Faster Spectral Density Estimation and Sparsification in the Nuclear Norm[1]

Yujia Jin YUJIAJIN@STANFORD.EDU
Stanford University
Ishani Karmarkar ISHANIK@STANFORD.EDU
Stanford University
Christopher Musco CMUSCO@NYU.EDU
New York University
Aaron Sidford SIDFORD@STANFORD.EDU
Stanford University
Apoorv Vikram Singh APOORV.SINGH@NYU.EDU
New York University

Editors: Shipra Agrawal and Aaron Roth

Abstract

We consider the problem of estimating the spectral density of a normalized graph adjacency matrix. Concretely, given an undirected graph $G = (V, E, w)$ with n nodes and positive edge weights $w \in \mathbb{R}^E_{>0}$, the goal is to return eigenvalue estimates $\widehat{\lambda}_1 \leq \cdots \leq \widehat{\lambda}_n$ such that

$$\frac{1}{n} \sum_{i \in \{1,\ldots,n\}} |\widehat{\lambda}_i - \lambda_i(N_G)| \leq \varepsilon,$$

where $\lambda_1(N_G) \leq \cdots \leq \lambda_n(N_G)$ are the eigenvalues of G's normalized adjacency matrix, N_G. This goal is equivalent to requiring that the Wasserstein-1 distance between the uniform distribution on $\lambda_1, \ldots, \lambda_n$ and the uniform distribution on $\widehat{\lambda}_1, \ldots, \widehat{\lambda}_n$ is less than ε.

We provide a randomized algorithm that achieves the guarantee above with $O(n\varepsilon^{-2})$ queries to a degree and neighbor oracle and in $O(n\varepsilon^{-3})$ time. This improves on previous state-of-the-art methods, including an $O(n\varepsilon^{-7})$ time algorithm from Braverman et al. (2022) and, for sufficiently small ε, a $2^{O(\varepsilon^{-1})}$ time method from Cohen-Steiner et al. (2018). To achieve this result, we introduce a new notion of graph sparsification, which we call *nuclear sparsification*. We provide an $O(n\varepsilon^{-2})$-query and $O(n\varepsilon^{-2})$-time algorithm for computing $O(n\varepsilon^{-2})$-sparse nuclear sparsifiers. We show that this bound is optimal in both its sparsity and query complexity, and we separate our results from the related notion of additive spectral sparsification. Of independent interest, we show that our sparsification method also yields the first *deterministic* algorithm for spectral density estimation that scales linearly with n (sublinear in the representation size of the graph).

Keywords: spectral density estimation, spectral sparsification, sublinear algorithm

References

Vladimir Braverman, Aditya Krishnan, and Christopher Musco. Sublinear time spectral density estimation. In *Proceedings of the 54th Annual ACM Symposium on Theory of Computing (STOC)*, pages 1144–1157, 2022.

David Cohen-Steiner, Weihao Kong, Christian Sohler, and Gregory Valiant. Approximating the spectrum of a graph. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1263–1271, 2018.

1. Extended abstract. Full version appears as [arXiv 2406.07521, v1]