

Completing any low-rank matrix, provably

Rachel Ward, University of Texas at Austin

We show that any n by n matrix of rank r can be exactly recovered from as few as $O(nr \log^2(n))$ randomly chosen elements, provided this random choice is made according to a specific biased distribution based on leverage scores of the underlying matrix. These results simplify and extend previous matrix completion results which rely on certain structural constraints, or incoherences, and the subset of elements is sampled uniformly at random. Our results are achieved by a new analysis based on matrix concentration bounds involving a matrix norm which computes the maximum of appropriately weighted row and column norms of the matrix. Finally, we show how our sampling methods and analysis give rise to new randomized algorithms for computing low-rank matrix approximations, and can provide tighter bounds in terms of the number of computations needed to achieve a certain accuracy in the matrix spectral norm. This is joint work with Yudong Chen, Srinadh Bhojanapalli and Sujay Sanghavi.