The Query Complexity of Witness Finding*

Akinori Kawachi[†] Benjamin Rossman[‡] Osamu Watanabe[§]

September 10, 2016

Abstract

We study the following information-theoretic *witness finding problem*: for a hidden nonempty subset W of $\{0,1\}^n$, how many non-adaptive randomized queries (yes/no questions about W) are needed to guess an element $x \in \{0,1\}^n$ such that $x \in W$ with probability $> 1/2$? Motivated by questions in complexity theory, we prove tight lower bounds with respect to a few different classes of queries:

- We show that the *monotone* query complexity of witness finding is $\Omega(n^2)$. This matches an $O(n^2)$ upper bound from the Valiant-Vazirani Isolation Lemma [8].
- We also prove a tight $\Omega(n^2)$ lower bound for the class of *NP queries* (queries defined by an NP machine with an oracle to W). This shows that the classic search-to-decision reduction of Ben-David, Chor, Goldreich and Luby [3] is optimal in a certain black-box model.
- Finally, we consider the setting where W is an affine subspace of $\{0,1\}^n$ and prove an $\Omega(n^2)$ lower bound for the class of *intersection queries* (queries of the form "$W \cap S \neq \emptyset$?" where S is a fixed subset of $\{0,1\}^n$). Along the way, we show that every monotone property defined by an intersection query has an exponentially sharp threshold in the lattice of affine subspaces of $\{0,1\}^n$.

1 Introduction

We initiate a study of the following information-theoretic search problem, parameterized by a family $\mathcal{W}$ of subsets of $\{0,1\}^n$ and a family $\mathcal{Q}$ of functions $\mathcal{W} \to \{\top, \bot\}$ (i.e. yes/no questions about elements of $\mathcal{W}$, which we refer to as "queries").

Question 1.1. *What is the minimum number of non-adaptive randomized queries from $\mathcal{Q}$ required to guess an element $x \in \{0,1\}^n$ such that $\mathbb{P}[x \in W] > 1/2$ for every nonempty $W \in \mathcal{W}$?*

Formally, Question 1.1 asks for a joint distribution $(\mathbf{Q}_1, \dots, \mathbf{Q}_m)$ on $\mathcal{Q}^m$ together with a function $f : \{\top, \bot\}^m \to \{0,1\}^n$ such that

$$\mathbb{P}[f(\mathbf{Q}_1(W), \dots, \mathbf{Q}_m(W)) \in W] > 1/2$$

for every nonempty $W \in \mathcal{W}$. We emphasize that randomized queries $\mathbf{Q}_1, \dots, \mathbf{Q}_m$ are non-adaptive, though not necessarily independent.[1]

We refer to Question 1.1 as the *witness finding problem* and to its answer, $m = m(\mathcal{W}, \mathcal{Q})$, as the *$\mathcal{Q}$-query complexity of $\mathcal{W}$-witness finding.* (We introduce the terminology "witness finding" to distinguish this

*This work was supported in part by the ELC (Exploring the Limits of Computation) project under KAKENHI, Grant Number 24106002, 2406008, and 2406009, KAKENHI Grant-in-Aid for Scientific Research (A) Grant Number 24240001, and JST ERATO Kawarabayashi Large Graph Project.

[†]Dept. of Mathematical and Computing Sciences, Tokyo Institute of Technology, Ookayama 2-12-1, Meguro-ku, Tokyo 152-8552, Japan

[‡]National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan

[§]Dept. of Mathematical and Computing Sciences, Tokyo Institute of Technology, Ookayama 2-12-1, Meguro-ku, Tokyo 152-8552, Japan

[1]That is, $\mathbf{Q}_1$ and $\mathbf{Q}_2$ may be dependent random variables. However, conditioned on $\mathbf{Q}_1 = Q_1$, $\mathbf{Q}_2$ cannot depend on the answer $Q_1(W) \in \{\top, \bot\}$.

information-theoretic problem from traditional computational search problems where the solution space is determined by an input, such as a boolean formula φ in the case of the search problem for SAT.) Note that $m(\mathcal{W}, \mathcal{Q})$ is monotone increasing with respect to $\mathcal{W}$ and monotone decreasing with respect to $\mathcal{Q}$. In this paper, we mainly study the setting where $\mathcal{W}$ is the set of all subsets of $\{0,1\}^n$. Here, to simplify notation, we simply write $m(\mathcal{Q})$ and speak of the *$\mathcal{Q}$-query complexity of witness finding*.

Our main results are tight lower bounds on $m(\mathcal{Q})$ for a few specific classes of queries (namely, *intersection queries*, *monotone queries* and *NP queries*). However, before defining these classes and stating our results formally, let us first dispense with the trivial cases where $\mathcal{Q}$ is the class All of all possible queries or the class Direct of *direct queries* of the form "$x \in W$?" where $x \in \{0,1\}^n$. It is easy to see that $m(\mathsf{All}) = n$ and $m(\mathsf{Direct}) = 2^n - 1$. Both lower bounds $m(\mathsf{All}) \geq n$ and $m(\mathsf{Direct}) \geq 2^n - 1$ follow from considering the random singleton witness set $\{\mathbf{x}\}$ where $\mathbf{x}$ is uniform in $\{0,1\}^n$. The upper bound $m(\mathsf{Direct}) \leq 2^n - 1$ is obvious, while the upper bound $m(\mathsf{All}) \leq n$ comes via deterministic queries $Q_1, \dots, Q_n$ where $Q_i(W)$ asks for the ith coordinate in the lexicographically minimal element of W.

