

abstract

Members Seminar
Topic:

Speaker:

Affiliation:

Date:

Time/Room:

A fundamental theorem in linear algebra is that any real $n \times d$ matrix has a singular value decomposition (SVD). Several important numerical linear algebra problems can be solved efficiently once the SVD of an input matrix is computed: e.g. least squares regression, low rank approximation, and computing preconditioners, just to name a few.

Unfortunately in many modern big data applications the input matrix is very large, so that computing the SVD is computationally expensive.

Following a line of work of [Sarlós, 2006] and [Clarkson and Woodruff, 2013] on using dimensionality reduction techniques for solving numerical linear algebra problems, we show that several such problems with input matrices of size $n \times d$ ($n \gg d$) can be automatically transformed into a new instance of the problem that has size very close to $d \times d$, and thus can be solved much more quickly. Furthermore, executing these transformations require time only linear, or nearly linear, in the number of non-zero entries in the input matrix. Our techniques are of independent interest in random matrix theory, and the main technical contribution of our work turns out to be an analysis of the smallest and largest eigenvalues of certain random matrices.

This talk is based on joint work with Huy Lê Nguyen (Princeton).