

Abstract

We study an iterative matrix conditioning algorithm due to Osborne (1960). The goal of the algorithm is to convert a square matrix into a *balanced* matrix where every row and corresponding column have the same norm. The original algorithm was proposed for balancing rows and columns in the L_2 norm, and it works by iterating over balancing a row-column pair in fixed round-robin order. Variants of the algorithm for other norms have been heavily studied and are implemented as standard preconditioners in many numerical linear algebra packages. Recently, Schulman and Sinclair (2015), in a first result of its kind for any norm, analyzed the rate of convergence of a variant of Osborne’s algorithm that uses the L_∞ norm and a different order of choosing row-column pairs. In this paper we study matrix balancing in the L_1 norm and other L_p norms. We show the following results for any matrix $A = (a_{ij})_{i,j=1}^n$, resolving in particular a main open problem mentioned by Schulman and Sinclair.

1. We analyze the iteration for the L_1 norm under a greedy order of balancing. We show that it converges to an ϵ -balanced matrix in $K = O(\min\{\epsilon^{-2} \log w, \epsilon^{-1} n^{3/2} \log(w/\epsilon)\})$ iterations that cost a total of $O(m + Kn \log n)$ arithmetic operations over $O(n \log(w/\epsilon))$ -bit numbers. Here m is the number of non-zero entries of A , and $w = \sum_{i,j} |a_{ij}|/a_{\min}$ with $a_{\min} = \min\{|a_{ij}| : a_{ij} \neq 0\}$.
2. We show that the original round-robin implementation converges to an ϵ -balanced matrix in $O(\epsilon^{-2} n^2 \log w)$ iterations totaling $O(\epsilon^{-2} mn \log w)$ arithmetic operations over $O(n \log(w/\epsilon))$ -bit numbers.
3. We show that a random implementation of the iteration converges to an ϵ -balanced matrix in $O(\epsilon^{-2} \log w)$ iterations using $O(m + \epsilon^{-2} n \log w)$ arithmetic operations over $O(\log(wn/\epsilon))$ -bit numbers.
4. We demonstrate a lower bound of $\Omega(1/\sqrt{\epsilon})$ on the convergence rate of any implementation of the iteration.
5. We observe, through a known trivial reduction, that our results for L_1 balancing apply to any L_p norm for all finite p , at the cost of increasing the number of iterations by only a factor of p .

We note that our techniques are very different from those used by Schulman and Sinclair.