In contrast to the non-adaptive setting, we remark that *adaptive* query complexity of the witness finding problem is not very interesting, as n adaptive queries (of a simple kind) are necessary and sufficient. To find a witness in W using adaptive queries, we first ask "Does W contain an element whose first coordinate is 1?" Depending on the answer (yes or no), we then ask "Does W contain an element whose first coordinate is (1 or 0) and whose second coordinate is 1?" In this way, we learn an element $x \in W$ using n adaptive deterministic queries. An easy information-theoretic argument shows that n is also a lower bound. Therefore, the adaptive query complexity of witness finding in n for any sufficiently non-trivial class $\mathcal{Q}$ (in particular, $\mathcal{Q}$ containing all *intersection queries*, which will define next).

1.1 Intersection Queries and Monotone Queries

The first class $\mathcal{Q}$ that we consider, for which the question of $m(\mathcal{Q})$ is nontrivial, is the class $\mathsf{Intersection}$ of *intersection queries* of the form "$S \cap W \neq \emptyset$?" for fixed $S \subseteq \{0,1\}^n$. As we now explain, the Valiant-Vazirani Isolation Lemma [8] gives an elegant upper bound of $m(\mathsf{Intersection}) = O(n^2)$. First, note that if W is a singleton $\{w\}$, then n non-adaptive intersection queries suffice to learn w: for $1 \leq i \leq n$, we ask "$S_i \cap W \neq \emptyset$?" where $S_i = \{x \in \{0,1\}^n : x_i = 0\}$. Moreover, by asking n additional intersection queries "$T_i \cap W \neq \emptyset$?" where $T_i = \{x \in \{0,1\}^n : x_i = 1\}$, we can learn whether or not W is a singleton, in addition to learning w in the event that $W = \{w\}$. The Valiant-Vazirani Isolation Lemma gives a distribution $\mathbf{X}$ on subsets of $\{0,1\}^n$ such that $\mathbb{P}[|W \cap \mathbf{X}| = 1] = \Omega(1/n)$ for every nonempty $W \subseteq \{0,1\}^n$. By taking $s = O(n)$ independent copies of $\mathbf{X}_1, \dots, \mathbf{X}_s$ of this distribution $\mathbf{X}$, we have $\mathbb{P}[\bigvee_{j=1}^{s} |W \cap \mathbf{X}_j| = 1] > 1/2$ for every nonempty $W \subseteq \{0,1\}^n$. We now get a witness finding procedure which makes $2ns = O(n^2)$ randomized intersection queries for sets $\mathbf{S}_{i,j} := S_i \cap \mathbf{X}_j$ and $\mathbf{T}_{i,j} := T_i \cap \mathbf{X}_j$. (By now the reader will have noticed our convention of designating random variables by bold letters.)

The present paper started out as an investigation into the question whether $O(n^2)$ is a tight upper bound on $m(\mathsf{Intersection})$. This question arose from work of Dell, Kabanets, van Melkebeek and Watanabe [7], who showed that the Valiant-Vazirani Isolation Lemma is optimal among so-called black-box isolation procedures:

Theorem 1.2 ([7]). *For every distribution $\mathbf{X}$ on subsets of $\{0,1\}^n$, there exists nonempty $W \subseteq \{0,1\}^n$ such that $\mathbb{P}[|\mathbf{X} \cap W| = 1] = O(1/n)$.*

Borrowing an idea from the proof of Theorem 1.2 (namely, a particular distribution on subsets of $\{0,1\}^n$), we were able to show $m(\mathsf{Intersection}) = \Omega(n^2)$. (Note that Theorem 1.2 can be derived from this lower bound, as any black-box isolation procedure with success probability $\omega(1/n)$ would show that $m(\mathsf{Intersection}) = o(n^2)$ by the argument sketched above.) As a natural next step, we considered the class of *monotone queries*, that is, $Q : \wp(\{0,1\}^n) \to \{\top, \bot\}$ such that $Q(W) = \top \Rightarrow Q(W') = \top$ for all $W \subseteq W' \subseteq \{0,1\}^n$. Note that intersection queries are monotone, hence $n \leq m(\mathsf{Monotone}) \leq m(\mathsf{Intersection}) = \Theta(n^2)$. Generalizing our lower bound for intersection queries, we were able to prove the stronger result:

Theorem 1.3. *The monotone query complexity of witness finding, $m(\mathsf{Monotone})$, is $\Omega(n^2)$.*

We present the proof of Theorem 1.3 in §2. The proof uses an entropy argument, which hinges on the threshold behavior of monotone queries (in particular, the theorem of Bollobás and Thomason [4]).

1.2 NP Queries

Another motivation for studying Question 1.1 comes from a question concerning search-to-decision reductions. In the context of SAT, a *search-to-decision reduction* is an algorithm which, given a boolean function $\varphi(x_1, \ldots, x_n)$, constructs a satisfying assignment $x \in \{0,1\}^n$ for φ (if one exists) using an oracle for the SAT decision problem. The standard P^{NP} search-to-decision reduction uses n adaptive deterministic queries. In the setting of non-adaptive randomized queries, Ben-David, Chor, Goldreich and Luby [3] (using the Valiant-Vazirani Isolation Lemma) gave a $\mathrm{BPP}^{\mathrm{NP}}_{||}$ search-to-decision reduction with $O(n^2)$ queries. ($\mathrm{BPP}^{\mathrm{NP}}_{||}$ is the class of BPP algorithms with non-adaptive (parallel) query access to an NP oracle.)

