



Statistical Algorithms and a Lower Bound for Detecting Planted Cliques

Vitaly Feldman Elena Grigorescu*[†] Lev Reyzin^{‡†}

Santosh S. Vempala[†] Ying Xiao[†]

vitaly@post.harvard.edu
Almaden Research Center
IBM
San Jose, CA 95120

elena-g@purdue.edu
Department of Computer Science
Purdue University
West Lafayette, IN 47907

lreyzin@math.uic.edu
Department of Mathematics, Statistics, and Computer Science
University of Illinois at Chicago
Chicago, IL 60607

{vempala,yxiao32}@cc.gatech.edu
School of Computer Science
Georgia Institute of Technology
Atlanta, GA 30332

*This material is based upon work supported by the National Science Foundation under Grant #1019343 to the Computing Research Association for the CIFellows Project.

[†]Research supported in part by NSF awards CCF-0915903 and CCF-1217793.

[‡]Research supported by a Simons Postdoctoral Fellowship.

Abstract

We introduce a framework for proving lower bounds on computational problems over distributions, based on a class of algorithms called *statistical algorithms*. For such algorithms, access to the input distribution is limited to obtaining an estimate of the expectation of any given function on a sample drawn randomly from the input distribution, rather than directly accessing samples. Most natural algorithms of interest in theory and in practice, e.g., moments-based methods, local search, standard iterative methods for convex optimization, MCMC and simulated annealing, are statistical algorithms or have statistical counterparts. Our framework is inspired by and generalizes the statistical query model in learning theory [34].

Our main application is a nearly optimal lower bound on the complexity of *any* statistical algorithm for detecting planted bipartite clique distributions (or planted dense subgraph distributions) when the planted clique has size $O(n^{1/2-\delta})$ for any constant $\delta > 0$. Variants of these problems have been assumed to be hard to prove hardness for other problems and for cryptographic applications. Our lower bounds provide concrete evidence of hardness, thus supporting these assumptions.

1 Introduction

We study the complexity of problems where the input consists of independent samples from an unknown distribution. Such problems are at the heart of machine learning and statistics (and their numerous applications) and also occur in many other contexts such as compressed sensing and cryptography. Many methods exist to estimate the sample complexity of such problems (e.g. VC dimension [48] and Rademacher complexity [6]). Proving lower bounds on the computational complexity of these problems has been much more challenging. The traditional approach to this is via reductions and finding distributions that can generate instances of some problem conjectured to be intractable (e.g., assuming $\text{NP} \neq \text{RP}$).

Here we present a different approach, namely showing that a broad class of algorithms, which we refer to as *statistical algorithms*, have high complexity, unconditionally. Our definition encompasses most algorithmic approaches used in practice and in theory on a wide variety of problems, including Expectation Maximization (EM) [17], local search, MCMC optimization [45, 25], simulated annealing [36, 49], first and second order methods for linear/convex optimization, [18, 7], k -means, Principal Component Analysis (PCA), Independent Component Analysis (ICA), Naïve Bayes, Neural Networks and many others (see [14] and [10] for proofs and many other examples). In fact, we are aware of only one algorithm that does not have a statistical counterpart: Gaussian elimination for solving linear equations over a field (e.g., mod 2).

Informally, statistical algorithms can access the input distribution only by asking for the value of any bounded real-valued function on a random sample from the distribution, or the average value of the function over a specified number of independent random samples. As an example, suppose we are trying to solve $\min_{u \in U} \mathbf{E}_{x \sim D}[f(x, u)]$ by gradient descent. Then the gradient of the objective function is (by interchanging the derivative with the integral)

$$\nabla_u \mathbf{E}_x[f(x, u)] = \mathbf{E}_x[\nabla_u(f(x, u))]$$

and can be estimated from samples; thus the algorithm can proceed without ever examining samples directly. A more involved example is Linear Programming. One version is the feasibility problem: find a nonzero w s.t. $a \cdot w \geq 0$ for all a in some set A . We can formulate this as

$$\max_w \mathbf{E}_a[\text{sign}(a \cdot w)]$$

and the distribution over a could be uniform over the set A if it is finite. This can be solved by a statistical algorithm [11, 18]. This is also the case for semidefinite programs and for conic optimization [7]. The key motivation for our definition of statistical algorithms is the empirical observation that almost all algorithms that work on explicit instances are already statistical in our sense or have natural statistical counterparts. Thus, proving lower bounds for statistical algorithms strongly indicates the need for new approaches even for explicit instances. We present the formal oracle-based definitions of statistical algorithms in Section 2.

The inspiration for our model is the *statistical query (SQ)* model in learning theory [34] defined as a restriction of Valiant's PAC learning model [46]. The primary goal of the restriction was to simplify the design of noise-tolerant learning algorithms. As was shown by Kearns and others in subsequent works, almost all classes of functions that can be learned efficiently can also be efficiently learned in the restricted SQ model. A notable and so far the only exception is the algorithm for learning parities, based on Gaussian elimination. As was already shown by Kearns [34], parities

require exponential time to learn in the SQ model. Further, Blum et al. [12] proved that the number of SQs required for weak learning (that is, for obtaining a non-negligible advantage over the random guessing) of a class of functions C over a fixed distribution D is characterized by a combinatorial parameter of C and D , referred to as SQ-DIM(C, D), the statistical query dimension.

Our notion of statistical algorithms generalizes SQ learning algorithms to any computational problem over distributions. For any problem over distributions we define a parameter of the problem that lower bounds the complexity of solving the problem by any statistical algorithm in the same way as SQ-DIM lower bounds the complexity of learning in the SQ model. Our techniques for proving lower bounds on statistical algorithms are also based on methods developed for lower-bounding the complexity of SQ algorithms. However, as we will describe later, they depart from the known techniques in a number of significant ways that are necessary for our more general definition and our applications.

We apply our techniques to the problem of detecting planted bipartite cliques/dense subgraphs. We now define these problems precisely and give some background.

Detecting Planted Cliques. In the planted clique problem, we are given a graph G whose edges are generated by starting with a random graph $G_{n,1/2}$, then “planting,” i.e., adding edges to form a clique on k randomly chosen vertices. Jerrum [30] and Kučera [37] introduced the planted clique problem as a potentially easier variant of the classical problem of finding the largest clique in a random graph [33]. A random graph $G_{n,1/2}$ contains a clique of size $2 \log n$ with high probability, and a simple greedy algorithm can find one of size $\log n$. Finding cliques of size $(2 - \epsilon) \log n$ is a hard problem for any $\epsilon > 0$. Planting a larger clique should make it easier to find one. The problem of finding the smallest k for which the planted clique can be detected in polynomial time has attracted significant attention. For $k \geq c\sqrt{n \log n}$, simply picking vertices of large degrees suffices [37]. Cliques of size $k = \Omega(\sqrt{n})$ can be found using spectral methods [2, 38, 15], via SDPs [20], nuclear norm minimization [3] or combinatorial methods [22, 16].

While there is no known polynomial-time algorithm that can detect cliques of size below the threshold of $\Omega(\sqrt{n})$, there is a quasipolynomial algorithm for any $k \geq 2 \log n$: enumerate subsets of size $2 \log n$; for each subset that forms a clique, take all common neighbors of the subset; one of these will be the planted clique. This is also the fastest known algorithm for any $k = O(n^{1/2-\delta})$, where $\delta > 0$.

Some evidence of the hardness of the problem was shown by Jerrum [30] who proved that a specific approach using a Markov chain cannot be efficient for $k = o(\sqrt{n})$. More evidence of hardness is given in [21], where it is shown that Lovász-Schrijver SDP relaxations, which include the SDP used in [20], cannot be used to efficiently find cliques of size $k = o(\sqrt{n})$. The problem has been used to generate cryptographic primitives [31], and as a hardness assumption [1, 28, 40].

We focus on the bipartite planted clique problem, where a bipartite $k \times k$ clique is planted in a random bipartite graph. A densest-subgraph version of the bipartite planted clique problem has been used as a hard problem for cryptographic applications [4]. All known bounds and algorithms for the k -clique problem can be easily adapted to the bipartite case. Therefore it is natural to suspect that new upper bounds on the planted k -clique problem would also yield new upper bounds for the bipartite case.

The starting point of our investigation for this problem is the property of the bipartite planted k -clique problem is that it has an equivalent formulation as a problem over distributions defined as follows.

Problem 1.1. Fix an integer k , $1 \leq k \leq n$, and a subset of k indices $S \subseteq \{1, 2, \dots, n\}$. The input distribution D_S on vectors $x \in \{0, 1\}^n$ is defined as follows: with probability $1 - (k/n)$, x is uniform over $\{0, 1\}^n$; and with probability k/n the k coordinates of x in S are set to 1, and the rest are uniform over their support. For an integer t , the **distributional planted k -biclique** problem with t samples is the problem of finding the unknown subset S using t samples drawn randomly from D_S .

One can view samples x_1, \dots, x_t as adjacency vectors of a bipartite graph as follows: the bipartite graph has n vertices on the right (with k marked as members of the clique) and t vertices on the left. Each of the t samples gives the adjacency vector of the corresponding vertex on the left. It is not hard to see that for $t = n$, conditioned on the event of getting exactly k samples with planted indices, we will get a random bipartite graph with a planted k -biclique (we prove the equivalence formally in Appendix A).

Another interesting approach for planted clique was proposed by Frieze and Kannan [24]. They gave a reduction from finding a planted clique in a random graph to finding a direction that maximizes a tensor norm; this was extended to general r in [13]. Specifically, they show that maximizing the r 'th moment (or the 2-norm of an r 'th order tensor) allows one to recover planted cliques of size $\tilde{\Omega}(n^{1/r})$. A related approach is to maximize the 3rd or higher moment of the distribution given by the distributional planted clique problem. For this approach it is natural to consider the following type of optimization algorithm: start with some unit vector u , then estimate the gradient at u (via samples), move along that direction and return to the sphere; repeat to reach an approximate local maximum. Unfortunately, over the unit sphere, the expected r 'th moment function can have (exponentially) many local maxima even for simple distributions. A more sophisticated approach [32] is through Markov chains or simulated annealing; it attempts to sample unit vectors from a distribution on the sphere which is heavier on vectors that induce a higher moment, e.g., u is sampled with density proportional to $e^{f(u)}$ where $f(u)$ is the expected r 'th moment along u . This could be implemented by a Markov chain with a Metropolis filter [39, 27] ensuring a proportional steady state distribution. If the Markov chain were to mix rapidly, that would give an efficient approximation algorithm because sampling from the steady state likely gives a vector of high moment. At each step, all one needs is to be able to estimate $f(u)$, which can be done by sampling from the input distribution.

As we will see presently, these statistical approaches will have provably high complexity. For the bipartite planted clique problem, statistical algorithms that can query expectations of arbitrary functions to within a small tolerance need $n^{\Omega(\log n)}$ queries to detect planted cliques of size $k < n^{\frac{1}{2}-\delta}$ for any $\delta > 0$. Even stronger exponential bounds apply for the more general problem of detecting planted dense subgraphs of the same size. These bounds match the current upper bounds. To describe these results precisely and discuss exactly what they mean for the complexity of these problems, we will need to define the notion of statistical algorithms, the complexity measures we use, and our main tool for proving lower bounds, a notion of statistical dimension of a set of distributions. We do this in the next section. In Section 3 we prove our general lower bound results and in Section 5 we estimate the statistical dimension of detecting planted bicliques and dense subgraphs.

2 Definitions and Overview

Here we formally define statistical algorithms and the key notion of statistical dimension, and then describe the resulting lower bounds in detail.

2.1 Problems over Distributions

We begin by formally defining the class of problems addressed by our framework.

Definition 2.1 (Search problems over distributions). *For a domain X , let \mathcal{D} be a set of distributions over X , let \mathcal{F} be a set of solutions and $\mathcal{Z} : \mathcal{D} \rightarrow 2^{\mathcal{F}}$ be a map from a distribution $D \in \mathcal{D}$ to a subset of solutions $\mathcal{Z}(D) \subseteq \mathcal{F}$ that are defined to be valid solutions for D . For $t > 0$ the distributional search problem \mathcal{Z} over \mathcal{D} and \mathcal{F} using t samples is to find a valid solution $f \in \mathcal{Z}(D)$ given access to t random samples from an unknown $D \in \mathcal{D}$.*

We note that this definition captures decision problems by having $\mathcal{F} = \{0, 1\}$. With slight abuse of notation, for a solution $f \in \mathcal{F}$, we denote by $\mathcal{Z}^{-1}(f)$ the set of distributions in \mathcal{D} for which f is a valid solution.

It is important to note that the number of available samples t can have a major influence on the complexity of the problem. First, for most problems there is a minimum t for which the problem is information-theoretically solvable. This value is often referred to as the *sample complexity* of the problem. But even for t which is larger than the sample complexity of the problem, having more samples can make the problem easier computationally. For example, in the context of attribute-efficient learning, there is a problem that is intractable with few samples (under cryptographic assumptions) but is easy to solve with a larger (but still polynomial) number of samples [43]. Our distributional planted biclique problem exhibits the same phenomenon.

2.2 Statistical Algorithms

The statistical query learning model of Kearns [34] is a restriction of the PAC model [46]. It captures algorithms that rely on empirical estimates of statistical properties of random examples of an unknown function instead of individual random examples (as in the PAC model of learning).

In the same spirit, for general search, decision and optimization problems over a distribution, we define statistical algorithms as algorithms that do not see samples from the distribution but instead have access to estimates of the expectation of any bounded function of a sample from the distribution.

Definition 2.2 (STAT oracle). *Let D be the input distribution over the domain X . For a tolerance parameter $\tau > 0$, $\text{STAT}(\tau)$ oracle is the oracle that for any query function $h : X \rightarrow [-1, 1]$, returns a value*

$$v \in [E_{x \sim D}[h(x)] - \tau, E_{x \sim D}[h(x)] + \tau].$$

The general algorithmic techniques mentioned earlier can all be expressed as algorithms using STAT oracle instead of samples themselves, in most cases in a straightforward way. We would also like to note that in the PAC learning model some of the algorithms, such as the Perceptron algorithm, did not initially appear to fall into the SQ framework but SQ analogues were later found for all known learning techniques except Gaussian elimination (for specific examples, see [34] and

[11]). We expect the situation to be similar even in the broader context of search problems over distributions.

The most natural realization of $\text{STAT}(\tau)$ oracle is one that computes h on $O(1/\tau^2)$ random samples from D and returns their average. Chernoff’s bound would then imply that the estimate is within the desired tolerance (with constant probability). However, if $h(x)$ is very biased (e.g. equal to 0 with high probability), it can be estimated with fewer samples. Our primary application requires a tight bound on the number of samples necessary to solve a problem over distributions. Therefore we define a stronger version of STAT oracle in which the tolerance is adjusted for the variance of the query function, that is the oracle returns the expectation to within the same tolerance (up to constant factors) as one expects to get from t samples. More formally, for a Boolean query h , $\text{VSTAT}(t)$ can return any value v for which the distribution $B(t, v)$ (sum of t independent Bernoulli variables with bias v) is statistically close to $B(t, E[h])$ (see Sec. 3.2 for more details on this correspondence).

