Abstract

We show that, assuming the (deterministic) Exponential Time Hypothesis, distinguishing between a graph with an induced k-clique and a graph in which all k-subgraphs have density at most $1 - \epsilon$, requires $n^{\tilde{\Omega}(\log n)}$ time. Our result essentially matches the quasi-polynomial algorithms of Feige and Seltser [FS97] and Barman [Bar14] for this problem, and is the first one to rule out an additive PTAS for Densest k-Subgraph. We further strengthen this result by showing that our lower bound continues to hold when, in the soundness case, even subgraphs smaller by a near-polynomial factor ($k' = k \cdot 2^{-\tilde{\Omega}(\log n)}$) are assumed to be at most $(1 - \epsilon)$-dense. Our reduction is inspired by recent applications of the "birthday repetition" technique [AIM14, BKW15]. Our analysis relies on information theoretical machinery and is similar in spirit to analyzing a parallel repetition of two-prover games in which the provers may choose to answer some challenges multiple times, while completely ignoring other challenges.