On the Rigidity of Vandermonde Matrices

Satyanarayana V. Lokam

Abstract

The rigidity function $\mathcal{R}_A(r)$ of a matrix A is the minimum number of entries of A that must be changed to reduce the rank of A to less than or equal to r. While almost all matrices have rigidity close to $(n-r)^2$, proving strong lower bounds on the rigidity of explicit matrices is a fundamental open question with several consequences in complexity theory.

A natural class of matrices expected to have high rigidity is that of Vandermonde matrices $V = (x_i^{j-1})_{1 \le i,j \le n}$. However, even when the x_i are algebraically independent, it was not known if $\mathcal{R}_V(r) = \Omega(n^2)$ for nonconstant r. We prove that if $r = O(\sqrt{n})$, then $\mathcal{R}_V(r) = \Omega(n^2)$, when the x_i are algebraically independent. Although not explicit, this provides a natural, n-dimensional manifold in the space of $n \times n$ matrices with $\Omega(n^2)$ rigidity for nonconstant r. Our proof is based on a technique due to Shoup and Smolensky.

For explicit Vandermonde matrices, the best known lower bound is $\mathcal{R}_V(r) = \Omega(n^2/r \ \log(n/r))$, which follows from a result of Shokrollahi, Spielman, and Stemann.