Definition 2.3 (VSTAT oracle). *Let D be the input distribution over the domain X . For a sample size parameter $t > 0$, $\text{VSTAT}(t)$ oracle is the oracle that for any query function $h : X \rightarrow \{0, 1\}$, returns a value $v \in [p - \tau, p + \tau]$, where $p = E_{x \sim D}[h(x)]$ and $\tau = \max \left\{ \frac{1}{t}, \sqrt{\frac{p(1-p)}{t}} \right\}$.*

Note that $\text{VSTAT}(t)$ always returns the value of the expectation within $1/\sqrt{t}$. Therefore it is stronger than $\text{STAT}(1/\sqrt{t})$ but weaker than $\text{STAT}(1/t)$. (STAT , unlike VSTAT , allows non-boolean functions but this is not a significant difference as any $[-1, 1]$ -valued query can be converted to a logarithmic number of $\{0, 1\}$ -valued queries).

The STAT and VSTAT oracles we defined can return any value within the given tolerance and therefore can make adversarial choices. We also aim to prove lower bounds against algorithms that use a potentially more benign, “unbiased” statistical oracle. The unbiased statistical oracle gives the algorithm the true value of a boolean query function on a randomly chosen sample. This model is based on the Honest SQ model in learning by Yang [50] (which itself is based on an earlier model of Jackson [29]).

Definition 2.4 (1-STAT oracle). *Let D be the input distribution over the domain X . The 1-STAT oracle is the oracle that given any function $h : X \rightarrow \{0, 1\}$, takes an independent random sample x from D and returns $h(x)$.*

Note that the 1-STAT oracle draws a fresh sample upon each time it is called. Without re-sampling each time, the answers of the 1-STAT oracle could be easily used to recover any sample bit-by-bit, making it equivalent to having access to random samples. The **query complexity** of an algorithm using 1-STAT is defined to be the number of calls it makes to the 1-STAT oracle. Note that the 1-STAT oracle can be used to simulate VSTAT (with high probability) by taking the average of $O(t)$ replies of 1-STAT for the same function h . While it might seem that access to 1-STAT gives an algorithm more power than access to VSTAT we will show that t samples from 1-STAT can be simulated using access to $\text{VSTAT}(O(t))$. This will allow us to translate our lower bounds on algorithms with access to VSTAT to query complexity lower bounds for unbiased statistical algorithms.

In the following discussion, we refer to algorithms using STAT or VSTAT oracles (instead of samples) as **statistical algorithms**. Algorithms using the 1-STAT oracle are henceforth called **unbiased statistical algorithms**.

2.3 Statistical Dimension

The main tool in our analysis is an information-theoretic bound on the complexity of statistical algorithms. Our definitions originate from the statistical query (SQ) dimension in learning theory [12] used to characterize SQ learning algorithms. Roughly speaking, the SQ dimension corresponds to the number of nearly uncorrelated labeling functions in a class (see Appendix B for the details of the definition and the relationship to our bounds).

We introduce a natural generalization and strengthening of this approach to search problems over arbitrary sets of distributions and prove lower bounds on the complexity of statistical algorithms based on the generalized notion. Our definition departs from SQ dimension in three aspects. (1) Our notion applies to any set of distributions whereas in the learning setting all known dimensions require fixing the distribution over the domain and only allow varying the labeling function. (2) Instead of relying on a bound on pairwise correlations, our dimension relies on a bound on average correlations in a large set of distributions. This weaker condition allows us to derive the tight bounds on the complexity of statistical algorithms for the planted k -biclique problem. (3) We show that our dimension also gives lower bounds for the stronger VSTAT oracle (without incurring a quadratic loss in the sample size parameter).

We now define our dimension formally. For two functions $f, g : X \rightarrow \mathcal{R}$ and a distribution D with probability density function $D(x)$, the inner product of f and g over D is defined as

$$\langle f, g \rangle_D \doteq \mathbf{E}_{x \sim D} [f(x)g(x)].$$

The norm of f over D is $\|f\|_D = \sqrt{\langle f, f \rangle_D}$. We remark that, by convention, the integral from the inner product is taken only over the support of D , i.e. for $x \in X$ such that $D(x) \neq 0$. Given a distribution D over X let $D(x)$ denote the probability density function of D relative to some fixed underlying measure over X (for example uniform distribution for discrete X or Lebesgue measure over \mathcal{R}^n). Our bound is based on the inner products between functions of the following form: $(D'(x) - D(x))/D(x)$ where D' and D are distributions over X . For this to be well-defined, we will only consider cases where $D(x) = 0$ implies $D'(x) = 0$, in which case $D'(x)/D(x)$ is treated as 1. To see why such functions are relevant to our discussion, note that for every real-valued function f over X ,

$$\begin{aligned} \mathbf{E}_{x \sim D'} [f(x)] - \mathbf{E}_{x \sim D} [f(x)] &= \mathbf{E}_{x \sim D} [D'(x)f(x)/D(x)] - \mathbf{E}_{x \sim D} [f(x)] \\ &= \left\langle \frac{D' - D}{D}, f \right\rangle_D. \end{aligned}$$

This means that the inner product of any function f with $(D' - D)/D$ is equal to the difference of expectations of f under the two distributions. We also remark that the quantity $\langle \frac{D'}{D} - 1, \frac{D'}{D} - 1 \rangle_D$ is known as the $\chi^2(D', D)$ distance and is widely used for hypothesis testing in statistics [41]. A key notion for our statistical dimension is the *average correlation* of a set of distributions \mathcal{D}' relative to a distribution D . We denote it by $\rho(\mathcal{D}', D)$ and define as follows:

$$\rho(\mathcal{D}', D) \doteq \frac{1}{m^2} \sum_{D_1, D_2 \in \mathcal{D}'} \left| \left\langle \frac{D_1}{D} - 1, \frac{D_2}{D} - 1 \right\rangle_D \right|.$$

Known lower bounds for SQ learning are based on bounding the pairwise correlations between functions. Bounds on pairwise correlations when strong enough easily imply bounds on the average correlation (see Lemma 3.12 for the proof). Bounding average correlation is more involved

technically but it appears to be necessary for the tight lower bounds that we give. In Section 3.3 we describe a pairwise-correlation version of our bounds. It is sufficient for some applications and generalizes the statistical query dimension from learning theory (see Appendix B for the details).

We are now ready to define the concept of statistical dimension.

Definition 2.5. For $\bar{\gamma} > 0$, domain X and a search problem \mathcal{Z} over a set of solutions \mathcal{F} and a class of distributions \mathcal{D} over X , let d be the largest integer such that there exists a reference distribution D over X and a finite set of distributions $\mathcal{D}_D \subseteq \mathcal{D}$ with the following property: for any solution $f \in \mathcal{F}$ the set $\mathcal{D}_f = \mathcal{D}_D \setminus \mathcal{Z}^{-1}(f)$ is non-empty and for any subset $\mathcal{D}' \subseteq \mathcal{D}_f$, where $|\mathcal{D}'| \geq |\mathcal{D}_f|/d$, $\rho(\mathcal{D}', D) < \bar{\gamma}$. We define the **statistical dimension** with average correlation $\bar{\gamma}$ of \mathcal{Z} to be d and denote it by $\text{SDA}(\mathcal{Z}, \bar{\gamma})$.

The statistical dimension with average correlation $\bar{\gamma}$ of a search problem over distributions gives a lower bound on the complexity of any (deterministic) statistical algorithm for the problem that uses queries to $\text{VSTAT}(1/(3\bar{\gamma}))$.

Theorem 2.6. Let X be a domain and \mathcal{Z} be a search problem over a set of solutions \mathcal{F} and a class of distributions \mathcal{D} over X . For $\bar{\gamma} > 0$ let $d = \text{SDA}(\mathcal{Z}, \bar{\gamma})$. Any statistical algorithm requires at least d calls to $\text{VSTAT}(1/(3\bar{\gamma}))$ oracle to solve \mathcal{Z} .

At a high level our proof works as follows. It is not hard to see that for any D , an algorithm that solves \mathcal{Z} needs to “distinguish” all distributions in \mathcal{D}_f from D for some solution f . Here by “distinguishing” we mean that the algorithm needs to ask a query g such that $\mathbf{E}_D[g]$ cannot be used as a response of $\text{VSTAT}(1/(3\bar{\gamma}))$ for $D' \in \mathcal{D}_f$. In the key component of the proof we show that if a query function g to $\text{VSTAT}(1/(3\bar{\gamma}))$ distinguishes between a distribution D and any distribution $D' \in \mathcal{D}'$, then \mathcal{D}' must have average correlation of at least $\bar{\gamma}$ relative to D . The condition that for any $|\mathcal{D}'| \geq |\mathcal{D}_f|/d$, $\rho(\mathcal{D}', D) < \bar{\gamma}$ then immediately implies that at least d queries are required to distinguish any distribution in \mathcal{D}_f from D .

In Section 3.1 we give a refinement of SDA which additionally requires that the set \mathcal{D}_f is not too small (and not just non-empty). This refined notion allows us to extend the lower bound to randomized statistical algorithms. In Section 3.2 we use the more refined notion of SDA to also show that (with high probability) one can simulate t samples of 1-STAT using access to $\text{VSTAT}(O(t))$. This implies that lower bounds on SDA imply lower bounds on the query complexity of any unbiased statistical algorithm (see Theorem 3.8 for details).

Finally, in Appendix B we show that our bounds generalize and strengthen the known results SQ learning. In particular, we show that Honest SQ learning is equivalent to SQ learning and generalize the standard lower bound based on SQ-DIM [12, 51] to randomized SQ learning algorithms.

2.4 Lower Bounds

We will prove lower bounds for the bipartite planted clique under both statistical oracles (VSTAT and 1-STAT).

Theorem 2.7. For any constant $\delta > 0$, any $k \leq n^{1/2-\delta}$ and $r > 0$, at least $n^{\Omega(\log r)}$ queries to $\text{VSTAT}(n^2/(rk^2))$ are required to solve the distributional planted bipartite k -clique. In particular, no polynomial-time statistical algorithm can solve the problem using queries to $\text{VSTAT}(o(n^2/k^2))$ and any statistical algorithm will require $n^{\Omega(\log n)}$ queries to $\text{VSTAT}(n^{2-\delta}/k^2)$.

This bound also applies to any randomized algorithm with probability of success being at least a (positive) constant. Note that this bound is essentially tight. For every vertex in the clique, the probability that the corresponding bit of a randomly chosen point is set to 1 is $1/2 + k/(2n)$ whereas for every vertex not in the clique, this probability is $1/2$. Therefore using n queries to $\text{VSTAT}(16n^2/k^2)$ (i.e., of tolerance $k/4n$) it is easy to detect the planted bipartite clique. Indeed, this can be done by using the query functions $h : \{0, 1, \}^n \rightarrow \{0, 1\}$, $h(x) = x_i$, for each $i \in [n]$. So, the answers of the VSTAT oracle represent the expected value of the i th bit over the sample.

There is also a statistical algorithm that uses $n^{O(\log n)}$ queries to $\text{VSTAT}(25n/k)$ (significantly higher tolerance) to find the planted set for any $k \geq \log n$. In fact, the same algorithm can be used for the standard planted clique problem that achieves complexity $n^{O(\log n)}$. We enumerate over subsets $T \subseteq [n]$ of $\log n$ indices and query $\text{VSTAT}(25n/k)$ with the function $g_T : \{0, 1\}^n \rightarrow \{0, 1\}$ defined as 1 if and only if the point has ones in all coordinates in T . Therefore, if the set T is included in the true clique then

$$\mathbf{E}_D[g_T] = \frac{k}{n} \cdot 1 + \left(1 - \frac{k}{n}\right) 2^{-\log n} \in \left[\frac{k}{n}, \frac{k+1}{n}\right].$$

With this expectation, $\text{VSTAT}(25n/k)$ has tolerance at most $\sqrt{k(k+1)/25n^2} \leq (k+1)/5n$ and will return at least $k/n - (k+1)/(5n) > 3k/(4n)$. If, on the other hand, at least one element of T is not from the planted clique, then $\mathbf{E}_D[g_T] \leq k/(2n) + 1/n$ and $\text{VSTAT}(25n/k)$ will return at most $(k+2)/(2n) + (k+2)/(5n) < 3k/(4n)$. Thus, we will know all $\log n$ -sized subsets of the planted clique and hence the entire clique.

Our bounds have direct implications for the average-case planted bipartite k -clique problem. An instance of this problem is a random $n \times n$ bipartite graph with a $k \times k$ bipartite clique planted randomly. In Appendix A we show that the average-case planted k -biclique is equivalent to our distributional planted k -biclique with n samples. For a statistical algorithm, n samples directly correspond to having access to $\text{VSTAT}(O(n))$. Our bounds show that this problem can be solved in polynomial time when $k = \Omega(\sqrt{n})$. At the same time, for $k \leq n^{1/2-\delta}$, any statistical algorithm will require $n^{\Omega(\log n)}$ queries to $\text{VSTAT}(n^{1+\delta})$.

We also give a bound for unbiased statistical algorithms.

Theorem 2.8. *For any constant $\delta > 0$ and any $k \leq n^{1/2-\delta}$, any unbiased statistical algorithm that succeeds with probability at least $2/3$ requires $\tilde{\Omega}(n^2/k^2)$ queries to solve the distributional planted bipartite k -clique problem.*

Each query of an unbiased statistical algorithm requires a new sample from D . Therefore this bound implies that any algorithm that does not reuse samples will require $\tilde{\Omega}(n^2/k^2)$ samples. To place this bound in context, we note that it is easy to detect whether a clique of size k has been planted using $\tilde{O}(n^2/k^2)$ samples (as before, to detect if a coordinate i is in the clique we can compute the average of x_i on $\tilde{O}(n^2/k^2)$ samples). Of course, finding all vertices in the clique would require reusing samples (which statistical algorithms cannot do). Note that $n^2/k^2 \leq n$ if and only if $k \geq \sqrt{n}$.

These bounds match the state of the art for the average-case planted bipartite k -clique and planted k -clique problems. One reason why we consider this natural is that some of the algorithms that solve the problem for $k = \Omega(\sqrt{n})$ are obviously statistical. For example, the key procedure of Feige and Ron [22] removes a vertex that has the lowest degree (in the current graph) and repeats until the remaining graph is a clique. The degree is the number of ones in a

column (or row) of the adjacency matrix and can be viewed of as an estimate of the expectation of 1 appearing in the corresponding coordinate of a random sample. Even the more involved algorithms for the problem, like finding the eigenvector with the largest eigenvalue or solving an SDP, have statistical analogues.

A closely related problem is the *planted densest subgraph* problem, where edges in the planted subset appear with higher probability than in the remaining graph. This is a variant of the densest k -subgraph problem, which itself is a natural generalization of k -clique that asks to recover the densest k -vertex subgraph of a given n -vertex graph [19, 35, 8, 9]. The conjectured hardness of its average case variant, the planted densest subgraph problem, has been used in public key encryption schemes [4] and in analyzing parameters specific to financial markets [5]. Our lower bounds extend in a straightforward manner to this problem and in addition become exponential as the density becomes (inverse-polynomially) close to $1/2$ (Corollary 5.9 and the remark following it).

