abstract

Video of this lecture

COMPUTER SCIENCE AND DISCRETE MATHEMATICS SEMINAR II

Topic:

Speaker:

Affiliation:

Date:

Time/Room:

Finding the longest increasing subsequence (LIS) is a classic algorithmic problem. Simple O(n log n) algorithms, based on dynamic programming, are known for solving this problem exactly on arrays of length n.

In this talk I'll discuss recent work of C. Seshadhri and myself, in which we consider the problem of approximating the LIS in time sublinear in the input size. We develop a polylogarithmic time randomized algorithm that for any constant $c > 0$, outputs an approximation to the length of the LIS that (with high probability) is accurate to within an additive cn.
Previously, the best known polylogarithmic time algorithms could achieve only an additive $n/2$ approximation.

It is easy to see that approximation based on uniform random sampling of the input can give a very poor approximation to the LIS. Our algorithm uses random sampling that is both non-uniform and adaptive (that is the choice of points to sample depends on values observed previously).