We are interested in lower bounds for the query complexity of search-to-decisions for SAT. Of course, any nontrivial lower bound would separate P from NP. However, we can consider a "black-box" setting where, instead of receiving a boolean formula $\varphi(x_1, \ldots, x_n)$ as input, the $\mathrm{BPP}^{\mathrm{NP}}_{||}$ algorithm (including both the BPP machine and the NP machine) are given input 1^n as well as an oracle to the set $\{x \in \{0,1\}^n : x$ is a satisfying assignment for $\varphi\}$. On inspection, it is clear that the reduction of Ben-David et al. (which is indifferent to the syntax of the boolean formula φ) carries over to this black-box setting. Thus, we have the upper bound:

Theorem 1.4 (follows from [3])**.** *There is a $\mathrm{BPP}^{\mathrm{NP}}_{||}$ algorithm which solves the black-box satisfiability search problem with $O(n^2)$ queries.*

Motivated by this connection to complexity theory, we next set our sights on the question whether $O(n^2)$ is tight in Theorem 1.4. To fit the question into the framework of Question 1.1, we define the class of *NP queries* as follows.

Definition 1.5. Informally, an *NP query* is a query Q given by an NP machine M with an oracle to W where $Q(W) = M^W(1^n)$ (i.e. $Q(W) = \top \Leftrightarrow M^W$ has an accepting computation on input 1^n). Formally, an *NP query* is a sequence $Q = (Q^1, Q^2, \ldots)$ of queries $Q^n : \wp(\{0,1\}^n) \to \{\top, \bot\})$ such that there exists a single NP machine $M^{()}$ (with an unspecified oracle) where $Q^n(W) = M^W(1^n)$ for every $W \subseteq \{0,1\}^n$. An *ensemble of NP queries* is a sequence $(Q_1, \ldots, Q_m)$ of NP queries given by NP machines $M_1, \ldots, M_m$ which have a common upper bound $t(n) = n^{O(1)}$ on their running time.

The NP query complexity of witness finding, $m(\mathsf{NP})$, gives a lower bound on the query complexity of $\mathrm{BPP}^{\mathrm{NP}}_{||}$ algorithms solving the black-box satisfiability search problem. Note that NP queries and monotone queries are incomparable: NP queries clearly need not be monotone, while it can be shown that the monotone "majority" query (defined by $Q_{\mathrm{maj}}(W) = \top$ iff $|W| \geq 2^{n-1}$) is not an NP query.[2] Nevertheless, we show that every NP query can be *well-approximated* by a monotone query (Lemma 3.2). Using this result together with our lower bound for $m(\mathsf{Monotone})$, we show:

Theorem 1.6. *The NP query complexity of witness finding, $m(\mathsf{NP})$, is $\Omega(n^2)$.*

Theorem 1.6 thus establishes the optimality of the search-to-decision reduction of Ben-David et al. in the black-box setting. The proof is presented in §3.

1.3 Affine Witness Sets

Finally, we consider the setting where $\mathcal{W}$ is the set of affine subspaces of $\{0,1\}^n$. (Recall that *affine subspaces* of $\{0,1\}^n$ are the solution sets of systems of linear equations in n variables over the 2-element field.) Here, for

[2] Due to uniformity issues, it does not make sense to compare the classes of NP queries and intersection queries. However, for a natural notion of *non-uniform NP queries*, every intersection query "$S \cap W \neq \emptyset$?" is a non-uniform NP query where the NP machine M hardwires S using 2^n advice bits, non-deterministically guesses $x \in S$ and simply verifies that $x \in W$ using one oracle call to W.

a class of queries $\mathcal{Q}$, we write $m_{\text{affine}}(\mathcal{Q})$ and speak of the *$\mathcal{Q}$-query complexity of affine witness finding*. While $m_{\text{affine}}(\mathcal{Q}) \le m(\mathcal{Q})$ by definition, intuitively the affine witness finding problem is easier because there are only $2^{O(n^2)}$ possibilities for W, as opposed to 2^{2^n}. One motivation for studying the affine setting comes from the observation that lower bounds on $m_{\text{affine}}(\mathsf{NP})$ imply lower bounds on the complexity of the black-box satisfiability search problem on *polynomial-size* boolean formulas, since every affine subspace of $\{0,1\}^n$ is the set of satisfying assignments to a polynomial-size boolean formula of n variables. While we were unable to prove any nontrivial lower bounds on $m_{\text{affine}}(\mathsf{Monotone})$ or $m_{\text{affine}}(\mathsf{NP})$, we did get a result for intersection queries:

Theorem 1.7. *The intersection query complexity of affine witness finding, $m_{\text{affine}}(\mathsf{Intersection})$, is $\Omega(n^2)$.*

The proof is presented in §4. Along the way, we show that every monotone property defined by an intersection query has an *exponentially sharp threshold* in the lattice of affine subspaces of $\{0,1\}^n$ (Theorem 4.5). This raises the question whether all monotone properties have an exponentially sharp threshold in the affine lattice (Question 4.1). We show that a positive answer to this question implies $m_{\text{affine}}(\mathsf{Monotone}) = \Omega(n^2)$.