3 Lower Bounds from Statistical Dimension

Here we prove the general lower bounds. In later sections, we will compute the parameters in these bounds for specific problems of interest.

3.1 Lower Bounds for Statistical Algorithms

We now prove Theorem 2.6 which is the basis of all our lower bounds. We will prove a stronger version of this theorem which also applies to randomized algorithms. For this version we need an additional parameter in the definition of SDA.

Definition 3.1. For $\bar{\gamma} > 0$, $\eta > 0$, domain X and a search problem \mathcal{Z} over a set of solutions \mathcal{F} and a class of distributions \mathcal{D} over X , let d be the largest integer such that there exists a reference distribution D over X and a finite set of distributions $\mathcal{D}_D \subseteq \mathcal{D}$ with the following property: for any solution $f \in \mathcal{F}$ the set $\mathcal{D}_f = \mathcal{D}_D \setminus \mathcal{Z}^{-1}(f)$ has size at least $(1 - \eta) \cdot |\mathcal{D}_D|$ and for any subset $\mathcal{D}' \subseteq \mathcal{D}_f$, where $|\mathcal{D}'| \geq |\mathcal{D}_f|/d$, $\rho(\mathcal{D}', D) < \bar{\gamma}$. We define the **statistical dimension** with average correlation $\bar{\gamma}$ and solution set bound η of \mathcal{Z} to be d and denote it by $\text{SDA}(\mathcal{Z}, \bar{\gamma}, \eta)$.

Note that for any $\eta < 1$, $\text{SDA}(\mathcal{Z}, \bar{\gamma}) \geq \text{SDA}(\mathcal{Z}, \bar{\gamma}, \eta)$ and for $1 > \eta \geq 1 - 1/|\mathcal{D}_D|$, we get $\text{SDA}(\mathcal{Z}, \bar{\gamma}) = \text{SDA}(\mathcal{Z}, \bar{\gamma}, \eta)$, where $|\mathcal{D}_D|$ is the set of distributions that maximizes $\text{SDA}(\mathcal{Z}, \bar{\gamma})$.

Theorem 3.2. Let X be a domain and \mathcal{Z} be a search problem over a set of solutions \mathcal{F} and a class of distributions \mathcal{D} over X . For $\bar{\gamma} > 0$ and $\eta \in (0, 1)$ let $d = \text{SDA}(\mathcal{Z}, \bar{\gamma}, \eta)$. Any randomized statistical algorithm that solves \mathcal{Z} with probability $\delta > \eta$ requires at least $\frac{\delta - \eta}{1 - \eta} d$ calls to $\text{VSTAT}(1/(3\bar{\gamma}))$.

Theorem 2.6 is obtained from Theorem 3.2 by setting $\delta = 1$ and using any $1 - 1/|\mathcal{D}_D| \leq \eta < 1$. Further, for any $\eta < 1$, $\text{SDA}(\mathcal{Z}, \bar{\gamma}) \geq \text{SDA}(\mathcal{Z}, \bar{\gamma}, \eta)$ and therefore for any $\eta < 1$, a bound on $\text{SDA}(\mathcal{Z}, \bar{\gamma}, \eta)$ can be used in Theorem 2.6 in place of bound on $\text{SDA}(\mathcal{Z}, \bar{\gamma})$.

To prove the claim we will need the following simple lemma that bounds the distance between any $p \in [0, 1]$ and p' which is returned by $\text{VSTAT}(t)$ on a query with expectation p in terms of p' .

Lemma 3.3. For an integer t and any $p \in [0, 1]$, let $p' \in [0, 1]$ be such that $|p' - p| \geq \tau = \max \left\{ \frac{1}{t}, \sqrt{\frac{p(1-p)}{t}} \right\}$. Then $|p' - p| \geq \sqrt{\frac{\min\{p', 1-p'\}}{3t}}$.

Proof. First note that our conditions and bounds do not change if we replace both p and p' with $1-p$ and $1-p'$, respectively. Therefore it is sufficient to prove the bound when $p \leq 1/2$. We know that $|p' - p| \geq \tau = \max\{1/t, \sqrt{p(1-p)/t}\}$. If $p' < p \leq 1/2$ then, certainly $|p' - p| \geq \sqrt{p(1-p)/t} \geq \sqrt{p'/(3t)}$. Otherwise (when $p < p'$), we have that $1-p \geq 1/2$ and

$$p' - p \geq \max \left\{ \frac{1}{t}, \sqrt{\frac{p(1-p)}{t}} \right\} \geq p(1-p) \geq \frac{p}{2}.$$

Hence $p \leq 2p'/3$ and $p' - p \geq p'/3$. We also know that $|p - p'| \geq \tau \geq 1/t$ and therefore $|p - p'| \geq \sqrt{\frac{p'}{3t}}$. \square

We are now ready to prove Theorem 3.2.

Proof of Theorem 3.2. We prove our lower bound by exhibiting a distribution over inputs (which are distributions over X) for which every deterministic statistical algorithm that solves \mathcal{Z} with probability δ (over the choice of input) requires at least $(\delta - \eta) \cdot d/(1 - \eta)$ calls to $\text{VSTAT}(1/(3\bar{\gamma}))$. The claim of the theorem will then follow by Yao's minimax principle [52].

Let D be the reference distribution and \mathcal{D}_D be a set of distributions for which the value d is achieved. Let \mathcal{A} be a deterministic statistical algorithm that uses q queries to $\text{VSTAT}(1/(3\bar{\gamma}))$ to solve \mathcal{Z} with probability δ over a random choice of a distribution from \mathcal{D}_D . Following an approach from [23], we simulate \mathcal{A} by answering any query $h : X \rightarrow \{0, 1\}$ of \mathcal{A} with value $\mathbf{E}_D[h(x)]$. Let h_1, h_2, \dots, h_q be the queries asked by \mathcal{A} in this simulation and let f be the output of \mathcal{A} .

By the definition of SDA, for $\mathcal{D}_f = \mathcal{D}_D \setminus \mathcal{Z}^{-1}(f)$ it holds that $|\mathcal{D}_f| \geq (1 - \eta)|\mathcal{D}_D|$ and for every $\mathcal{D}' \subseteq \mathcal{D}_f$, either $\rho(\mathcal{D}', D) < \bar{\gamma}$ or $|\mathcal{D}'| \leq |\mathcal{D}_f|/d$. Let the set $\mathcal{D}^+ \subseteq \mathcal{D}_D$ be the set of distributions on which \mathcal{A} is successful. Let $\mathcal{D}_f^+ = \mathcal{D}_f \cap \mathcal{D}^+$ and we denote these distributions by $\{D_1, D_2, \dots, D_m\}$. We note that $\mathcal{D}_f^+ = \mathcal{D}^+ \setminus (\mathcal{D}_D \setminus \mathcal{D}_f)$ and therefore

$$m = |\mathcal{D}_f^+| \geq |\mathcal{D}^+| - |\mathcal{D}_D \setminus \mathcal{D}_f| \geq \delta|\mathcal{D}_D| - |\mathcal{D}_D \setminus \mathcal{D}_f| = \frac{\delta|\mathcal{D}_D| - |\mathcal{D}_D \setminus \mathcal{D}_f|}{|\mathcal{D}_D| - |\mathcal{D}_D \setminus \mathcal{D}_f|} |\mathcal{D}_f| \geq \frac{\delta - \eta}{1 - \eta} |\mathcal{D}_f|. \quad (1)$$

To lower bound q , we use a generalization of an elegant argument of Szörényi [44]. For every $k \leq q$, let A_k be the set of all distributions D_i such that

$$\left| \mathbf{E}_D[h_k(x)] - \mathbf{E}_{D_i}[h_k(x)] \right| > \tau_{i,k} \doteq \max \left\{ \frac{1}{t}, \sqrt{\frac{p_{i,k}(1-p_{i,k})}{t}} \right\},$$

where we use t to denote $1/(3\bar{\gamma})$ and $p_{i,k}$ to denote $\mathbf{E}_{D_i}[h_k(x)]$. To prove the desired bound we first prove the following two claims:

1. $\sum_{k \leq q} |A_k| \geq m$;
2. for every k , $|A_k| \leq |\mathcal{D}_f|/d$.

Combining these two implies that $q \geq d \cdot m/|\mathcal{D}_f|$. By inequality (1), $q \geq \frac{\delta - \eta}{1 - \eta} \cdot d$ giving the desired lower bound.

In the rest of the proof for conciseness we drop the subscript D from inner products and norms. To prove the first claim we assume, for the sake of contradiction, that there exists $D_i \notin \cup_{k \leq q} A_k$.

Then for every $k \leq q$, $|\mathbf{E}_D[h_k(x)] - \mathbf{E}_{D_i}[h_k(x)]| \leq \tau_{i,k}$. This implies that the replies of our simulation $\mathbf{E}_D[h_k(x)]$ are within $\tau_{i,k}$ of $\mathbf{E}_{D_i}[h_k(x)]$. By the definition of \mathcal{A} and $\text{VSTAT}(t)$, this implies that f is a valid solution for \mathcal{Z} on D_i , contradicting the condition that $D_i \in \mathcal{D}_f^+ \subseteq \mathcal{D}_D \setminus Z^{-1}(f)$.

To prove the second claim, suppose that for some $k \in [d]$, $|A_k| > |\mathcal{D}_f|/d$. Let $p_k = \mathbf{E}_D[h_k(x)]$ and assume that $p_k \leq 1/2$ (when $p_k > 1/2$ we just replace h_k by $1 - h_k$ in the analysis below). First we note that:

$$\mathbf{E}_{D_i}[h_k(x)] - \mathbf{E}_D[h_k(x)] = \mathbf{E}_D \left[\frac{D_i(x)}{D(x)} h_k(x) \right] - \mathbf{E}_D[h_k(x)] = \left\langle h_k, \frac{D_i}{D} - 1 \right\rangle = p_{i,k} - p_k.$$

Let $\hat{D}_i(x) = \frac{D_i(x)}{D(x)} - 1$, (where the convention is that $\hat{D}_i(x) = 0$ if $D(x) = 0$). We will next show upper and lower bounds on the following quantity

$$\Phi = \left\langle h_k, \sum_{D_i \in A_k} \hat{D}_i \cdot \text{sign}\langle h_k, \hat{D}_i \rangle \right\rangle.$$

By Cauchy-Schwartz we have that

$$\begin{aligned} \Phi^2 &= \left\langle h_k, \sum_{D_i \in A_k} \hat{D}_i \cdot \text{sign}\langle h_k, \hat{D}_i \rangle \right\rangle^2 \leq \|h_k\|^2 \cdot \left\| \sum_{D_i \in A_k} \hat{D}_i \cdot \text{sign}\langle h_k, \hat{D}_i \rangle \right\|^2 \\ &\leq \|h_k\|^2 \cdot \left(\sum_{D_i, D_j \in A_k} |\langle \hat{D}_i, \hat{D}_j \rangle| \right) \\ &\leq \|h_k\|^2 \cdot \rho(A_k, D) \cdot |A_k|^2. \end{aligned} \quad (2)$$

As before, we also have that

$$\begin{aligned} \Phi^2 &= \left\langle h_k, \sum_{D_i \in A_k} \hat{D}_i \cdot \text{sign}\langle h_k, \hat{D}_i \rangle \right\rangle^2 = \left(\sum_{D_i \in A_k} \langle h_k, \hat{D}_i \rangle \cdot \text{sign}\langle h_k, \hat{D}_i \rangle \right)^2 \\ &\geq \left(\sum_{D_i \in A_k} |p_{i,k} - p_k| \right)^2. \end{aligned} \quad (3)$$

To evaluate the last term of this inequality we use the fact that $|p_{i,k} - p_k| \geq \tau_{i,k} = \max\{1/t, \sqrt{p_{i,k}(1 - p_{i,k})/t}\}$ and Lemma 3.3 to obtain that for every $D_i \in A_k$,

$$|p_k - p_{i,k}| \geq \sqrt{\frac{\min\{p_k, 1 - p_k\}}{3t}} = \sqrt{\frac{p_k}{3t}}. \quad (4)$$

By substituting equation (4) into (3) we get that $\Phi^2 \geq \frac{p_k}{3t} \cdot |A_k|^2$.

We note that, h_k is a $\{0, 1\}$ -valued function and therefore $\|h_k\|^2 = p_k$. Substituting this into equation (2) we get that $\Phi^2 \leq p_k \cdot \rho(A_k, D) \cdot |A_k|^2$. By combining these two bounds on Φ^2 we obtain that $\rho(A_k, D) \geq 1/(3t) = \bar{\gamma}$ which contradicts the definition of SDA. \square

We remark that if the statistical algorithm was using the fixed tolerance $\text{STAT}(\tau)$ for $\tau = \sqrt{\bar{\gamma}}$ then $\Phi^2 \geq \left(\sum_{D_i \in A_k} |p_{i,k} - p_k|\right)^2 \geq \tau^2 |A_k|^2$ and the proof could be obtained by directly combining equations (2) and (3) to get a contradiction. This also allows us to eliminate the factor of 3 in the bound and the assumption that queries are $\{0, 1\}$ (it suffices that $\|h_k\|^2 \leq 1$). We obtain the following variant of our main bound.

Theorem 3.4. *Let X be a domain and \mathcal{Z} be a search problem over a set of solutions \mathcal{F} and a class of distributions \mathcal{D} over X . For $\bar{\gamma} > 0$, let $d = \text{SDA}(\mathcal{Z}, \bar{\gamma})$. Any statistical algorithm requires at least d calls to $\text{STAT}(\sqrt{\bar{\gamma}})$ oracle to solve \mathcal{Z} .*

3.2 Lower Bounds for Unbiased Statistical Algorithms

Next we address lower bounds on algorithms that use the 1-STAT oracle. We recall that the 1-STAT oracle returns the value of a function on a single randomly chosen point. To estimate the expectation of a function, an algorithm can simply query this oracle multiple times with the same function and average the results. A lower bound for this oracle directly translates to a lower bound on the number of samples that any statistical algorithm must use.

We note that responses of 1-STAT do not have the room for the possibly adversarial deviation afforded by the tolerance of the STAT and VSTAT oracles. The ability to use these slight deviations in a coordinated way is used crucially in our lower bounds against VSTAT and in known lower bounds for SQ learning algorithms. This makes proving lower bounds against unbiased statistical algorithms harder and indeed lower bounds for Honest SQ learning (which is a special case of unbiased statistical algorithms) required a substantially more involved argument than lower bounds for the regular SQ model [51].

Our lower bounds for unbiased statistical algorithms use a different approach. We show a direct simulation of 1-STAT oracle using VSTAT oracle. This simulation implies that the use of an unbiased statistical oracle does not give additional power to the algorithm and allows us derive lower bounds on unbiased statistical algorithms from Theorem 3.2.

