On the Effectiveness of Single Vector Krylov Methods for Low-Rank Approximation

Christopher Musco
New York University
cmusco@nyu.edu

Krylov subspace methods are a ubiquitous tool for computing near-optimal rank k approximations of large matrices. While "large block" Krylov methods with block size at least k give the best known theoretical guarantees, block size one (a single vector) or a small constant is often preferred in practice. Despite their popularity, we lack theoretical bounds on the performance of such "small block" Krylov methods for low-rank approximation.

In this talk I will talk about a recent result that addresses this gap between theory and practice by proving that small block Krylov methods essentially match all known low-rank approximation guarantees for large block methods. Via a black-box reduction we show, for example, that the standard single vector Krylov method run for t iterations obtains the same spectral norm and Frobenius norm error bounds as a Krylov method with block size $\ell \geq k$ run for $O(t/\ell)$ iterations, up to a logarithmic dependence on the smallest gap between sequential singular values. That is, for a given number of matrix-vector products, single vector methods are essentially as effective as the best choice of large block size.

By combining our result with tail-bounds on eigenvalue gaps in random matrices, we prove that the dependence on the smallest singular value gap can be eliminated if the input matrix is perturbed by a small random matrix.

Joint work with Raphael Meyer (New York University) and Cameron Musco (University of Massachusetts Amherst).