we will touch upon the different aspects of this interesting story, and present some of the state-of-the-art solutions.

Orateur: PALPANAS, Themis (LIPADE - Paris Descartes University)