Theorem 3.5. *Let \mathcal{Z} be a search problem and let \mathcal{A} be a (possibly randomized) unbiased statistical algorithm that solves \mathcal{Z} with probability at least δ using m samples from 1-STAT. For any $\delta' \in (0, 1/2]$, there exists a statistical algorithm \mathcal{A}' that uses at most m queries to $\text{VSTAT}(m/\delta'^2)$ and solves \mathcal{Z} with probability at least $\delta - \delta'$.*

Our proof relies on a simple simulation. Given query $h_1 : X \rightarrow \{0, 1\}$ from \mathcal{A} to 1-STAT, we make the same query h_1 to $\text{VSTAT}(t)$ for $t = m/\delta'^2$. Let p'_1 be the response. We flip a coin with bias p'_1 (that is one that outputs 1 with probability p'_1 and 0 with probability $1 - p'_1$) and return it to the algorithm. We do the same for the remaining $m - 1$ queries which we denote by h_2, h_3, \dots, h_m . We then prove that true m samples of 1-STAT and our simulated coin flips are statistically close. This implies that success probability of the simulated algorithm is not much worse than that of the unbiased statistical algorithm.

In our proof we will, for simplicity and without loss of generality, assume that $\text{VSTAT}(t)$ always outputs a value in the interval $[1/t, 1 - 1/t]$. We can always replace a value v returned by $\text{VSTAT}(t)$ by v' which is the closest value to v in the above interval. It is easy to see that if v is a valid answer of $\text{VSTAT}(t)$ then so is v' .

We will need the following lemmas for our proof. The first one bounds the total variation distance between two distributions in terms of the expected ratio of probability density functions.

Lemma 3.6. Let D_1 and D_2 be two distribution over a domain X of finite¹ size such that $D_2(x)$ is non-vanishing. Denote the total variation distance between D_1 and D_2 by $\Delta_{TV}(D_1, D_2)$. Then $\Delta_{TV}(D_1, D_2) \leq \sqrt{\rho}/2$, where $\rho = \mathbf{E}_{D_1} \left[\frac{D_1(x)}{D_2(x)} \right] - 1$.

Proof. The key observation is that the χ^2 -divergence between D_1 and D_2 is exactly the expected ratio minus 1.

$$\rho = \mathbf{E}_{D_1} \left[\frac{D_1(x)}{D_2(x)} \right] - 1 = \mathbf{E}_{D_2} \left[\frac{D_1^2(x)}{D_2^2(x)} \right] - 1 = \mathbf{E}_{D_2} \left[\frac{D_1^2(x)}{D_2^2(x)} - 2 \frac{D_1(x)}{D_2(x)} + 1 \right] = \mathbf{E}_{D_2} \left[\left(\frac{D_1(x)}{D_2(x)} - 1 \right)^2 \right].$$

By Jensen's inequality this implies that

$$\mathbf{E}_{D_2} \left[\left| \frac{D_1(x)}{D_2(x)} - 1 \right| \right] \leq \sqrt{\mathbf{E}_{D_2} \left[\left(\frac{D_1(x)}{D_2(x)} - 1 \right)^2 \right]} = \sqrt{\rho}.$$

Finally,

$$\Delta_{TV}(D_1, D_2) = \frac{1}{2} \sum_{x \in X} |D_1(x) - D_2(x)| = \frac{1}{2} \mathbf{E}_{D_2} \left[\left| \frac{D_1(x)}{D_2(x)} - 1 \right| \right] \leq \frac{\sqrt{\rho}}{2}.$$

□

The second lemma is that if p' is an answer of $\text{VSTAT}(t)$ for a query with h , such that $\mathbf{E}_D[h] = p$, then the expected ratio of density functions of $B(1, p)$ and $B(1, p')$ is small.

Lemma 3.7. For an integer t and $p \in [0, 1]$ let $p' \in [1/t, 1-1/t]$ such that $|p' - p| \leq \max\{\frac{1}{t}, \sqrt{\frac{p(1-p)}{t}}\}$. Then

$$\mathbf{E}_{b \sim B(1, p)} \left(\frac{\Pr[B(1, p) = b]}{\Pr[B(1, p') = b]} \right) = 1 + \frac{2}{t}.$$

Proof. If $b = 1$, the ratio is p/p' and when $b = 0$, then it is $(1-p)/(1-p')$. Thus, the expected ratio is

$$\frac{p^2}{p'} + \frac{(1-p)^2}{1-p'} = 1 + \frac{(p-p')^2}{p'(1-p')}.$$

We can assume without loss of generality that $p \leq 1/2$.

Now if $p' > p$ then $p(1-p) \leq p \leq p' \leq 2p'(1-p')$. Otherwise (when, $p' \leq p$), we know that $p \geq p' \geq 1/t$ and $p-p' \leq 1/t$. Therefore $p \leq 2p'$ and $p(1-p) \leq 2p'(1-p')$. Further $1-1/t \geq p' \geq 1/t$ and therefore $2p'(1-p') \geq 1/t$. This implies that

$$\max \left\{ \frac{1}{t}, \sqrt{\frac{p(1-p)}{t}} \right\} \leq \max \left\{ \frac{1}{t}, \sqrt{\frac{2p'(1-p')}{t}} \right\} \leq \sqrt{\frac{2p'(1-p')}{t}}.$$

By using this bound in the ratio equation we get that

$$1 + \frac{(p-p')^2}{p'(1-p')} \leq 1 + \frac{2p'(1-p')}{p'(1-p')} \leq 1 + \frac{2}{t}.$$

□

¹This assumption is simply for convenience of notation. It holds in our applications.

We can now complete the proof of Theorem 3.5.

Proof of Theorem 3.5. We simulate \mathcal{A} using $\text{VSTAT}(t)$ as described above. We now prove that for any algorithm the total variation distance between the true answers of 1-STAT and the simulated distribution is at most δ' . Formally, let R denote the set of all outcomes of \mathcal{A} 's random bits and for $r \in R$, let \mathcal{A}^r denote the execution of \mathcal{A} when its random bits are set to r . let $\Pi_{\mathcal{A}}$ denote the distribution over the m bits obtained by the algorithm \mathcal{A} when it is run with 1-STAT oracle. Similarly, let $\Pi'_{\mathcal{A}}$ denote the distribution over $\{0, 1\}^m$ obtained by running the algorithm \mathcal{A} simulated using $\text{VSTAT}(t)$ as above. By definition, $\Pi_{\mathcal{A}} = \mathbf{E}_{r \in R} \Pi_{\mathcal{A}^r}$ and similarly $\Pi'_{\mathcal{A}} = \mathbf{E}_{r \in R} \Pi'_{\mathcal{A}^r}$. This implies that,

$$\Delta_{TV}(\Pi_{\mathcal{A}}, \Pi'_{\mathcal{A}}) \leq \mathbf{E}_{r \in R} [\Delta_{TV}(\Pi_{\mathcal{A}^r}, \Pi'_{\mathcal{A}^r})] \leq \max_{r \in R} \Delta_{TV}(\Pi_{\mathcal{A}^r}, \Pi'_{\mathcal{A}^r}).$$

The algorithm \mathcal{A}^r is deterministic and it is therefore sufficient to prove the bound on total variation distance for deterministic algorithms. For conciseness we assume henceforth that \mathcal{A} is deterministic.

For any $i \in [m]$ let $\Pi_{\mathcal{A}_i}$ denote the probability distribution on the first i samples of \mathcal{A} executed with 1-STAT. For $j \leq i$ let z^j denote the first j bits of z . Let $\Pi_{\mathcal{A}_i}(z \mid z^{i-1})$ denote the probability that the first i samples of \mathcal{A} executed with 1-STAT oracle are equal to z conditioned on the probability that the first $i-1$ samples are equal to z^{i-1} . We define $\Pi'_{\mathcal{A}_i}(z)$ and $\Pi'_{\mathcal{A}_i}(z \mid z^{i-1})$ analogously. We also denote by h_z the query that \mathcal{A} asks after getting z as the response to first i samples and let $p_z = \mathbf{E}_D[h_z]$. Let p'_z denote the response of $\text{VSTAT}(t)$ on h_z .

For $i \in [m]$ and any $z \in \{0, 1\}^i$, $\Pi_{\mathcal{A}_i}(z \mid z^{i-1}) = \Pr[B(1, p_{z^{i-1}}) = z_i]$ and hence $\Pi_{\mathcal{A}_i}(z) = \Pi_{\mathcal{A}_{i-1}}(z^{i-1}) \Pr[B(1, p_{z^{i-1}}) = z_i]$. Similarly, $\Pi'_{\mathcal{A}_i}(z \mid z^{i-1}) = \Pr[B(1, p'_{z^{i-1}}) = z_i]$ and $\Pi'_{\mathcal{A}_i}(z) = \Pi'_{\mathcal{A}_{i-1}}(z^{i-1}) \Pr[B(1, p'_{z^{i-1}}) = z_i]$. This implies that:

$$\begin{aligned} \mathbf{E}_{z \sim \Pi_{\mathcal{A}_i}} \left[\frac{\Pi_{\mathcal{A}_i}(z)}{\Pi'_{\mathcal{A}_i}(z)} \right] &= \mathbf{E}_{z \sim \Pi_{\mathcal{A}_i}} \left[\frac{\Pi_{\mathcal{A}_{i-1}}(z^{i-1}) \Pr[B(1, p_{z^{i-1}}) = z_i]}{\Pi'_{\mathcal{A}_{i-1}}(z^{i-1}) \Pr[B(1, p'_{z^{i-1}}) = z_i]} \right] \\ &= \mathbf{E}_{z \sim \Pi_{\mathcal{A}_i}} \left[\frac{\Pi_{\mathcal{A}_{i-1}}(z^{i-1})}{\Pi'_{\mathcal{A}_{i-1}}(z^{i-1})} \cdot \mathbf{E}_{b \sim B(1, p_{z^{i-1}})} \left[\frac{\Pr[B(1, p_{z^{i-1}}) = b]}{\Pr[B(1, p'_{z^{i-1}}) = b]} \right] \right]. \end{aligned}$$

Now by Lemma 3.7, this implies that for any z of length $i \in [m]$,

$$\mathbf{E}_{z \sim \Pi_{\mathcal{A}_i}} \left[\frac{\Pi_{\mathcal{A}_i}(z)}{\Pi'_{\mathcal{A}_i}(z)} \right] \leq \mathbf{E}_{y \sim \Pi_{\mathcal{A}_{i-1}}} \left[\frac{\Pi_{\mathcal{A}_i}(y)}{\Pi'_{\mathcal{A}_i}(y)} \right] \cdot \left(1 + \frac{2}{t} \right).$$

Applying this iteratively we obtain that

$$\mathbf{E}_{z \sim \Pi_{\mathcal{A}}} \left[\frac{\Pi_{\mathcal{A}}(z)}{\Pi'_{\mathcal{A}}(z)} \right] \leq \left(1 + \frac{2}{t} \right)^m \leq e^{2m/t}.$$

By our definition, $t = m/\delta'^2 \geq 4m$. Therefore, $2m/t \leq 1/2$ and hence $e^{2m/t} \leq 1 + 4m/t$. By Lemma 3.6, we get that $\Delta_{TV}(\Pi_{\mathcal{A}}, \Pi'_{\mathcal{A}}) \leq \sqrt{(1 + 4m/t - 1)/2} = \sqrt{m/t} = \delta'$. This implies that the success probability of \mathcal{A} using the simulated oracle is at least $\delta - \delta'$. \square

We now combine Theorems 3.2 and 3.5 to obtain the following lower bound for unbiased statistical algorithms.

Theorem 3.8. *Let X be a domain and \mathcal{Z} be a search problem over a set of solutions \mathcal{F} and a class of distributions \mathcal{D} over X . For $\bar{\gamma} > 0$ and $\eta \in (0, 1)$, let $d = \text{SDA}(\mathcal{Z}, \bar{\gamma}, \eta)$. Any (possibly randomized) unbiased statistical algorithm that solves \mathcal{Z} with probability δ requires at least m calls to 1-STAT for*

$$m = \min \left\{ \frac{d(\delta - \eta)}{2(1 - \eta)}, \frac{(\delta - \eta)^2}{12\bar{\gamma}} \right\}.$$

In particular, if $\eta \leq 1/6$ then any algorithm with success probability of at least $2/3$ requires at least $\min\{d/4, 1/48\bar{\gamma}\}$ samples from 1-STAT.

Proof. Assuming the existence of an unbiased statistical algorithm using less than m samples we apply Theorem 3.5 for $\delta' = (\delta - \eta)/2$ to simulate the algorithm using VSTAT. The bound on m ensures that the resulting algorithm uses less than $\frac{d(\delta - \eta)}{2(1 - \eta)}$ queries to $\text{VSTAT}(\frac{1}{3\bar{\gamma}})$ and has success probability of at least $(\delta + \eta)/2$. By substituting these parameters into Theorem 3.2 we obtain a contradiction. \square

To conclude we formally state the trivial reduction in the other direction, namely that $\text{VSTAT}(t)$ oracle can be simulated using 1-STAT oracle. It has been observed in the context of learning [50] that, given a boolean query function h one can obtain an estimate of $\mathbf{E}_D[h]$ using $t = O(\log(1/\delta)/\tau^2)$ samples which with probability at least $1 - \delta$ will be within τ of $\mathbf{E}_D[h]$. Using the multiplicative Chernoff bound, it is not hard to see that a stronger claim also holds: $O(t \log(1/\delta))$ samples are sufficient to estimate $p = \mathbf{E}_D[h]$ within tolerance guaranteed by $\text{VSTAT}(t)$, namely $\max\{1/t, \sqrt{p(1-p)/t}\}$.

We also allow real-valued query functions in our model but any such query function can be replaced by $\lceil \log(1/\tau) \rceil + 2$ boolean queries each of tolerance $\tau/2$. A query i computes bit i of $1 + h(x) \in [0, 2]$ so only $\lceil \log(1/\tau) \rceil + 2$ bits are necessary to get the value of $h(x)$ within $\tau/2$. Combining these two observations gives us the following theorem.

Theorem 3.9. *Let \mathcal{Z} be a search problem and let \mathcal{A} be a statistical algorithm that solved \mathcal{Z} using q queries to $\text{VSTAT}(t)$. For any $\delta > 0$, there exists an unbiased statistical algorithm \mathcal{A}' that uses at most $O(qt \cdot \log(t) \cdot \log(q/\delta))$ samples and solves \mathcal{Z} with probability at least $1 - \delta$.*

3.3 Statistical Dimension Based on Pairwise Correlations

In addition to SDA which is based on average correlation we introduce a simpler notion based on pairwise correlations. It is sufficient for some of our applications and is easy to relate to SQ-DIM used in learning (as we do in Appendix B).

Definition 3.10. *For $\gamma, \beta > 0$, domain X and a search problem \mathcal{Z} over a set of solutions \mathcal{F} and a class of distributions \mathcal{D} over X , let m be the maximum integer such that there exists a reference distribution D over X and a finite set of distributions $\mathcal{D}_D \subseteq \mathcal{D}$ such that for any solution $f \in \mathcal{F}$, there exists a set of m distributions $\mathcal{D}_f = \{D_1, \dots, D_m\} \subseteq \mathcal{D}_D \setminus \mathcal{Z}^{-1}(f)$ with the following property:*

$$\left| \left\langle \frac{D_i}{D} - 1, \frac{D_j}{D} - 1 \right\rangle_D \right| \leq \begin{cases} \beta & \text{for } i = j \in [m] \\ \gamma & \text{for } i \neq j \in [m]. \end{cases}$$

*We define the **statistical dimension** with pairwise correlations (γ, β) of \mathcal{Z} to be m and denote it by $\text{SD}(\mathcal{Z}, \gamma, \beta)$.*

As a corollary of Theorem 2.6 we obtain the following lower bound in terms of SD.