1.4 From Threshold Behavior to Lower Bounds

The lower bound arguments of Theorems 1.3 and 1.7 follow a common pattern and hinge on the threshold behavior of the queries considered. For any fixed monotone query Q, if we pick W uniform at random from all subsets of $\{0,1\}^n$ of a given size 2^k, then $Q(W)$ jumps from almost surely $\bot$ to almost surely $\top$ very quickly as a function of k. By choosing k uniformly in $[n]$, we show that the random variable $Q(W)$ has entropy $O(1/n)$. On the other hand, we show that any random variable $z \in \{0,1\}^n$ satisfying $\mathbb{P}[z \in W] > 1/2$ has entropy $\Omega(n)$. These two inequalities (together with Yao's principle) imply a lower bound of $\Omega(n^2)$ on $m(\mathsf{Monotone})$. We get a lower bound on $m_{\text{affine}}(\mathsf{Intersection})$ by a similar argument where W is uniformly distributed among affine subspaces of $\{0,1\}^n$ of dimension k.

2 Lower Bound for Monotone Queries

In this section, we prove Theorem 1.3 ($m(\mathsf{Monotone}) = \Omega(n^2)$) using an information-theoretic argument. Let $H : [0,1] \to [0,1]$ denote the binary entropy function $H(p) := p\log(1/p) + (1-p)\log(1/(1-p))$. For finite random variables $\mathbf{X}$ and $\mathbf{Y}$, entropy $\mathbb{H}(\mathbf{X})$ and relative entropy $\mathbb{H}(\mathbf{X} \mid \mathbf{Y})$ are defined by

$$\mathbb{H}(\mathbf{X}) := \sum_{x \in \text{Supp}(\mathbf{X})} \mathbb{P}[\mathbf{X} = x] \cdot \log(1/\,\mathbb{P}[\mathbf{X} = x]),$$

$$\mathbb{H}(\mathbf{X} \mid \mathbf{Y}) := \sum_{y \in \text{Supp}(\mathbf{Y})} \mathbb{P}[\mathbf{Y} = y] \cdot \mathbb{H}(\mathbf{X} \mid \mathbf{Y} = y).$$

(Here $\mathbb{H}(\mathbf{X} \mid \mathbf{Y} = y)$ is the entropy of the marginal distribution of $\mathbf{X}$ conditioned on $\mathbf{Y} = y$.) We assume familiarity with the basic properties of entropy, namely the chain rule $\mathbb{H}(\mathbf{X},\mathbf{Y}) = \mathbb{H}(\mathbf{X}) + \mathbb{H}(\mathbf{Y} \mid \mathbf{X})$, the fact that $\mathbb{H}(f(\mathbf{X})) \le \mathbb{H}(\mathbf{X})$ for every deterministic function f of $\mathbf{X}$, and the fact $\mathbb{H}(\mathbf{X}) \le \log|\text{Supp}(\mathbf{X})|$ with equality iff $\mathbf{X}$ is uniform (for more background, see [6]).

Our lower bound uses a standard averaging argument (Yao's principle) to invert the role of randomness in the definition of $m(\mathcal{W},\mathcal{Q})$. For completeness, the proof is included in Appendix A.

Lemma 2.1. *Suppose $\mathbf{W}$ is a random variable on $\mathcal{W} \setminus \{\emptyset\}$ such that for all $Q_1,\dots,Q_m \in \mathcal{Q}$ and every function $f : \{\top,\bot\}^m \to \{0,1\}^n$,*

$$\mathbb{P}[f(Q_1(\mathbf{W}),\dots,Q_m(\mathbf{W})) \in \mathbf{W}] \le 1/2.$$

Then the $\mathcal{Q}$-query complexity of $\mathcal{W}$-witness finding is greater than m.

We now define a particular random subset $\mathbf{W}$ of $\{0,1\}^n$. First, for each $k \in \{1,\dots,n\}$, let $\mathbf{W}_k$ be the random subset of $\{0,1\}^n$ that contains each $x \in \{0,1\}^n$ independently with probability 2^{k-n}. Note that $\mathbf{W}_n = \{0,1\}^n$ with probability 1. We also define $\mathbf{W}_0 := \emptyset$ (with probability 1). Note that $\mathbf{W}_k$ strictly stochastically dominates $\mathbf{W}_{k-1}$ for all $k \in \{1,\dots,n\}$. In other words, for every monotone property Q of subsets of $\{0,1\}^n$, we have $\mathbb{P}[\mathbf{W}_0 \in Q] \le \mathbb{P}[\mathbf{W}_1 \in Q] \le \dots \le \mathbb{P}[\mathbf{W}_n \in Q]$; moreover, if Q is non-trivial (i.e. $\emptyset \notin Q$ and $\{0,1\}^n \in Q$), then

$$0 = \mathbb{P}[\mathbf{W}_0 \in Q] < \mathbb{P}[\mathbf{W}_1 \in Q] < \dots < \mathbb{P}[\mathbf{W}_n \in Q] = 1.$$

Finally, we define $\mathbf{W}$ as the random set $\mathbf{W}_{\mathbf{k}}$ where $\mathbf{k}$ is a uniform random integer in the set $\{1,\dots,n/2\}$.[3] That is $\mathbf{W}$ is obtained by first sampling $k \leftarrow \mathbf{k}$ and then sampling from $\mathbf{W}_k$. (A similar distribution was considered by Dell et al. [7] in proving an upper bound of $O(1/n)$ on the success probability of black-box isolation procedures.)

The following lemma is a special case of the Bollobás-Thomason Theorem [4] (informally, "every monotone increasing property of subsets of a fixed set has a threshold function"). For completeness, a simple self-contained proof is included in Appendix B.

