Randomized Joint Diagonalization of Symmetric Matrices

Haoze He
École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
haoze.he@epfl.ch

Given a family of nearly commuting symmetric matrices, we consider the task of computing an orthogonal matrix that nearly diagonalizes every matrix in the family. In this work, we propose and analyze randomized joint diagonalization (RJD) for performing this task. RJD applies a standard eigenvalue solver to random linear combinations of the matrices. Unlike existing optimization-based methods, RJD is simple to implement and leverages existing high-quality linear algebra software packages. Our main novel contribution is to prove robust recovery: Given a family that is ϵ-close to a commuting family, RJD jointly diagonalizes this family, with high probability, up to an error of norm $\mathcal{O}(\epsilon)$. No other existing method is known to enjoy such a universal robust recovery guarantee. We also discuss how the algorithm can be further improved by deflation techniques and demonstrate its state-of-the-art performance by numerical experiments with synthetic and real-world data.

Joint work with Daniel Kressner (EPFL).