Corollary 3.11. *Let X be a domain and \mathcal{Z} be a search problem over a set of solutions \mathcal{F} and a class of distributions \mathcal{D} over X . For $\gamma, \beta > 0$, let $m = \text{SD}(\mathcal{Z}, \gamma, \beta)$. For any $\tau > 0$, any statistical algorithm requires at least $m(\tau^2 - \gamma)/(\beta - \gamma)$ calls to the $\text{STAT}(\tau)$ oracle to solve \mathcal{Z} . In particular, if for $m > 0$, $\text{SD}(\mathcal{Z}, \gamma = \frac{m^{-2/3}}{2}, \beta = 1) \geq m$ then at least $m^{1/3}/2$ calls of tolerance $m^{-1/3}$ to the STAT oracle are required to solve \mathcal{Z} .*

Proof. Take $d = m(\tau^2 - \gamma)/(\beta - \gamma)$; we will prove that $\text{SDA}(\mathcal{Z}, \tau^2) \geq d$ and apply Theorem 3.4. Consider a set of distributions $\mathcal{D}' \subseteq \mathcal{D}_f = \{D_1, D_2, \dots, D_m\} \subseteq \mathcal{D}_D \setminus Z^{-1}(f)$, where $|\mathcal{D}'| \geq |\mathcal{D}_D \setminus Z^{-1}(f)|/d \geq m/d = (\beta - \gamma)/(\tau^2 - \gamma)$:

$$\begin{aligned} \rho(\mathcal{D}', D) &= \frac{1}{|\mathcal{D}'|^2} \sum_{D_1, D_2 \in \mathcal{D}'} \left| \left\langle \frac{D_1}{D} - 1, \frac{D_2}{D} - 1 \right\rangle_D \right| \\ &\leq \frac{1}{|\mathcal{D}'|^2} (|\mathcal{D}'|\beta + (|\mathcal{D}'|^2 - |\mathcal{D}'|)\gamma) \\ &\leq \gamma + \frac{\beta - \gamma}{|\mathcal{D}'|} \\ &\leq \tau^2 \end{aligned}$$

□

Using the same technique we can also directly bound SDA in terms of SD.

Lemma 3.12. *Let X be a domain and \mathcal{Z} be a search problem over a set of solutions \mathcal{F} and a class of distributions \mathcal{D} over X . For $\gamma, \beta > 0$, let $m = \text{SD}(\mathcal{Z}, \gamma, \beta)$. Then for every $\gamma' > \gamma$, $\text{SDA}(\mathcal{Z}, \gamma') \geq \frac{m(\gamma' - \gamma)}{\beta - \gamma}$.*

4 Warm-up: MAX-XOR-SAT

In this section, we demonstrate our techniques on a warm-up problem, MAX-XOR-SAT. For this problem, it is sufficient to use pairwise correlations, rather than average correlations.

The MAX-XOR-SAT problem over a distribution is defined as follows.

Problem 4.1. *For $\epsilon \geq 0$, the ϵ -approximate **MAX-XOR-SAT** problem is defined as follows. Given samples from some unknown distribution D over XOR clauses on n variables, find an assignment x that maximizes up to additive ϵ the probability a random clause drawn from D is satisfied.*

In the worst case, it is known that MAX-XOR-SAT is NP-hard to approximate to within $1/2 - \delta$ for any constant δ [26]. In practice, local search algorithms such as WalkSat [42] are commonly applied as heuristics for maximum satisfiability problems. We give strong evidence that the distributional version of MAX-XOR-SAT is hard for algorithms that locally seek to improve an assignment by flipping variables as to satisfy more clauses, giving some theoretical justification for the observations of [42]. Moreover, our proof even applies to the case when there exists an assignment that satisfies all the clauses generated by the target distribution.

The bound we obtain can be viewed as a restatement of the known lower bound for learning parities using statistical query algorithms (indeed, the problem of learning parities is a special case of our distributional MAX-XOR-SAT).

Theorem 4.2. For any $\delta > 0$, any statistical algorithm requires at least $\tau^2(2^n - 1)$ queries to $\text{STAT}(\tau)$ to solve $(\frac{1}{2} - \delta)$ -approximate MAX-XOR-SAT. In particular, at least $2^{n/3}$ queries of tolerance $2^{-n/3}$ are required.

For a vector $c \in \{0, 1\}^n$ we define the parity function $\chi_c(x)$ as usual $\chi_c(x) \doteq -(-1)^{c \cdot x}$. Further, let D_c to be the uniform distribution over the set $S_c = \{x \mid \chi_c(x) = 1\}$.

Lemma 4.3. For $c \in \{0, 1\}^n$, $c \neq \bar{0}$ and the uniform distribution U over $\{-1, 1\}^n$, the following hold:

$$1) \quad \mathbf{E}_{x \sim D_c} [\chi_{c'}(x)] = \begin{cases} 1 & \text{if } c = c' \\ 0 & \text{otherwise.} \end{cases} \quad 2) \quad \mathbf{E}_{x \sim U} [\chi_c(x)\chi_{c'}(x)] = \begin{cases} 1 & \text{if } c = c' \\ 0 & \text{otherwise.} \end{cases}$$

Proof. To show Part 1) note that if $c = c'$ then $\mathbf{E}_{x \in S_c} [\chi_c(x)] = 1$. If $c \neq c' \neq \bar{0}$ then it is easy to see that $|S_c \cap S_{c'}| = |S_c|/2 = |S_{c'}|/2$ and so $\mathbf{E}_{x \in S_c} [\chi_{c'}(x)] = \sum_{x \in S_c \cap S_{c'}} 1 + \sum_{x \in S_c \setminus S_{c'}} (-1) = 0$. Part 2) states the well-known fact that the parity functions are uncorrelated relative to the uniform distribution. \square

These two facts will imply that when $D = U$ (the uniform distribution) and the D_i 's consist of the D_c 's, we can set $\gamma = 0$ and $\beta = 1$, when considering the statistical dimension of the problems presented in the following sections.

We are ready to prove a lower bound for the MAX-XOR-SAT problem. Let D be a distribution over XOR clauses $c \in \{0, 1\}^n$. We interpret $c_i = 1$ as variable i appearing in c and otherwise not; for simplicity, no variables are negated in the clauses. The problem is to find an assignment $x \in \{0, 1\}^n$ that maximizes the expected number of satisfied XOR clauses. We now give the statistical dimension of this problem, from which Theorem 4.2 follows.

Theorem 4.4. For the MAX-XOR-SAT problem, let $\mathcal{F} = \{\chi_x\}_{x \in \{0, 1\}^n}$, let \mathcal{D} be the set of all distributions over clauses $c \in \{0, 1\}^n$, and for any $\delta > 0$, let \mathcal{Z} be the $(\frac{1}{2} - \delta)$ -approximate MAX-XOR-SAT problem defined over \mathcal{F} and \mathcal{D} . Then $\text{SD}(\mathcal{Z}, 0, 1) \geq 2^n - 1$.

Proof. Maximizing the expected number of satisfied clauses is equivalent to maximizing the quantity

$$\max_{x \in \{0, 1\}^n} \mathbf{E}_{c \sim D} [\chi_x(c)].$$

This proof is a fairly direct application of Lemma 4.3 to the definition of statistical dimension.

For the conditions in Definition 3.10, for each of the 2^n possible assignments to x let D_x be the uniform distribution over the clauses $c \in \{0, 1\}^n$ such that $\chi_c(x) = 1$.

Because $\chi_c(x)$ is symmetric in x and c , the conditions in Definition 3.10, with $\beta = 1$ and $\gamma = 0$, which follow from Lemma 4.3, are satisfied for the 2^n distributions D_c , with $D = U$. Because $\chi_c(x) = 1$ when assignment x satisfies clause c and -1 otherwise, we need to scale the approximation term by $1/2$ when measuring the *fraction* of satisfied clauses \square

5 Planted Clique

5.1 Statistical dimension of planted clique

We now prove the lower bound claimed in Theorem 2.7 on the problem of detecting a planted k -clique in the given distribution on vectors from $\{0, 1\}^n$ as defined above.

For a subset $S \subseteq [n]$, let D_S be the distribution over $\{0, 1\}^n$ with a planted clique on the subset S . Let $\{S_1, \dots, S_m\}$ be the set of all $\binom{n}{k}$ subsets of $[n]$ of size k . For $i \in [m]$ we use D_i to denote D_{S_i} . The reference distribution in our lower bounds will be the uniform distribution over $\{0, 1\}^n$ and let \hat{D}_S denote $D_S/D - 1$. In order to apply our lower bounds based on statistical dimension with average correlation we now prove that for the planted clique problem average correlations of large sets must be small. We start with a lemma that bounds the correlation of two planted clique distributions relative to the reference distribution D as a function of the overlap between the cliques:

Lemma 5.1. For $i, j \in [m]$,

$$\langle \hat{D}_i, \hat{D}_j \rangle_D \leq \frac{2^\lambda k^2}{n^2},$$

where $\lambda = |S_i \cap S_j|$.

Proof. For the distribution D_i , we consider the probability $D_i(x)$ of generating the vector x . Then,

$$D_i(x) = \begin{cases} \left(\frac{n-k}{n}\right) \frac{1}{2^n} + \left(\frac{k}{n}\right) \frac{1}{2^{n-k}} & \text{if } \forall s \in S_i, x_s = 1 \\ \left(\frac{n-k}{n}\right) \frac{1}{2^n} & \text{otherwise.} \end{cases}$$

Now we compute the vector $\hat{D}_i = \frac{D_i}{D} - 1$:

$$\frac{D_i}{D} - 1 = \begin{cases} \frac{k2^k}{n} - \frac{k}{n} & \text{if } \forall s \in S_i, x_s = 1 \\ -\frac{k}{n} & \text{otherwise.} \end{cases}$$

We then bound the inner product:

$$\begin{aligned} \langle \hat{D}_i, \hat{D}_j \rangle_D &\leq \frac{2^{n-2k+\lambda}}{2^n} \left(\frac{k2^k}{n} - \frac{k}{n}\right)^2 + 2 \left(\frac{2^{n-k}}{2^n}\right) \left(\frac{k2^k}{n} - \frac{k}{n}\right) \left(-\frac{k}{n}\right) + \left(-\frac{k}{n}\right)^2 \\ &\leq \frac{2^\lambda k^2}{n^2} \end{aligned}$$

□

We now give a bound on the average correlation of any \hat{D}_S with a large number of distinct clique distributions.

Lemma 5.2. For $\delta > 0$ and $k \leq n^{1/2-\delta}$, let $\{S_1, \dots, S_m\}$ be the set of all $\binom{n}{k}$ subsets of $[n]$ of size k and $\{D_1, \dots, D_m\}$ be the corresponding distributions on $\{0, 1\}^n$. Then for any integer $\ell \leq k$, set S of size k and subset $A \subseteq \{S_1, \dots, S_m\}$ where $|A| \geq 4(m-1)/n^{2\ell\delta}$,

$$\frac{1}{|A|} \sum_{S_i \in A} \langle \hat{D}_S, \hat{D}_i \rangle < 2^{\ell+2} \frac{k^2}{n^2}.$$

Proof. In this proof we first show that if the total number of sets in A is large then most of sets in A have a small overlap with S . We then use the bound on the overlap of most sets to obtain a bound on the average correlation of D_S with distributions for sets in A .

Formally, we let $\alpha = \frac{k^2}{n^2}$ and using Lemma 5.1 get the bound $\langle \hat{D}_i, \hat{D}_j \rangle \leq 2^{|S_i \cap S_j|} \alpha$. Summing over $S_i \in A$,

$$\sum_{S_i \in A} \langle \hat{D}_S, \hat{D}_i \rangle \leq \sum_{S_i \in A} 2^{|S \cap S_i|} \alpha.$$

For any set $A \subseteq \{S_1, \dots, S_m\}$ of size t this bound is maximized when the sets of A include S , then all sets that intersect S in $k-1$ indices, then all sets that intersect S in $k-2$ indices and so on until the size bound t is exhausted. We can therefore assume without loss of generality that A is defined in precisely this way.

Let

$$T_\lambda = \{S_i \mid |S \cap S_i| = \lambda\}$$

denote the subset of all k -subsets that intersect with S in exactly λ indices. Let λ_0 be the smallest λ for which $A \cap T_\lambda$ is non-empty. We first observe that for any $1 \leq j \leq k-1$,

$$\frac{|T_j|}{|T_{j+1}|} = \frac{\binom{k}{j} \binom{n-k}{k-j}}{\binom{k}{j+1} \binom{n-k}{k-j-1}} = \frac{(j+1)(n-2k+j+1)}{(k-j-1)(k-j)} \geq \frac{(j+1)(n-2k)}{k(k+1)} \geq \frac{(j+1)n^{2\delta}}{2}.$$

By applying this equation inductively we obtain,

$$|T_j| \leq \frac{2^j \cdot |T_0|}{j! \cdot n^{2\delta j}} < \frac{2^j \cdot (m-1)}{j! \cdot n^{2\delta j}}$$

and

$$\sum_{k \geq \lambda \geq j} |T_\lambda| < \sum_{k \geq \lambda \geq j} \frac{2^\lambda \cdot (m-1)}{\lambda! \cdot n^{2\delta \lambda}} \leq \frac{4(m-1)}{n^{\delta j}}.$$

By definition of λ_0 , $|A| \leq \sum_{j \geq \lambda_0} |T_j| < 4(m-1)/n^{2\delta \lambda_0}$. In particular, if $|A| \geq 4(m-1)/n^{2\ell\delta}$ then $n^{2\ell\delta}/4 < n^{2\delta \lambda_0}/4$ or $\lambda_0 < \ell$. Now we can conclude that

$$\begin{aligned} \sum_{S_i \in A} \langle \hat{D}_S, \hat{D}_i \rangle &\leq \sum_{j=\lambda_0}^k 2^j |T_j \cap A| \alpha \\ &\leq \left(2^{\lambda_0} |T_{\lambda_0} \cap A| + \sum_{j=\lambda_0+1}^k 2^j |T_j| \right) \alpha \\ &\leq \left(2^{\lambda_0} |T_{\lambda_0} \cap A| + 2 \cdot 2^{\lambda_0+1} |T_{\lambda_0+1}| \right) \alpha \\ &\leq 2^{\lambda_0+2} |A| \alpha < 2^{\ell+2} |A| \alpha. \end{aligned}$$

To derive the last inequality we need to note that for every $j \geq 0$, $2^j |T_j| > 2(2^{j+1} |T_{j+1}|)$ we can therefore telescope the sum. \square

Lemma 5.2 gives a simple way to bound the statistical dimension with average correlation of the planted bipartite k -clique problem.

