Abstract

Properties definable in first-order logic are algorithmically interesting for both theoretical and pragmatic reasons. Many of the most studied algorithmic problems, such as Hitting Set and Orthogonal Vectors, are first-order, and the first-order properties naturally arise as relational database queries. A relatively straightforward algorithm for evaluating a property with $k+1$ quantifiers takes time $O(m^k)$ and, assuming the Strong Exponential Time Hypothesis (SETH), some such properties require $O(m^{k-\epsilon})$ time for any $\epsilon > 0$. (Here, m represents the size of the input structure, i.e. the number of tuples in all relations.) We give algorithms for every first-order property that improves this upper bound to $m^k/2^{\Theta(\sqrt{\log n})}$, i.e., an improvement by a factor more than any poly-log, but less than the polynomial required to refute SETH. Moreover, we show that further improvement is *equivalent* to improving algorithms for sparse instances of the well-studied Orthogonal Vectors problem. Surprisingly, both results are obtained by showing completeness of the Sparse Orthogonal Vectors problem for the class of first-order properties under fine-grained reductions. To obtain improved algorithms, we apply the fast Orthogonal Vectors algorithm of (A. Abboud, R. Williams, and H. Yu. *More applications of the polynomial method to algorithm design.* In Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 218-230. SIAM, 2015), (T. M. Chan and R. Williams. *Deterministic APSP, Orthogonal Vectors, and More: Quickly derandomizing Razborov-Smolensky.* In Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1246-1255. SIAM, 2016). While fine-grained reductions (reductions that closely preserve the conjectured complexities of problems) have been used to relate the hardness of disparate specific problems both within P and beyond, this is the first such completeness result for a standard complexity class.