Lemma 2.2. *Let Q be a non-trivial monotone increasing property of subsets of $\{0,1\}^n$. For $k \in \{0,\dots,n\}$, let $p_k := \mathbb{P}[\mathbf{W}_k$ has property $Q]$. Let $\theta \in \{0,\dots,k-1\}$ be the unique index such that $p_\theta \le 1/2 < p_{\theta+1}$. Then*

$$p_{\theta-i} \le 2^{-i} \ln 2 \qquad \text{for all } 0 \le i \le \theta, \tag{1}$$

$$p_{\theta+i+1} \ge 1 - 2^{-2^i} \qquad \text{for all } 0 \le i \le n-\theta-1, \tag{2}$$

$$H(p_k) \le (|\theta-k|+1)/2^{|\theta-k|-1} \qquad \text{for all } 0 \le k \le n. \tag{3}$$

Using Lemma 2.2(3), we prove a sharp bound on the relative entropy $Q(\mathbf{W} \mid \mathbf{k})$ for all monotone queries Q.

Lemma 2.3. *$\mathbb{H}(Q(\mathbf{W}) \mid \mathbf{k}) = O(1/n)$ for every monotone query Q.*

Proof. If Q is identically $\perp$ or $\top$, then the statement is trivial (as $\mathbb{H}(Q(\mathbf{W}) \mid \mathbf{k}) = 0$). So assume Q is a non-trivial monotone query and let $p_0,\dots,p_n$ and θ be as in Lemma 2.2. Then

$$\begin{aligned}\mathbb{H}(Q(\mathbf{W}) \mid \mathbf{k}) &= \sum_{k=0}^{n/2} \mathbb{P}[\mathbf{k}=k] \cdot \mathbb{H}(Q(\mathbf{W}_k)) \\ &= \frac{2}{n}\sum_{k=1}^{n/2} H(p_k) \;\le\; \frac{2}{n}\sum_{k=1}^{n/2} \frac{|\theta-k|+1}{2^{|\theta-k|-1}} \;\le\; \frac{4}{n}\sum_{i=0}^{\infty}\frac{i+1}{2^{i-1}} \;\le\; \frac{32}{n}.\end{aligned}$$

∎

The next lemma relates the entropy of an arbitrary random variable $\mathbf{z}$ on $\{0,1\}^n$ to the probability that $\mathbf{z} \in \mathbf{W}$.

Lemma 2.4. *For every random variable $\mathbf{z}$ on $\{0,1\}^n$ (not necessarily independent of $\mathbf{W}$),*

$$\mathbb{P}[\mathbf{z} \in \mathbf{W}] \le \frac{4}{n}\mathbb{H}(\mathbf{z}) + \frac{1}{2^{n/4}}.$$

Proof. Define $S \subseteq \{0,1\}^n$ by $S := \{x \in \{0,1\}^n : \mathbb{P}[\mathbf{z}=x] \ge 2^{-n/4}\}$. Note that

$$\mathbb{P}[\mathbf{z} \in \mathbf{W}] \le \mathbb{P}[\mathbf{z} \notin S] + \mathbb{P}[S \cap \mathbf{W} \ne \emptyset].$$

We bound each these righthand probabilities. First, by definition of S and $\mathbb{H}(\mathbf{z})$,

$$\mathbb{P}[\mathbf{z} \notin S] \;=\; \sum_{x \in \{0,1\}^n \setminus S} \mathbb{P}[\mathbf{z}=x] \;\le\; \sum_{x \in \{0,1\}^n \setminus S} \mathbb{P}[\mathbf{z}=x]\frac{\log(1/\mathbb{P}[\mathbf{z}=x])}{n/4} \;\le\; \frac{4}{n}\mathbb{H}(\mathbf{z}).$$

[3] We treat $n/2$ as an integer (i.e. an abbreviation for $\lfloor n/2 \rfloor$).

(Here we used $x \notin S \Rightarrow \mathbb{P}[\mathbf{z} = x] < 2^{-n/4} \Rightarrow \log(1/\mathbb{P}[\mathbf{z} = x]) > n/4$.) Finally, noting that $|S| \leq 2^{n/4}$ and $\mathbb{P}[x \in \mathbf{W}] < 2^{-n/2}$ for all $x \in \{0,1\}^n$, we have

$$\mathbb{P}[\mathbf{W} \cap S \neq \emptyset] \;\leq\; \sum_{x \in S} \mathbb{P}[x \in \mathbf{W}] \;<\; \frac{1}{2^{n/4}}. \qquad \square$$

Combining Lemmas 2.3 and 2.4, we get our main lemma:

Lemma 2.5. *For all monotone queries $Q_1, \ldots, Q_m$ and every function $f : \{\top, \bot\}^m \to \{0,1\}^n$,*

$$\mathbb{P}[f(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W})) \in \mathbf{W}] \leq O(m/n^2) + o(1).$$

Proof. By standard entropy inequalities,

$$\begin{aligned}
\mathbb{H}(f(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W}))) &\leq \mathbb{H}(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W})) \\
&\leq \mathbb{H}(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W}), \mathbf{k}) \\
&= \mathbb{H}(\mathbf{k}) + \mathbb{H}(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W}) \mid \mathbf{k}) \\
&\leq \mathbb{H}(\mathbf{k}) + \mathbb{H}(Q_1(\mathbf{W}) \mid \mathbf{k}) + \cdots + \mathbb{H}(Q_m(\mathbf{W}) \mid \mathbf{k}).
\end{aligned}$$