Theorem 5.3. *For $\delta > 0$ and $k \leq n^{1/2-\delta}$ let \mathcal{Z} the planted bipartite k -clique problem. Then for any $\ell \leq k$, $\text{SDA}(\mathcal{Z}, 2^{\ell+2} k^2/n^2, 1/\binom{n}{k}) \geq n^{2\ell\delta}/4$.*

Proof. Let $\{S_1, \dots, S_m\}$ be the set of all $\binom{n}{k}$ subsets of $[n]$ of size k and $\mathcal{D} = \{D_1, \dots, D_m\}$ be the corresponding distributions on $\{0, 1\}^n$. For every solution $S \in \mathcal{F}$, $\mathcal{Z}^{-1}(S) = D_S$ and let $\mathcal{D}_S = \mathcal{D} \setminus \{D_S\}$. Note that $|\mathcal{D}_S| = \binom{n}{k} - 1$ and therefore $|\mathcal{D}_S| \geq (1 - 1/\binom{n}{k})|\mathcal{D}|$.

Let \mathcal{D}' be a set of distributions $\mathcal{D}' \subseteq \mathcal{D}_S$ such that $|\mathcal{D}'| \geq 4(m-1)/n^{2\ell\delta}$. Then by Lemma 5.2, for every $S_i \in \mathcal{D}'$,

$$\frac{1}{|\mathcal{D}'|} \sum_{S_j \in \mathcal{D}'} \langle \hat{D}_i, \hat{D}_j \rangle < 2^{\ell+2} \frac{k^2}{n^2}.$$

In particular, $\rho(\mathcal{D}', D) < 2^{\ell+2} \frac{k^2}{n^2}$. By the definition of SDA, this means that $\text{SDA}(\mathcal{Z}, 2^{\ell+2} k^2/n^2, 1/\binom{n}{k}) \geq n^{2\ell\delta}/4$. \square

Combining Theorems 3.2 and 5.3 gives Theorem 2.7 by taking $\ell = \log(r)$. Theorems 3.8 and 5.3 imply the sample complexity lower bound stated in Theorem 2.8 by taking $\ell \geq 4/\delta$.

5.2 Generalized Planted Densest Subgraph

We will now show the lower bound on detecting a planted densest subset, a generalization of the planted clique problem.

Problem 5.4. Fix $0 < q \leq p \leq 1$. For $1 \leq k \leq n$, let $S \subseteq \{1, 2, \dots, n\}$ be a set of k vertex indices and D_S be a distribution over $\{0, 1\}^n$ such that when $x \sim D_S$, with probability $1 - (k/n)$ the entries of x are independently chosen according to a q biased Bernoulli variable, and with probability k/n the k coordinates in S are independently chosen according to a p biased Bernoulli variable, and the rest are independently chosen according to a q biased Bernoulli variable. The **generalized planted bipartite densest k -subgraph** problem is to find the unknown subset S given access to samples from D_S .

Note that $p = 1, q = 1/2$ is precisely the planted bipartite clique problem. For this generalized problem, we will take D , the reference distribution, to be that of n independent Bernoulli variables with bias q . Before we give our results for this problem, we have to fix some further notation: for a $x \in \{0, 1\}^n$, we define $\|x\|_1 = \sum x_i$ (ie the number of 1's in x); similarly for $\|\bar{x}\|_1 = \sum 1 - x_i$ (the number of 0's in x). We will denote the restriction of a set by subscripting so that x_S is x restricted to the subset $S \subseteq [n]$.

First, we give a computation of the correlation. This is a generalized version of Lemma 5.1.

Lemma 5.5. Fix $0 < q \leq p \leq 1$. For $i, j \in [m]$,

$$\langle \hat{D}_i, \hat{D}_j \rangle_D = \left(\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^\lambda - 1 \right) \frac{k^2}{n^2},$$

where $\lambda = |S_i \cap S_j|$.

Proof. For any x , we have $D(x) = q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1}$. For $D_i(x)$:

$$\begin{aligned} D_i(x) &= \Pr[x|\text{planted}] \Pr[\text{planted}] + \Pr[x|\text{not planted}] \Pr[\text{not planted}] \\ &= \frac{k}{n} p^{\|x_{S_i}\|_1} (1-p)^{\|\bar{x}_{S_i}\|_1} q^{\|x_{\bar{S}_i}\|_1} (1-q)^{\|\bar{x}_{\bar{S}_i}\|_1} + \left(1 - \frac{k}{n} \right) q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1} \end{aligned}$$

For $D_i(x)/D(x) - 1$, we have:

$$\begin{aligned} \frac{D_i(x)}{D(x)} - 1 &= \frac{k}{n} \cdot \frac{p^{\|x_{S_i}\|_1} (1-p)^{\|\bar{x}_{S_i}\|_1} q^{\|x_{S_i}\|_1} (1-q)^{\|\bar{x}_{S_i}\|_1}}{q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1}} - \frac{k}{n} \\ &= \frac{k}{n} \left(\frac{p}{q}\right)^{\|x_{S_i}\|_1} \left(\frac{1-p}{1-q}\right)^{\|\bar{x}_{S_i}\|_1} - \frac{k}{n} \end{aligned}$$

Now, for S_j where $|S_i \cap S_j| = \lambda$, we want to compute:

$$\langle \hat{D}_i, \hat{D}_j \rangle_D = \left(\frac{k}{n}\right)^2 \sum_{x \in \{0,1\}^n} q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1} \left[\left(\frac{p}{q}\right)^{\|x_{S_i}\|_1} \left(\frac{1-p}{1-q}\right)^{\|\bar{x}_{S_i}\|_1} - 1 \right] \left[\left(\frac{p}{q}\right)^{\|x_{S_j}\|_1} \left(\frac{1-p}{1-q}\right)^{\|\bar{x}_{S_j}\|_1} - 1 \right]$$

There are three types of terms in the product in the summand. We deal with all these terms by repeated applications of the Binomial theorem. The first term illustrates this approach:

$$\sum_{x \in \{0,1\}^n} q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1} = (q + (1-q))^n = 1$$

The second type of term is given by:

$$\begin{aligned} &\sum_{x \in \{0,1\}^n} q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1} \left(\frac{p}{q}\right)^{\|x_{S_i}\|_1} \left(\frac{1-p}{1-q}\right)^{\|\bar{x}_{S_i}\|_1} \\ &= \sum_{x \in \{0,1\}^n} q^{\|x_{S_i}\|_1} (1-q)^{\|\bar{x}_{S_i}\|_1} p^{\|x_{S_i}\|_1} (1-p)^{\|\bar{x}_{S_i}\|_1} \\ &= \sum_{y \in \{0,1\}^{|S_i|}} p^{\|y\|_1} (1-p)^{\|\bar{y}\|_1} \sum_{z \in \{0,1\}^{|S_i^c|}} p^{\|z\|_1} (1-p)^{\|\bar{z}\|_1} \\ &= 1 \end{aligned}$$

The third type of term is more complicated – using the above trick, we can restrict x to $T = S_i \cup S_j$ because the sum take over the remaining x_i yields 1.

$$\begin{aligned} &\sum_{x \in \{0,1\}^n} q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1} \left[\left(\frac{p}{q}\right)^{\|x_{S_i}\|_1} \left(\frac{1-p}{1-q}\right)^{\|\bar{x}_{S_i}\|_1} \right] \left[\left(\frac{p}{q}\right)^{\|x_{S_j}\|_1} \left(\frac{1-p}{1-q}\right)^{\|\bar{x}_{S_j}\|_1} \right] \\ &= \sum_{x \in \{0,1\}^{|T|}} q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1} \left[\left(\frac{p}{q}\right)^{\|x_{S_i}\|_1} \left(\frac{1-p}{1-q}\right)^{\|\bar{x}_{S_i}\|_1} \right] \left[\left(\frac{p}{q}\right)^{\|x_{S_j}\|_1} \left(\frac{1-p}{1-q}\right)^{\|\bar{x}_{S_j}\|_1} \right] \end{aligned}$$

Similarly, we can sum x over components associated with S_i/S_j and S_j/S_i . Hence, the sum sim-

plifies:

$$\begin{aligned}
& \sum_{x \in \{0,1\}^n} q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1} \left[\left(\frac{p}{q} \right)^{\|x_{S_i}\|_1} \left(\frac{1-p}{1-q} \right)^{\|\bar{x}_{S_i}\|_1} \right] \left[\left(\frac{p}{q} \right)^{\|x_{S_j}\|_1} \left(\frac{1-p}{1-q} \right)^{\|\bar{x}_{S_j}\|_1} \right] \\
&= \sum_{x \in \{0,1\}^{|S_i \cap S_j|}} q^{\|x\|_1} (1-q)^{\|\bar{x}\|_1} \left[\left(\frac{p}{q} \right)^{\|S_i \cap S_j\|_1} \left(\frac{1-p}{1-q} \right)^{\|\overline{S_i \cap S_j}\|_1} \right]^2 \\
&= \sum_{x \in \{0,1\}^{|S_i \cap S_j|}} \left(\frac{p^2}{q} \right)^{\|S_i \cap S_j\|_1} \left(\frac{(1-p)^2}{1-q} \right)^{\|\overline{S_i \cap S_j}\|_1} \\
&= \left(\frac{p^2}{q} + \frac{(1-p)^2}{1-q} \right)^\lambda \\
&= \left(\frac{k}{n} \right)^2 \left[\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^\lambda \right]
\end{aligned}$$

Combining these three calculations yields:

$$\langle \hat{D}_i, \hat{D}_j \rangle_D = \left(\frac{k}{n} \right)^2 \left[\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^\lambda - 1 \right]$$

□

Next, in analogy with Lemma 5.2, we give a bound on average correlation for sufficiently many distributions.

Lemma 5.6. *Fix $0 < q \leq p \leq 1$. For $\delta > 0$ and $k \leq n^{1/2-\delta}$, let $\{S_1, \dots, S_m\}$ be the set of all $\binom{n}{k}$ subsets of $[n]$ of size k and $\{D_1, \dots, D_m\}$ be the corresponding distributions on $\{0,1\}^n$. Let $n^{2\delta} \geq 1 + \frac{(p-q)^2}{q(1-q)}$. Then for any integer $\ell \leq k$, set S of size k and subset $A \subseteq \{S_1, \dots, S_m\}$ where $|A| \geq 4(m-1)/n^{2\ell\delta}$,*

$$\frac{1}{|A|} \sum_{S_i \in A} \langle \hat{D}_S, \hat{D}_i \rangle < \frac{2k^2}{n^2} \left(\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^{\ell+1} - 1 \right).$$

Proof. We proceed as in the proof of Lemma 5.2. The only difference is that we must verify that the tail terms in the following sum are geometrically decreasing:

$$\sum_{j=\lambda_0+1}^k \left(\frac{(p-q)^2}{q(1-q)} \right)^j |T_j|$$

In particular, it suffices to show that the succeeding term is at most half of the current term:

$$\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^j |T_j| \geq 2 \left(1 + \frac{(p-q)^2}{q(1-q)} \right)^{j+1} |T_{j+1}|$$

which is equivalent to:

$$\frac{|T_j|}{|T_{j+1}|} \geq 1 + \frac{(p-q)^2}{q(1-q)}$$

From the proof of Lemma 5.2, we know that the left hand side is lower bounded, and then we can apply our hypothesis:

$$\frac{|T_j|}{|T_{j+1}|} \geq \frac{(j+1)n^{2\delta}}{2} \geq 1 + \frac{(p-q)^2}{q(1-q)}$$

To conclude, with the appropriate choice of λ_0 , we obtain

$$\begin{aligned} \sum_{S_i \in A} \langle \hat{D}_S, \hat{D}_i \rangle &\leq \sum_{j=\lambda_0}^k \alpha \left(\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^j - 1 \right) |T_j \cap A| \\ &\leq \left(|T_{\lambda_0} \cap A| \left(\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^{\lambda_0} - 1 \right) + \sum_{j=\lambda_0+1}^k |T_j| \left(1 + \frac{(p-q)^2}{q(1-q)} \right)^j - \sum_{j=\lambda_0+1}^k |T_j| \right) \alpha \\ &\leq \left(\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^{\lambda_0} |T_{\lambda_0} \cap A| + 2 \cdot |T_{\lambda_0+1}| \left(1 + \frac{(p-q)^2}{q(1-q)} \right)^{\lambda_0+1} - \sum_{j=\lambda_0+1}^k |T_j| \right) \alpha \\ &\leq 2\alpha |A| \left(\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^{\lambda_0+1} - 1 \right) \end{aligned}$$

□

We can now give a bound on statistical dimension SDA:

Theorem 5.7. *Fix $0 < q \leq p \leq 1$. For $\delta > 0$ and $k \leq n^{1/2-\delta}$ let \mathcal{Z} be the generalized planted bipartite densest subgraph problem. Then for any $\ell \leq k$,*

$$\text{SDA} \left(\mathcal{Z}, \frac{2k^2}{n^2} \left(\left(1 + \frac{(p-q)^2}{q(1-q)} \right)^{\ell+1} - 1 \right), \frac{1}{\binom{n}{k}} \right) \geq n^{2\ell\delta}/4.$$

provided $n^{2\delta} \geq 1 + (p-q)^2/q(1-q)$.

Proof. The proof is analogous to that of Theorem 5.3. □

This SDA bound also yields lower bounds for the VSTAT oracle:

Corollary 5.8. *Fix $0 < q \leq p \leq 1$. For any constant $\delta > 0$, any $k \leq n^{1/2-\delta}$, $\ell \leq k$, at least $n^{\Omega(\ell)}$ queries to $VSTAT(n^2(1+(p-q)^2/q(1-q))^{-\ell-1}/k^2)$ are required to solve the distributional planted bipartite densest k -subgraph with density q provided that $n^{2\delta} \geq 1 + (p-q)^2/q(1-q)$.*

One is often interested in the case when $q = 1/2$ and $p > q$ (the classical planted densest k -subgraph problem); we give the following results for this setting. We now observe that our bound becomes exponential when p approaches $1/2$.

Corollary 5.9. *For any constant $\delta > 0$, any $k \leq n^{1/2-\delta}$, $\ell \leq k$ and density parameter $p = 1/2 + \alpha$, at least $n^{\Omega(\ell)}$ queries to $VSTAT(n^2/(\ell\alpha^2k^2))$ are required to solve the distributional planted bipartite densest k -subgraph with density p .*

For example consider the setting $k = n^{1/3}$ and $\alpha = n^{-1/4}$. It is not hard to see that for this setting the problem can be solved on a random bipartite graph with n vertices on both sides (in exponential time). Our lower bound for this setting implies that at least $n^{\Omega(n^{1/3})}$ queries to $VSTAT(n^{3/2})$ will be required.

With appropriate choices of parameter settings, we get the following corollary for 1-STAT.

