Dimensionality Reduction: Theoretical Analysis of Practical Measures

Nova Fandina, Hebrew University of Jerusalem

Metric embedding has been an extremely fruitful and influential field within Mathematics and Theoretical Computer Science over the past few decades, where the vast majority of work has been
devoted to analyzing the worst case behavior of embeddings. Metric embedding is also an important tool in real-world applications within many fields. Yet, in nearly all practical applications the fundamental criterion for measuring the quality of the embedding is its average performance over all pairs, where the measure of quality per pair is often the distortion, the square distortion, and similar related notions. Some important and vastly used criteria of this type originated within the well-known Multi-Dimensional Scaling framework. Yet, no theoretical studies have thus far attempted at providing
rigorous analysis of these criteria. We make the first step in this direction while focusing on the
fundamental task of dimensionality reduction, where the target is k-dimensional Euclidean space. We provide almost tight upper and lower bounds on the q-moments of each measure. In particular, we establish ties between the most common practical measurement criteria and show that various techniques, developed within the TCS/Math community, for sake of worst-case analysis, can be used to provide the first
rigorous analysis of these measures.

This is joint work with Yair Bartal and Nova Fandina,