Since $\mathbb{H}(\mathbf{k}) = \log(n/2)$ and $\mathbb{H}(Q_i(\mathbf{W}) \mid \mathbf{k}) = O(1/n)$ for all i by Lemma 2.3, we have

$$\mathbb{H}(f(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W}))) \leq O(m/n) + \log n.$$

Since $f(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W}))$ is a random variable on $\{0,1\}^n$, we can apply Lemma 2.4 to get

$$\begin{aligned}
\mathbb{P}[f(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W})) \in \mathbf{W}] &\leq \frac{4}{n}\mathbb{H}(f(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W}))) + \frac{1}{2^{n/4}} \\
&\leq O(m/n^2) + \frac{4\log n}{n} + \frac{1}{2^{n/4}} \\
&= O(m/n^2) + o(1). \qquad \square
\end{aligned}$$

Finally, we prove the main theorem of this section.

Theorem 2.3. (restated) *The monotone query complexity of witness finding, $m(\mathsf{Monotone})$, is $\Omega(n^2)$.*

Proof. Let $m = m(\mathsf{Monotone})$. By Lemma 2.1, there exist monotone queries $Q_1, \ldots, Q_m$ and a function $f : \{\top, \bot\}^m \to \{0,1\}^n$ such that

$$\mathbb{P}[f(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W})) \in \mathbf{W} \mid \mathbf{W} \neq \emptyset] > 1/2.$$

Note that $\mathbb{P}[\mathbf{W} = \emptyset] = (2/n)\sum_{k=1}^{n/2} \mathbb{P}[\mathbf{W}_k = \emptyset] = (2/n)\sum_{k=1}^{n/2}(1-2^{n-k})^{2^n} \leq (2/n)\sum_{k=1}^{n/2}\exp(2^{-k}) = O(1/n)$. Therefore, $\mathbb{P}[\mathbf{W} \neq \emptyset] = 1 - o(1)$. By Lemma 2.5,

$$\begin{aligned}
\mathbb{P}[f(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W})) \in \mathbf{W} \mid \mathbf{W} \neq \emptyset] &= \frac{\mathbb{P}[f(Q_1(\mathbf{W}), \ldots, Q_m(\mathbf{W})) \in \mathbf{W}]}{\mathbb{P}[\mathbf{W} \neq \emptyset]} \\
&\leq O(m/n^2) + o(1).
\end{aligned}$$

It follows that $1/2 < O(m/n^2) + o(1)$ and hence $m = \Omega(n^2)$. $\square$

3 Lower Bound for NP Queries

In this section, we prove Theorem 1.6 ($m(\mathsf{NP}) = \Omega(n^2)$). The main idea in the proof involves showing that every NP query is well-approximated by a monotone query. First, we give a normal form for NP queries.

Lemma 3.1. *For every NP query Q, there exists a sequence $(A_1, B_1), \dots, (A_s, B_s)$ where $A_i, B_i \subseteq \{0,1\}^n$ and $|A_i|, |B_i| \le n^{O(1)}$ and $A_i \cap B_i = \emptyset$ such that for all $W \subseteq \{0,1\}^n$,*

$$Q(W) = \top \iff \bigvee_{i=1}^{s} (A_i \subseteq W) \wedge (B_i \cap W = \emptyset).$$

Proof. Let $M^{()}$ be the nondeterministic Turing machine (with an unspecified oracle) which defines Q, that is, $Q(W) = M^W(1^n)$. Let $t = n^{O(1)}$ be the maximum running time of $M^{()}$. For each accepting computation of $M^{()}$ on input 1^n, there is a sequence $\sigma = ((x_1, y_1), \dots, (x_{t'}, y_{t'})) \in (\{0,1\}^n \times \{\top, \bot\})^{t'}$, $t' \le t$, such that the computation makes oracle calls $x_1, \dots, x_{t'}$ and receives answers $y_1, \dots, y_{t'}$. Let $A_\sigma := \{x_i : y_i = \top\}$ and $B_\sigma := \{x_i : y_i = \bot\}$ and note that $|A_\sigma|, |B_\sigma| \le t' \le t$ and $A_\sigma \cap B_\sigma = \emptyset$. Let $(A_1, B_1), \dots, (A_s, B_s)$ enumerate pairs (A_σ, B_σ) over all σ corresponding to accepting computations of $M^{()}$. This sequence $(A_1, B_1), \dots, (A_s, B_s)$ satisfies the conditions of the lemma. □

The next lemma gives the approximation of NP queries by monotone queries. Let $\mathbf{W}$ continue to denote the random subset of $\{0,1\}^n$ defined in the previous section.

Lemma 3.2. *For every NP query Q, there is a monotone query Q^+ such that* $\mathbb{P}[Q(\mathbf{W}) \neq Q^+(\mathbf{W})] = 2^{-\Omega(n)}$.