Corollary 5.10. *For constants $c, \delta > 0$, density $1/2 < p \leq 1/2 + 1/n^c$, and $k \leq n^{1/2-\delta}$, any unbiased statistical algorithm that succeeds with probability at least $2/3$, requires $\tilde{\Omega}((n^{2+2c})/k^2)$ queries to find a planted densest subgraph of size k .*

Proof. By applying Theorem 3.8 and choosing $\ell = O(\log \log(n))$, it suffices to estimate $1/(48\bar{\gamma})$ to obtain the desired lower bound. We have:

$$\begin{aligned} \frac{1}{48\bar{\gamma}} &= \frac{n^2}{48k^2} \cdot \frac{1}{(2(p^2 + (1-p)^2)^\ell - 1)} \\ &\leq \frac{n^2}{48k^2} \frac{1}{(1 + 4/n^{2c})^\ell - 1} \\ &\leq \frac{n^{2+2c}}{24lk^2} \end{aligned}$$

where the last inequality holds for sufficiently large n . □

Acknowledgments

We thank Benny Applebaum, Avrim Blum, Uri Feige, Ravi Kannan, Michael Kearns, Robi Krauthgamer, Moni Naor, Jan Vondrak, and Avi Wigderson for insightful comments and helpful discussions.

References

- [1] N. Alon, A. Andoni, T. Kaufman, K. Matulef, R. Rubinfeld, and N. Xie. Testing k -wise and almost k -wise independence. In *Proceedings of STOC*, pages 496–505, 2007.
- [2] N. Alon, M. Krivelevich, and B. Sudakov. Finding a large hidden clique in a random graph. In *SODA*, pages 594–598, 1998.
- [3] B. P. W. Ames and S. A. Vavasis. Nuclear norm minimization for the planted clique and biclique problems. *Math. Program.*, 129(1):69–89, 2011.
- [4] B. Applebaum, B. Barak, and A. Wigderson. Public-key cryptography from different assumptions. In *STOC*, pages 171–180, 2010.
- [5] S. Arora, B. Barak, M. Brunnermeier, and R. Ge. Computational complexity and information asymmetry in financial products (extended abstract). In *ICS*, pages 49–65, 2010.

- [6] P. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- [7] A. Belloni, R. M. Freund, and S. Vempala. An efficient rescaled perceptron algorithm for conic systems. *Math. Oper. Res.*, 34(3):621–641, 2009.
- [8] A. Bhaskara, M. Charikar, E. Chlamtac, U. Feige, and A. Vijayaraghavan. Detecting high log-densities: an $o(n^{1/4})$ approximation for densest k -subgraph. In *STOC*, pages 201–210, 2010.
- [9] A. Bhaskara, M. Charikar, A. Vijayaraghavan, V. Guruswami, and Y. Zhou. Polynomial integrality gaps for strong sdp relaxations of densest k -subgraph. In *SODA*, pages 388–405, 2012.
- [10] A. Blum, C. Dwork, F. McSherry, and K. Nissim. Practical privacy: the SuLQ framework. In *Proceedings of PODS*, pages 128–138, 2005.
- [11] A. Blum, A. M. Frieze, R. Kannan, and S. Vempala. A polynomial-time algorithm for learning noisy linear threshold functions. *Algorithmica*, 22(1/2):35–52, 1998.
- [12] A. Blum, M. L. Furst, J. C. Jackson, M. J. Kearns, Y. Mansour, and S. Rudich. Weakly learning dnf and characterizing statistical query learning using fourier analysis. In *STOC*, pages 253–262, 1994.
- [13] S. Brubaker and S. Vempala. Random tensors and planted cliques. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 5687, pages 406–419. 2009.
- [14] C. Chu, S. Kim, Y. Lin, Y. Yu, G. Bradski, A. Ng, and K. Olukotun. Map-reduce for machine learning on multicore. In *Proceedings of NIPS*, pages 281–288, 2006.
- [15] A. Coja-Oghlan. Graph partitioning via adaptive spectral techniques. *Combinatorics, Probability & Computing*, 19(2):227–284, 2010.
- [16] Y. Dekel, O. Gurel-Gurevich, and Y. Peres. Finding hidden cliques in linear time with high probability. In *Proceedings of ANALCO*, pages 67–75, 2011.
- [17] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.
- [18] J. Dunagan and S. Vempala. A simple polynomial-time rescaling algorithm for solving linear programs. *Math. Program.*, 114(1):101–114, 2008.
- [19] U. Feige. Relations between average case complexity and approximation complexity. In *IEEE Conference on Computational Complexity*, page 5, 2002.
- [20] U. Feige and R. Krauthgamer. Finding and certifying a large hidden clique in a semirandom graph. *Random Struct. Algorithms*, 16(2):195–208, 2000.
- [21] U. Feige and R. Krauthgamer. The probable value of the lovász–schrijver relaxations for maximum independent set. *SIAM Journal on Computing*, 32(2):345–370, 2003.

- [22] U. Feige and D. Ron. Finding hidden cliques in linear time. In *Proceedings of AofA*, pages 189–204, 2010.
- [23] V. Feldman. A complete characterization of statistical query learning with applications to evolvability. *Journal of Computer System Sciences*, 78(5):1444–1459, 2012.
- [24] A. M. Frieze and R. Kannan. A new approach to the planted clique problem. In *FSTTCS*, pages 187–198, 2008.
- [25] A. E. Gelfand and A. F. M. Smith. Sampling based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85:398–409, 1990.
- [26] J. Håstad. Some optimal inapproximability results. *J. ACM*, 48:798–859, July 2001.
- [27] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- [28] E. Hazan and R. Krauthgamer. How hard is it to approximate the best nash equilibrium? *SIAM J. Comput.*, 40(1):79–91, 2011.
- [29] J. Jackson. On the efficiency of noise-tolerant PAC algorithms derived from statistical queries. *Annals of Mathematics and Artificial Intelligence*, 39(3):291–313, Nov. 2003.
- [30] M. Jerrum. Large cliques elude the metropolis process. *Random Struct. Algorithms*, 3(4):347–360, 1992.
- [31] A. Juels and M. Peinado. Hiding cliques for cryptographic security. *Des. Codes Cryptography*, 20(3):269–280, 2000.
- [32] R. Kannan. personal communication.
- [33] R. Karp. Probabilistic analysis of graph-theoretic algorithms. In *Proceedings of Computer Science and Statistics 12th Annual Symposium on the Interface*, page 173, 1979.
- [34] M. Kearns. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM*, 45(6):983–1006, 1998.
- [35] S. Khot. Ruling out ptas for graph min-bisection, densest subgraph and bipartite clique. In *FOCS*, pages 136–145, 2004.
- [36] S. Kirkpatrick, D. G. Jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [37] L. Kucera. Expected complexity of graph partitioning problems. *Discrete Applied Mathematics*, 57(2-3):193–212, 1995.
- [38] F. McSherry. Spectral partitioning of random graphs. In *FOCS*, pages 529–537, 2001.
- [39] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

- [40] L. Minder and D. Vilenchik. Small clique detection and approximate nash equilibria. 5687:673–685, 2009.
- [41] K. Pearson. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *Philosophical Magazine, Series 5*, 50(302):157–175, 1900.
- [42] B. Selman, H. Kautz, and B. Cohen. Local search strategies for satisfiability testing. In *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 521–532, 1995.
- [43] R. Servedio. Computational sample complexity and attribute-efficient learning. *Journal of Computer and System Sciences*, 60(1):161–178, 2000.
- [44] B. Szörényi. Characterizing statistical query learning: Simplified notions and proofs. In *ALT*, pages 186–200, 2009.
- [45] M. Tanner and W. Wong. The calculation of posterior distributions by data augmentation (with discussion). *Journal of the American Statistical Association*, 82:528–550, 1987.
- [46] L. G. Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, 1984.
- [47] L. G. Valiant. Evolvability. *J. ACM*, 56(1), 2009.
- [48] V. Vapnik and A. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probab. and its Applications*, 16(2):264–280, 1971.
- [49] V. Černý. Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. *Journal of Optimization Theory and Applications*, 45(1):41–51, Jan. 1985.
- [50] K. Yang. On learning correlated boolean functions using statistical queries. In *Proceedings of ALT*, pages 59–76, 2001.
- [51] K. Yang. New lower bounds for statistical query learning. *J. Comput. Syst. Sci.*, 70(4):485–509, 2005.
- [52] A. Yao. Probabilistic computations: Toward a unified measure of complexity. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pages 222–227, 1977.

A Average-case vs Distributional Planted Bipartite Clique

In this section we show the equivalence between the average-case planted biclique problem (where cliques of size $k \times k$ are planted in bipartite graphs of size $n \times n$) and the distributional biclique problem (with n examples from $\{0, 1\}^n$ where there is a set of k coordinates that contain a plant). The idea of both the reductions is to obtain examples from one distribution given examples from the other distribution. This is not immediate due to the fact that in the distributional biclique problem the plant does necessarily have the same size on the left set of vertices as on the right set. To remedy this, we will need to essentially search over plants of smaller and smaller sizes on one side, by replacing possibly planted vertices with common ones.

Definition A.1. (*Average-case planted bipartite clique (biclique) problem (APBC(n, k))*) Given integers $1 \leq k \leq n$, consider the following distribution $\mathcal{D}_{avg}(n, k)$ on bipartite graphs on $[n] \times [n]$ vertices. Pick two random sets of k vertices each from left/right respectively, say S_1 and S_2 . Plant a bipartite clique on $S_1 \times S_2$ and add an edges between each remaining pairs (i, j) with probability $1/2$. The problem is to find the $S_1 \times S_2$ biclique given as input a random bipartite graph G on $[n] \times [n]$ vertices chosen according to $\mathcal{D}_{avg}(n, k)$.

We recall the definition of the distributional biclique problem from the introduction.

Definition A.2. (*Distributional planted biclique problem (DPBC(n, k))*) Given integers $1 \leq k \leq n$ and $S \subset [n]$ a set of k vertices, consider the distribution D_S over $\{0, 1\}^n$, such that when $x \sim D_S$, w.p. $1 - (k/n)$ the entries of x are chosen uniformly and independently from $\{0, 1\}$ and with probability k/n the k coordinates in S are set to 1 and the rests are chosen uniformly and independently from $\{0, 1\}$. The problem is to find the unknown subset S given inputs from $\{0, 1\}^n$ chosen according to D_S .

Proposition A.3. Suppose that there is an algorithm that solves APBC(n', k') in time $T'(n', k')$ and outputs the correct answer w.p. $p'(n', k')$. Then there exists an algorithm that solves DPBC(n, k) in time $T(n, k) = O(nT'(n, k/2))$ and outputs the correct answer with probability $p(n, k) = (1 - 2e^{-k/8})p'(n, k/2)$.

Proof. We will think of the distribution $\mathcal{D}_{avg}(n, k')$ on graphs as a distribution on their respective adjacency matrices from $\{0, 1\}^{n \times n}$. Given k and n , and access to n samples from $D_S \subset \{0, 1\}^n$ for some set S of size k , we will design an algorithm that finds S and which makes $O(n)$ calls to the algorithm $A_{n, k'}$ that solves instances of APBC(n, k') for some values of k' .

Let M be the $n \times n$ binary matrix whose rows are the n samples from D_S . First apply a random permutation $\pi : [n] \rightarrow [n]$ to the columns of M to obtain M' (this will ensure that the planted set is uniformly distributed among the n coordinates, which is necessary in order to obtain instances distributed according to $\mathcal{D}_{avg}(n, k)$.)

In what follows we will denote by $k' \times k$ a clique with k' vertices on the left and k vertices on the right. Note that M' has a $k' \times k$ planted clique for some k' that is distributed according to a binomial distribution of n Bernoulli trials with probability of success k/n (i.e. $B(n, k/n)$).

By a multiplicative Chernoff bound, $Pr[k/2 < k' < k] \geq 1 - 2e^{-k/8}$. From now on we'll condition on this event occurring.

Suppose that $k/2 < k' \leq k$. We aim at obtaining instances of APBC(n, k') but recall that there is a plant of size k on the columns. To obtain only k' planted columns (i.e. vertices on the right of the biclique), we will be replacing sequentially the columns by random columns whose entries are chosen uniformly from $\{0, 1\}$. That is, pick one after the other n uniformly random columns of M' (without replacement) and replace them by random $\{0, 1\}^n$ vectors. Let M'_1, M'_2, \dots, M'_n be the sequence of matrices obtained this way. Since we started with a plant of size $k' \times k$ and we ended with no plant it must have been that one of the matrices in the sequence had a plant of size $k' \times k'$. Let N be that matrix. To obtain the $k' \times k'$ plant we run $A_{n, k'}$ on all the matrices M'_i , $i \in [n]$, and in particular since N is a matrix distributed according to $\mathcal{D}_{avg}(n, k')$, $A_{n, k'}$ will output the $k' \times k'$ plant on $S_1 \times S_2$ from N w.p. $p'(n, k')$. Note that we would only recover $k' \leq k$ coordinates in this way (after applying the inverse permutation π^{-1} to the set S_2), and we need to recover the entire set S of k indices. To recover the remaining planted indices, one can look at the actual examples from the matrix M corresponding to rows indexed from S_1 . Then we can output all the indices

$i \in [n]$ such that the value of each of these rows at i is 1. Since there was a plant of size k , the set output must contain S . For simplicity we can assume that the given samples contain a unique set of k planted positions.

If $k < k' < 2k$ then we apply a similar argument. Now we can aim at finding plants of size $k \times k$. This time we need to reduce the size of the plant indexed by rows of M' (corresponding to the left vertices of the clique). As before we will sequentially replace random rows with vectors uniformly chosen in $\{0, 1\}^n$, one at a time until we have replaced all of them. Running $A_{n,k}$ on each of the n instances obtained this way lets one recover the plant $S_1 \times S_2$ w.p. $p'(k, n)$. If this is successful the correct set S is $\pi^{-1}(S_2)$.

We can assume that it is harder to find smaller plants than larger ones, and so $T'(n, k') \leq T'(k/2, n)$ and $p'(n, k') \leq p'(n, k/2)$. Therefore the running time of our algorithm is $T(n, k) = O(nT'(n, k/2))$, and its success probability is $p(n, k) = (1 - 2e^{-k/8})p'(n, k/2)$. \square

Proposition A.4. *Suppose that there is an algorithm that solves DPBC(n, k) that runs in time $T(n, k)$ and outputs the correct answer with probability $p(n, k)$. Then there exists an algorithm that solves APBC(n', k') in time $T'(n', k') = O(nT(n', k'))$ and outputs the correct answer w.p. $p'(n', k') \geq (1 - 2e^{-k'/8})(p(n', k') - 2e^{-k'/8})$.*

Proof. Let $A_{n,k}$ denote the algorithm for DPBC(n, k) which takes n samples chosen according to some D_S and outputs the planted set S w.p. $p(n, k)$ (for any S of size k). Since the number of samples that witness the plant is distributed according to $B(n, k/n)$, as before, by Chernoff's bound this number is some k_1 such that $k/2 \leq k_1 \leq 2k$ w.p. $1 - 2e^{-k/8}$. Therefore the distribution of the n samples (viewed as $n \times n$ binary matrices) has at most $2e^{-k/8}$ mass on matrices M that witness the plants on k_1 rows for $k_1 \notin [-k/2, 2k]$. Therefore, conditioned on the algorithm seeing samples M s.t. $k/2 \leq k_1 \leq 2k$, $A_{n,k}$ has success probability $\geq p(n, k) - 2e^{-k/8}$.

