Fast Monte-Carlo Algorithms for Matrix Multiplication

Petros Drineas *   Ravi Kannan†

August 22, 2000

Abstract

In certain applications (i.e. information retrieval), the data consists of an $m \times n$ matrix A and it is of interest to compute (or approximate) $A^T \cdot A$, with m and n larger than 10^4.

We present two relatively simple randomized algorithms that compute a "description" to an approximation P to $A^T \cdot A$. The approximation is "described" as a sum of rank-one matrices, namely $P = \sum_{k=1}^{s} r^{(k)} {r^{(k)}}^T$, where $r^{(k)} \in \mathcal{R}^n$ and s is some constant. The first algorithm explicitly computes all $r^{(k)}$, while the second algorithm "describes" the $r^{(k)}$ as a matrix-vector product. If we allow a preprocessing time equal to reading all non-zero elements of A once, the first algorithm runs in $O(n)$ time, while the second algorithm runs in O(1) time. We measure the error of the approximation by giving a bound for $e = \|A^T \cdot A - P\|_F$, namely $e = O(\frac{1}{\sqrt{s}}\|A\|_F^2)$. We remind that $\|A\|_F^2 = \sum_{i,j} A_{ij}^2$.

The above algorithms can be generalized to compute the product of any two matrices. They are based on some older results of ours that approximate the singular value decomposition of a matrix. Also, a more straight-forward algorithm that approximates the product of two matrices is described.

*Department of Computer Science, Yale University, New Haven, CT 06520, drineas@cs.yale.edu.

†Department of Computer Science, Yale University, New Haven, CT 06520, kannan@cs.yale.edu.