Near-Optimal Entrywise Sampling of Numerically Sparse Matrices

Vladimir Braverman*
Johns Hopkins University
vova@cs.jhu.edu

Robert Krauthgamer†
Weizmann Institute of Science
robert.krauthgamer@weizmann.ac.il

Aditya Krishnan‡
Johns Hopkins University
akrish23@jhu.edu

Shay Sapir
Weizmann Institute of Science
shay.sapir@weizmann.ac.il

Abstract

Many real-world data sets are sparse or almost sparse. One method to measure this for a matrix $A \in \mathbb{R}^{n\times n}$ is the *numerical sparsity*, denoted $\mathsf{ns}(A)$, defined as the minimum $k \geq 1$ such that $\|a\|_1/\|a\|_2 \leq \sqrt{k}$ for every row and every column a of A. This measure of a is smooth and is clearly only smaller than the number of non-zeros in the row/column a.

The seminal work of Achlioptas and McSherry [2007] has put forward the question of approximating an input matrix A by entrywise sampling. More precisely, the goal is to quickly compute a sparse matrix $\tilde{A}$ satisfying $\|A - \tilde{A}\|_2 \leq \epsilon\|A\|_2$ (i.e., additive spectral approximation) given an error parameter $\epsilon > 0$. The known schemes sample and rescale a small fraction of entries from A.

We propose a scheme that sparsifies an almost-sparse matrix A — it produces a matrix $\tilde{A}$ with $O(\epsilon^{-2}\mathsf{ns}(A) \cdot n \ln n)$ non-zero entries with high probability. We also prove that this upper bound on $\mathsf{nnz}(\tilde{A})$ is *tight* up to logarithmic factors. Moreover, our upper bound improves when the spectrum of A decays quickly (roughly replacing n with the stable rank of A). Our scheme can be implemented in time $O(\mathsf{nnz}(A))$ when $\|A\|_2$ is given. Previously, a similar upper bound was obtained by Achlioptas et al. [2013] but only for a restricted class of inputs that does not even include symmetric or covariance matrices. Finally, we demonstrate two applications of these sampling techniques, to faster approximate matrix multiplication, and to ridge regression by using sparse preconditioners.

References

D. Achlioptas and F. McSherry. Fast computation of low-rank matrix approximations. *Journal of the ACM (JACM)*, 54(2):9–es, 2007.

D. Achlioptas, Z. S. Karnin, and E. Liberty. Near-optimal entrywise sampling for data matrices. In *Advances in Neural Information Processing Systems*, pages 1565–1573, 2013.

*This research was supported in part by NSF CAREER grant 1652257, NSF grant 1934979, ONR Award N00014-18-1-2364 and the Lifelong Learning Machines program from DARPA/MTO.

†Work partially supported by ONR Award N00014-18-1-2364, the Israel Science Foundation grant #1086/18, and a Minerva Foundation grant.

‡This research was supported in part by NSF CAREER grant 1652257, ONR Award N00014-18-1-2364 and the Lifelong Learning Machines program from DARPA/MTO.