Proof. Let $(A_1, B_1), \dots, (A_s, B_s)$ be as in Lemma 3.1. Define Q^+ by

$$Q^+(W) = \top \overset{\text{def}}{\iff} \bigvee_{i=1}^{s} (A_i \subseteq W).$$

Clearly, Q^+ is a monotone query and $Q(W) \Rightarrow Q^+(W)$ (i.e. $Q(W) = \top$ implies $Q^+(W) = \top$). We have

$$\begin{aligned}
\mathbb{P}\Big[Q(\mathbf{W}) \neq Q^+(\mathbf{W})\Big] &= \mathbb{P}\Big[\neg Q(\mathbf{W}) \wedge Q^+(\mathbf{W})\Big] \\
&= \mathbb{P}\Big[\Big(\bigwedge_{i=1}^{s} (A_i \nsubseteq \mathbf{W}) \vee (B_i \cap \mathbf{W} \neq \emptyset)\Big) \wedge \Big(\bigvee_{i=1}^{s} (A_i \subseteq \mathbf{W})\Big)\Big] \\
&\le \mathbb{P}\Big[\bigvee_{i=1}^{s} (B_i \cap \mathbf{W} \neq \emptyset) \wedge (A_i \subseteq \mathbf{W}) \wedge \bigwedge_{j=1}^{i-1} (A_j \nsubseteq \mathbf{W})\Big] \\
&\le \max_i \mathbb{P}\Big[B_i \cap \mathbf{W} \neq \emptyset \;\Big|\; (A_i \subseteq \mathbf{W}) \wedge \bigwedge_{j=1}^{i-1} (A_j \nsubseteq \mathbf{W})\Big], \qquad (4)
\end{aligned}$$

where this last inequality is justified by the fact that events $\{(A_i \subseteq \mathbf{W}) \wedge \bigwedge_{j=1}^{i-1} (A_j \nsubseteq \mathbf{W})\}$ are mutually exclusive over $i \in \{1, \dots, s\}$.

Now fix i which maximizes (4). We claim that

$$\mathbb{P}\Big[B_i \cap \mathbf{W} \neq \emptyset \;\Big|\; (A_i \subseteq \mathbf{W}) \wedge \bigwedge_{j=1}^{i-1} (A_j \nsubseteq \mathbf{W})\Big] \le \mathbb{P}\Big[B_i \cap \mathbf{W} \neq \emptyset\Big]. \qquad (5)$$

This may be seen as follows. For $1 \le k \le n/2$, write $\mathbf{X}_k, \mathbf{Y}_k, \mathbf{Z}_k$ for events

$$\mathbf{X}_k := \{B_i \cap \mathbf{W}_k \neq \emptyset\}, \quad \mathbf{Y}_k := \{A_i \subseteq \mathbf{W}_k\}, \quad \mathbf{Z}_k := \{\textstyle\bigwedge_{j=1}^{i-1} \bigvee_{y \in A_j \setminus A_i} (y \notin \mathbf{W}_k)\}.$$

First, note that $\mathbf{Y}_k \wedge \bigwedge_{j=1}^{i-1}(A_j \not\subseteq \mathbf{W}_k)$ is equivalent to $\mathbf{Y}_k \wedge \mathbf{Z}_k$. Next, note that $(\mathbf{X}_k, \mathbf{Z}_k)$ is independent of $\mathbf{Y}_k$ (by the independence of events $\{x \in \mathbf{W}_k\}$ over $x \in \{0,1\}^n$ and the fact that $A_i \cap B_i = \emptyset$). Therefore, $\mathbb{P}[\mathbf{X}_k \,|\, \mathbf{Y}_k \wedge \mathbf{Z}_k] = \mathbb{P}[\mathbf{X}_k \,|\, \mathbf{Z}_k]$. We now use that fact that events $\mathbf{X}_k$ and $\mathbf{Z}_k$ are respectively *increasing* and *decreasing* in the lattice of subsets of $\{0,1\}^n$. Since $\mathbf{W}_k$ is a product distribution ($\mathbf{W}_k$ contains each element of $\{0,1\}^n$ independently with probability 2^{k-n}), it follows that events $\mathbf{X}_k$ and $\mathbf{Z}_k$ are negatively correlated, that is, $\mathbb{P}[\mathbf{X}_k \,|\, \mathbf{Z}_k] \le \mathbb{P}[\mathbf{X}_k]$.[4] Therefore, $\mathbb{P}[\mathbf{X}_k \,|\, \mathbf{Y}_k \wedge \mathbf{Z}_k] \le \mathbb{P}[\mathbf{X}_k]$ for all $1 \le k \le n/2$. It follows that $\mathbb{P}[\mathbf{X_k} \,|\, \mathbf{Y_k} \wedge \mathbf{Z_k}] \le \mathbb{P}[\mathbf{X_k}]$; note that this inequality is equivalent to (5).

Having shown (5), we next observe

$$\mathbb{P}[B_i \cap \mathbf{W} \neq \emptyset] \;\le\; \sum_{x \in B_i} \mathbb{P}[x \in \mathbf{W}] \;\le\; \frac{|B_i|}{2^{n/2}} \;=\; \frac{n^{O(1)}}{2^{n/2}} \;=\; 2^{-\Omega(n)}. \tag{6}$$

Stringing together inequalities (4), (5) and (6), we conclude that $\mathbb{P}[Q(\mathbf{W}) \neq Q^+(\mathbf{W})] = 2^{-\Omega(n)}$. □

Using this approximation of NP queries by monotone queries, we prove:

Theorem 3.6. (restated) *The NP query complexity of witness finding, $m(\mathsf{NP})$, is $\Omega(n^2)$.*

Proof. Let $m = m(\mathsf{NP})$. By Lemma 2.1, there exist NP queries $Q_1, \dots, Q_m$ and a function $f : \{\top, \bot\}^m \to \{0,1\}^n$ such that

$$\mathbb{P}[f(Q_1(\mathbf{W}), \dots, Q_m(\mathbf{W})) \in \mathbf{W} \mid \mathbf{W} \neq \emptyset] > 1/2.$$

