Randomized algorithms for trace estimation

Alice Cortinovis
École Polytechnique Fédérale de Lausanne, Lausanne (Switzerland)
alice.cortinovis@epfl.ch

Computing or estimating the trace of a large symmetric matrix A is a problem that arises in a wide variety of applications, such as Frobenius norm estimation, log determinant computations, triangle counting in graphs, and lattice quantum chromodynamics.

The Hutchinson's trace estimator is a popular technique to approximate the trace of a large-scale matrix A by computing the average of quadratic forms X^TAX for many samples of a random vector X. In this talk, we present new tail bounds that apply to symmetric indefinite matrices A for Rademacher and Gaussian random vectors [1].

The Hutch++ algorithm [2] is a more efficient trace estimation algorithm that reduces the variance of the Hutchinson's trace estimator by combining it with the randomized singular value decomposition, which obtains a low-rank approximation of A by multiplying the matrix with some random vectors. In this talk, we present an improved version of Hutch++ which aims at minimizing the computational cost – that is, the number of matrix-vector multiplications with A – needed to achieve a trace estimate with a target accuracy [3].

Joint work with D. Kressner *and* D. Persson.

References

[1] A. Cortinovis and D. Kressner, *On randomized trace estimates for indefinite matrices with an application to determinants*, Foundations of Computational Mathematics, (2021).

[2] R. A. Meyer, C. Musco, C. Musco, and D. P. Woodruff, *Hutch++: Optimal stochastic trace estimation*, In Symposium on Simplicity in Algorithms (SOSA), pages 142–55. SIAM (2021).

[3] D. Persson, A. Cortinovis, and D. Kressner, *Improved variants of the Hutch++ algorithm for trace estimation*, arXiv preprint arXiv:2109.10659, (2021).