Finding a low-rank basis in a matrix subspace

Y. Nakatsukasa[1], T. Soma[2], and A. Uschmajew[3]

[1]University of Tokyo, nakatsukasa@mist.i.u-tokyo.ac.jp
[2]University of Tokyo, tasuku soma@mist.i.u-tokyo.ac.jp
[3]University of Bonn, uschmajew@ins.uni-bonn.de

For a given matrix subspace, how can we find a basis that consists of low-rank matrices? This problem is a generalization of the sparse vector problem. When the subspace is spanned by rank-one matrices, a solution is equivalent to a tensor CP decomposition. If the information on the rank-one basis is not given in advance, or if the space is spanned by matrices of higher rank, the situation is not as straightforward. By standard arguments from matroid theory, the problem can be in theory solved using a greedy algorithm. In this work we present a practical algorithm that mimics this greedy procedure. It finds basis elements one by another in two stages, by first estimating a minimal rank by applying soft singular value thresholding to a nuclear norm relaxation, and then computing a matrix with that rank using the method of alternating projections. Given the hardness of the problem, our method provides surprisingly reliable results in a number of experiments. Potential applications include data compression beyond the classical truncated SVD, computation of "low-rank" eigenvectors to simple or even multiple eigenvalues, image separation, and others.

References

[1] Y. Nakatsukasa, T. Soma, A. Uschmajew, *Finding a low-rank basis in a matrix subspace*. Preprint, 2015