Consider the following algorithm for APBC(n, k') that takes as input an adjacency matrix G chosen randomly according to $\mathcal{D}_{avg}(n, k')$, and so, it contains a $k' \times k'$ plant $S_1 \times S_2$ that needs to be found.

Pick k s.t. $k'/2 < k < 2k'$ according to $B(n, k/n)$. From G , which has a plant of size $k' \times k'$ we'd like to obtain samples M with $k \times k'$ plants, or $k' \times k$ plants.

If $k'/2 < k \leq k'$, as in the proof of Proposition A.3 we can obtain a sample with a plant of size $k \times k'$ by sequentially replacing the n rows by random $\{0, 1\}^n$ vectors. As before, after running $A_{n,k'}$ on each of these n matrices the $k \times k'$ plant $S'_1 \times S_2$ will be output. To actually complete the plant on the left side, it is enough to look at the S_2 indices and output all the rows i s.t. all the entries at i and $j \in S_2$ take value 1. This takes $nT(n, k')$ time complexity.

If $k' < k < 2k'$ then $k/2 < k' < k$ and from G we can generate a sample matrix M with a plant on $k' \times k$ just as before but where now we replace the columns by random columns. This takes $nT(n, k) = O(nT(n, k'))$ time.

Therefore, $T'(n, k') = O(nT(n, k'))$ and $p'(n, k') = (1 - 2e^{-k'/8})(p(n, k) - 2e^{-k/8})$, since we conditioned on the event that $k'/2 < k \leq 2k'$. \square

B Relationship to Statistical Queries in Learning

We will now use Corollary 3.11 to demonstrate that our work generalizes the notion of statistical query dimension and the statistical query lower bounds from learning theory. In an instance of a PAC learning problem, the learner has access to random examples of an unknown boolean function $f' : X' \rightarrow \{-1, 1\}$ from a set of boolean functions \mathcal{C} (whenever necessary, we use $'$ to distinguish variables from the identically named ones in the context of general search problems). A random example is a pair including a point and its label $(x', c(x'))$ such that x' is drawn randomly from an unknown distribution D' . For $\epsilon > 0$, the goal of an ϵ -accurate learning algorithm is to find, with high probability, a boolean hypothesis h' for which $\Pr_{x' \sim D'}[h'(x') \neq f'(x')] \leq \epsilon$.

A statistical query (SQ) learning algorithm [34] has access to a statistical query oracle for the unknown function f' and distribution D' in place of random examples. A query to the SQ oracle is a function $\phi : X' \times \{-1, 1\} \rightarrow [-1, 1]$ that depends on both the example x' and its label ℓ . To such a query the oracle returns a value v which is within τ of $\mathbf{E}_{D'}[\phi(x', c(x'))]$, where τ is the tolerance parameter. A SQ algorithm does not depend on the randomness of examples and hence must always succeed.

Blum et al. [12] defined the *statistical query dimension* or SQ-DIM of a set of functions \mathcal{C} and distribution D' over X' as follows (we present a simplification and strengthening due to Yang [51]).

Definition B.1 ([12]). *For a concept class \mathcal{C} and distribution D' , $SQ-DIM(\mathcal{C}, D') = d'$ if d' is the largest value for which there exist d' functions $c_1, c_2, \dots, c_{d'} \in \mathcal{C}$ such that for every $i \neq j$, $|\langle c_i, c_j \rangle_{D'}| \leq 1/d'$.*

Blum et al. [12] proved that if a class of functions is learnable using only a polynomial number of statistical queries of inverse polynomial tolerance then its statistical query dimension is polynomial. Yang [51] strengthened their result and proved the following bound (see [44] for a simpler proof).

Theorem B.2 ([51]). *Let \mathcal{C} be a class of functions and D' be a distribution over X' and let $d' = SQ-DIM(\mathcal{C}, D')$. Then any SQ learning algorithm for \mathcal{C} over D' that makes q queries of tolerance $1/d'^{1/3}$ and outputs an ϵ -accurate hypothesis for $\epsilon \leq 1/2 - 1/(2d'^{1/3})$ satisfies that $q \geq d'^{1/3}/2 - 1$.*

In this result the distribution D' is fixed and known to the learner (such learning is referred to as *distribution-specific*) and it can be used to lower bound the complexity of learning \mathcal{C} even in a weak sense. Specifically, when the learning algorithm is only required to output a hypothesis h' such that $\Pr_{x' \sim D'}[h'(x') \neq c(x')] \leq 1/2 + \gamma'$ for some inverse polynomial γ' (or $\epsilon \leq 1/2 - \gamma'$).

It is not hard to see that we can cast this learning problem as a search problem over distributions, and by doing so, we will obtain that our statistical dimension implies a lower bound on learning which is slightly stronger than that of Yang [51].

Let $\mathcal{L} = (\mathcal{C}, D', \epsilon)$ be an instance of a distribution-specific learning problem of a class of functions \mathcal{C} over distribution D' to accuracy $1 - \epsilon$. We define the following search problem $\mathcal{Z}_{\mathcal{L}}$ over distributions. The domain is all the labeled points or $X = X' \times \{-1, 1\}$. When the target function equals $c \in \mathcal{C}$ the learning algorithm gets samples from the distribution D_c over X , where $D_c(x', c(x')) = D'(x')$ and $D_c(x', -c(x')) = 0$. Therefore we define the set of distributions over which we search to be $\mathcal{D}_{\mathcal{L}} = \{D_c \mid c \in \mathcal{C}\}$. Note that STAT oracle for D_c with tolerance τ is equivalent to the statistical query oracle for c over D' with tolerance τ . We can take the set of solutions $\mathcal{F}_{\mathcal{L}}$ over which a learning algorithm searches optimizes to be the set of all boolean functions over X' (an efficient learning algorithm can only output circuits of polynomial size but this distinction

is not important for our information-theoretic bounds). We define $\mathcal{Z}_{\mathcal{L}}$ to be the problem: given random samples from $D_c \in \mathcal{D}_{\mathcal{L}}$ find a hypothesis $f \in \mathcal{F}_{\mathcal{L}}$ that has accuracy of at least $1 - \epsilon$ on c . This is equivalent to saying $\mathbf{E}_{(x', \ell) \sim D_c} [f(x') = \ell] \geq 1 - \epsilon$.

We claim that SQ-DIM(\mathcal{C}, D')-based lower bound given in Theorem B.2 is special (and substantially simpler) case of our statistical dimension lower bound for $\mathcal{Z}_{\mathcal{L}}$ (Cor. 3.11).

Theorem B.3. *Let \mathcal{C} be a class of functions and D' be a distribution over X' and let $d' = \text{SQ-DIM}(\mathcal{C}, D')$. Denote by $\mathcal{L} = (\mathcal{C}, D', \epsilon)$ the instance of learning \mathcal{C} over D' for $\epsilon = 1/2 - 1/(2d'^{1/3})$. Then*

$$\text{SD}(\mathcal{Z}_{\mathcal{L}}, \gamma = 1/d', \beta = 1) \geq \left(d' - \frac{1}{1/d'^{2/3} - 1/d'} \right).$$

Proof. Let $c_1, c_2, \dots, c_{d'}$ be the almost uncorrelated functions in \mathcal{C} implied by the definition of SQ-DIM(\mathcal{C}, D'). We define the reference distribution D as the distribution for which for every $(x', \ell) \in X$, $D(x', \ell) = D'(x')/2$. We note that this ensures that $D(x', \ell)$ is non-vanishing only when $D'(x')$ is non-vanishing and hence the function $(\frac{D_c}{D} - 1)$ will be well-defined for all $c \in \mathcal{C}$. For every $c \in \mathcal{C}$, we have

$$\frac{D_c(x', c(x'))}{D(x', c(x'))} - 1 = 2 - 1 = 1 \quad \text{and} \quad \frac{D_c(x', -c(x'))}{D(x', -c(x'))} - 1 = 0 - 1 = -1.$$

Therefore, $\frac{D_c(x', \ell)}{D(x', \ell)} = \ell \cdot c(x')$. This implies that for any $c_i, c_j \in \mathcal{C}$,

$$\left\langle \frac{D_{c_i}}{D} - 1, \frac{D_{c_j}}{D} - 1 \right\rangle_D = \mathbf{E}_D[\ell \cdot c_i(x') \cdot \ell \cdot c_j(x')] = \mathbf{E}_{D'}[c_i(x') \cdot c_j(x')] = \langle c_i, c_j \rangle_{D'}.$$

Hence

1. for any $c \in \mathcal{C}$, $\left\| \frac{D_c(x)}{D(x)} - 1 \right\|_D^2 = 1$;
2. for any $i \neq j \leq d'$, $\left\langle \frac{D_{c_i}(x)}{D(x)} - 1, \frac{D_{c_j}(x)}{D(x)} - 1 \right\rangle_D \leq 1/d'$.

These properties imply that d' functions in \mathcal{C} give d' distributions in $\mathcal{D}_{\mathcal{L}}$ which are almost uncorrelated (relative to D). This is essentially the condition required to obtain a lower bound of d' on $\text{SD}(\mathcal{Z}_{\mathcal{L}}, 1/d', 1)$. The only issue is that we need to exclude distributions for which any given $f \in \mathcal{F}_{\mathcal{L}}$ is a valid solution (that is, f $(1 - \epsilon)$ -approximates c). We claim that it is easy to bound the number of such distributions for a fixed $f(x')$. First, note that the condition of $1 - \epsilon$ approximation of f for c states that $\mathbf{E}_{(x', \ell) \sim D_c} [f(x') = \ell] \geq 1 - \epsilon$ and therefore

$$\mathbf{E}_{(x', \ell) \sim D_c} [f(x') \cdot \ell] \geq 1 - 2\epsilon = 1/d'^{1/3}.$$

On the other hand, $\mathbf{E}_{(x', \ell) \sim D} [f(x') \cdot \ell] = 0$ and therefore $\mathbf{E}_{(x', \ell) \sim D_c} [f(x') \cdot \ell] - \mathbf{E}_{(x', \ell) \sim D} [f(x') \cdot \ell] \geq 1/d'^{1/3}$. This implies that if we view $f(x') \cdot \ell$ as a query function then expectations of the query function relative to D_c and D differ by at least $\tau = 1/d'^{1/3}$. In the proof of Theorem 3.4 we proved that the set of distributions that satisfies this condition has average correlation of at least τ^2 relative to D . In Corollary 3.11 we proved that at most $(\beta - \gamma)/(\tau^2 - \gamma)$ distributions with pairwise correlations (γ, β) can have average correlation of at least τ^2 . For our parameters this gives a bound of $\frac{1}{1/d'^{2/3} - 1/d'} = \frac{d'}{d'^{1/3} - 1}$ distributions. Hence for $m = d' - \frac{d'}{d'^{1/3} - 1}$ we obtain that for every $f \in \mathcal{F}_{\mathcal{L}}$ there exist m distributions $D_1, \dots, D_m \subseteq \{D_{c_1}, \dots, D_{c_{d'}}\} \setminus \mathcal{Z}_{\mathcal{L}}^{-1}(f)$ such that

1. for any $i \leq m$, $\left\| \frac{D_i(x)}{D(x)} - 1 \right\|_D = 1$;
2. for any $i \neq j \leq m$, $\left\langle \frac{D_i(x)}{D(x)} - 1, \frac{D_j(x)}{D(x)} - 1 \right\rangle_D \leq 1/d'$.

□

Applying Corollary 3.11, we get the following lower bound, which is twice larger than the $d^{1/3}/2 - 1$ bound of Yang [51].

Corollary B.4. *Let \mathcal{C} be a class of functions and D' be a distribution over X' , let $d' = \text{SQ-DIM}(\mathcal{C}, D')$ and let $\epsilon = 1/2 - 1/(2d^{1/3})$. Then any SQ learning algorithm requires at least $d^{1/3} - 2$ queries of tolerance $1/d^{1/3}$ to ϵ -accurately learn \mathcal{C} over D' .*

We also note that, unlike previous bounds for SQ algorithms our technique immediately gives lower bounds against randomized algorithms. To see this we note that in the proof of Theorem B.3 we have that for every $f \in \mathcal{F}_{\mathcal{L}}$, $|\mathcal{Z}_{\mathcal{L}}^{-1}(f)| \leq \frac{d'}{d^{1/3}-1}$. This implies a solution set size bound $\eta = \frac{d'}{d^{1/3}-1}/d' = \frac{1}{d^{1/3}-1}$. Now, using Lemma 3.12 together with Theorem 3.8, it is easy to obtain the following corollary.

Corollary B.5. *Let \mathcal{C} be a class of functions and D' be a distribution over X' , let $d' = \text{SQ-DIM}(\mathcal{C}, D')$ and let $\epsilon = 1/2 - 1/(2d^{1/3})$. Then any randomized SQ learning algorithm that succeeds with probability at least $1/2$, requires at least $d^{1/3}/2 - 2$ queries of tolerance $1/d^{1/3}$ to ϵ -accurately learn \mathcal{C} over D' .*

B.1 Honest Statistical Queries

We now turn to the Honest SQ model [29, 50], which inspired our notion of unbiased statistical algorithms. In the Honest SQ model, the learner has access to an HSQ oracle and can again evaluate queries which are a function of the data points and their labels. As in our 1-STAT oracle, the queries are evaluated on an “honest” sample drawn from the target distribution. More precisely, the HSQ oracle accepts a function $\phi : X' \times \{-1, 1\} \rightarrow \{-1, 1\}$ and a sample size $t > 0$, draws $x'_1, \dots, x'_t \sim D'$, and returns the value $\frac{1}{t} \sum_{i=1}^t \phi(x'_i, c(x'_i))$. The total sample count of an algorithm is the sum of the sample sizes it passes to HSQ.

We note that using our one-sample-per-query-function oracle 1-STAT one can simulate estimation of queries from larger samples in the straightforward way while obtaining the same sample complexity. Therefore HSQ is equivalent to our 1-STAT oracle.

We first observe that our direct simulation in Theorem 3.5 implies that the Honest SQ learning model is equivalent (up to polynomial factors) to the SQ learning model. We are not aware of this observation having been made before (although Valiant [47] implicitly uses it to show that evolvable concept classes are also learnable in the SQ model).

In addition, by adapting Theorem 3.8 to this setting as we did in Corollary B.5 (that is using Lemma 3.12 and the bound on the solution set size) we can directly derive sample complexity bounds on honest statistical query algorithms for learning.

Corollary B.6. *Let \mathcal{C} be a class of functions, D' be a distribution over X' , $d' = \text{SQ-DIM}(\mathcal{C}, D')$ and $\epsilon = 1/2 - 1/(2d^{1/3})$. Then any Honest SQ algorithm that, with probability at least $2/3$, ϵ -accurately learns \mathcal{C} over D' requires $\Omega(\sqrt{d'})$ queries.*

This gives a simpler proof for the bound of Yang [51] up to polynomial factors (Yang gets a bound of $\Omega(d'/\log d')$ using a stronger $1/d'^3$ upper bound on correlations).