NP-hardness of LAROS

Xuan Vinh Doan and Stephen Vavasis

2013 March 13

This writeup is an unpublished supplement to "Finding approximately rank-one submatrices with the nuclear norm and ℓ_1-norm" by the same authors. Refer to the main paper for the background and bibliography. For this writeup, $\|\cdot\|$ denotes the spectral norm. We define the LAROS problem as follows: Given an $m \times n$ matrix A and parameter $\theta > 0$, find $I \subset \{1, \ldots, m\}$, $J \subset \{1, \ldots, n\}$ to minimize $\|A(I, J)\|^{-1} + \theta|I| \cdot |J|$.

Theorem 1 *The LAROS problem is NP-hard.*

Proof. We start with the following preliminary claim, which is proved via a line of calculus.

Claim 1: The function $f(x) = x^{-1/2} + \theta x$ defined on $\{x \in \mathbf{R} : x > 0\}$ is strictly convex and has a unique minimizer at $x = (2\theta)^{-2/3}$.

The next claim follows from the fact that $\|H\| \leq \|H\|_F$.

Claim 2: Let H be an $M \times N$ matrix of all 0's and 1's. Then $\|H\|^{-1} \geq (MN)^{-1/2}$ with equality achieved if and only if H is all 1's.

Now consider the following problem: Given a bipartite graph $G = (U, V, E)$ and integer e, does G contain a biclique with a total of e edges? This problem is known to be NP-hard according to Peeters. We can reduce this problem to LAROS as follows. Given (U, V, E) and e, choose θ so that $(2\theta)^{-2/3} = e$. Let A be the 0-1 matrix that encodes the (U, V) relationship of G. Consider minimizing $\|A(I, J)\|^{-1} + \theta|I| \cdot |J|$. Suppose $I \subset \{1, \ldots, m\}$, $J \subset \{1, \ldots, n\}$ and $|I| \cdot |J| = x$. Then we have the following inequalities:

$$\begin{aligned} \|A(I,J)\|^{-1} + \theta|I| \cdot |J| &= \|A(I,J)\|^{-1} + \theta x \\ &\geq x^{-1/2} + \theta x \\ &\geq e^{-1/2} + \theta e, \end{aligned}$$

where the second line follows from Claim 2 and the third line follows from Claim 1. Furthermore, according to the claims, these two inequalities are tight iff $x = e$ and $A(I, J)$ is all 1's. Therefore, the minimum value of the objective function is $e^{-1/2} + \theta e$ if and only if a biclique with e edges exists in G.

There are two technicalities that should be addressed. The first is that the parameter θ, which is an input to the LAROS problem, would generally be irrational according to our construction, but usual Turing Machine formality requires it to be a rational number with a polynomially bounded number of bits. The proof can be modified by noting that it is not necessary for e to be the unique optimizer of the objective function $f(x)$ over all positive reals; it suffices for e to be the unique optimizer over all positive integers. In this case, there is an interval of θ's that establishes the same result, and a rational value with a polynomially bounded number of bits can be found in this interval. Details are left to the reader.

A second issue is that the above proof seems to show the NP-hardness of exactly solving LAROS, but in practice it could be solved only approximately since the solution is an irrational number. It is possible to extend the proof to show NP-hardness even of an approximate solution to LAROS by noting that Claim 2 can be strengthened. The stronger version states that if H is a 0-1 matrix that is not the matrix of all 1's, then $\|H\|^{-1} \geq (MN-1)^{-1/2}$. This means that there is a discrete gap between the optimizer of LAROS in the case that the biclique exists versus the case that it does not. Thus, any approximate solution close enough to the optimizer smaller than this gap still solves the max biclique instance Therefore, approximate solution to LAROS is also NP-hard. Details are left to the reader.