# abstract

COMPUTER SCIENCE/DISCRETE MATH, I
Topic:

Speaker:

Affiliation:

Date:

Time/Room:

A low-distortion embedding between two metric spaces is a mapping which preserves the distances between each pair of points, up to a small factor called distortion. Low-distortion embeddings have recently found numerous applications in computer science.

Most of the known embedding results are "absolute", that is, of the form: any metric Y from a given class of metrics C can be embedded into a metric X with low distortion c. This is beneficial if one can guarantee low distortion for all metrics Y in C. However, in many situations, the worst-case distortion is too large to be meaningful. For example, if X is a line metric, then even very simple metrics (an n-point star or an n-point cycle) are embeddable into X only with distortion linear in n. Nevertheless, embeddings into the line (or into low-dimensional spaces) are important for many applications.

A solution to this issue is to consider "relative" (or "approximation") embedding problems, where the goal is to design an ($a$-approximation) algorithm which, given any metric X from C as an input, finds an embedding of X into Y which has distortion a*c_Y(X), where c_Y(X) is the best possible distortion of an embedding of X into Y. In this talk I will present several algorithms, as well as hardness results, for relative embedding problems.