Let $Q_1^+, \dots, Q_m^+$ be monotone queries approximating $Q_1, \dots, Q_m$ as in Lemma 3.2. We have

$$\begin{aligned} &\mathbb{P}[f(Q_1^+(\mathbf{W}), \dots, Q_m^+(\mathbf{W})) \in \mathbf{W}] \\ &\qquad \ge \mathbb{P}[f(Q_1(\mathbf{W}), \dots, Q_m(\mathbf{W})) \in \mathbf{W}] - \sum_{i=1}^{m} \mathbb{P}[Q_i(\mathbf{W}) \neq Q_i^+(\mathbf{W})] \\ &\qquad = \Omega(1) - \frac{m}{2^{\Omega(n)}}. \end{aligned}$$

On the other hand, by Lemma 2.5,

$$\mathbb{P}[f(Q_1^+(\mathbf{W}), \dots, Q_m^+(\mathbf{W})) \in \mathbf{W}] \le O(m/n^2) + o(1).$$

It follows that $\Omega(1) - m2^{-\Omega(n)} \le O(m/n^2) + o(1)$, which is only possible if $m = \Omega(n^2)$. □

4 Affine Witness Sets

At this point, we have shown that $m(\mathsf{Intersection})$, $m(\mathsf{Monotone})$ and $m(\mathsf{NP})$ are all $\Theta(n^2)$ by a combination of our lower bound (Theorems 1.3 and 1.6) and the upper bounds mentioned in §1. We now turn our attention to the setting of affine witness sets. We would like to prove lower bounds on $m_{\text{affine}}(\mathsf{Intersection})$, $m_{\text{affine}}(\mathsf{Monotone})$ and $m_{\text{affine}}(\mathsf{NP})$ using similar information-theoretic arguments. We begin by considering the natural affine analogue of the random witness set $\mathbf{W}$. For all $0 \le k \le n$, let $\mathbf{A}_k$ be the uniform random k-dimensional subspace of $\{0,1\}^n$. Let $\mathbf{k}$ be uniform in $\{1, \dots, n/2\}$ (as before) and let $\mathbf{A} := \mathbf{A_k}$.

Unfortunately, when we attempt to repeat the argument in §2, we get stuck at Lemma 2.2 (the Bollobás-Thomason Theorem). In particular, in order to have an appropriate version of Lemma 2.2(3) in the affine setting, we need a positive answer the following question:

Question 4.1. *Let Q be a non-trivial monotone increasing property of affine subspaces of $\{0,1\}^n$. For all $0 \le k \le n$, let $p_k := \mathbb{P}[\mathbf{A}_k$ has property $Q]$. Let θ be the unique index such that $p_\theta \le 1/2 < p_{\theta+1}$. Is it necessarily true that $\min\{p_k, 1-p_k\} \le 2^{-|\theta-k|+O(1)}$ for all k?*

[4]For any set I (in our case, $I = \{0,1\}^n$), if μ is a product distribution on $\{0,1\}^I$ (equivalently, the lattice of subsets of I) and $A, B \subseteq \{0,1\}^I$ such that A is monotone increasing and B is monotone decreasing, then $\mu(A|B) \le \mu(A)$. This is a special case of the FKG inequality (see Ch. 6 of [1]).

In other words, Question 4.1 asks whether every monotone property has an *exponentially sharp threshold* in the lattice of affine subspaces of $\{0,1\}^n$.

Remark 4.2. We can ask a similar question with respect to the lattice $\mathcal{L}_n$ of linear subspaces of $\{0,1\}^n$ (we suspect that the answer is the same). Writing $\mathcal{P}_n$ (resp. $\mathcal{P}_{2^n}$) for the lattice of subsets of $[n]$ (resp. $\{0,1\}^n$), note that $\mathcal{L}_n$ has an ambiguous status in relation to $\mathcal{P}_n$ and $\mathcal{P}_{2^n}$: on the one hand, $\mathcal{L}_n$ is the "q-analogue" of $\mathcal{P}_n$; on the other hand, $\mathcal{L}_n$ is a subset (in fact, a sub-meet-semilattice) of $\mathcal{P}_{2^n}$. Using a q-analogue of the Kruskal-Katona Theorem due to Chowdhury and Patkos [5], we can show that $p_k \leq 2^{-\Omega(\theta/k)}$ for all $k < \theta$ and $1 - p_k \leq 2^{-\Omega((n-\theta)/(n-k))}$ for all $k > \theta$. This shows that the threshold behavior of monotone properties in $\mathcal{L}_n$ scales at least like monotone properties in $\mathcal{P}_n$. The linear version of Question 4.1 asks whether the threshold behavior of monotone properties in $\mathcal{L}_n$ in fact scales like monotone properties in $\mathcal{P}_{2^n}$.

If the answer to Question 4.1 is "yes", then we get $m_{\text{affine}}(\mathsf{Monotone}) = \Omega(n^2)$ by using the same information-theoretic argument as in our proof of Theorem 1.3 in §2. While we were unable to answer Question 4.1 for general monotone queries, we give a positive answer in the special case where Q is an intersection query (Theorem 4.5). This, in turn, leads to Theorem 1.7, the lower bound $m_{\text{affine}}(\mathsf{Intersection}) = \Omega(n^2)$. The proof of Theorem 4.5 requires two technical lemmas, which we state next.

Lemma 4.3. *Let $S \subseteq \{0,1\}^n$ and let $\mathbf{H}$ be a uniform random affine hyperplane (i.e. $(n-1)$-dimensional subspace) in $\{0,1\}^n$. Then for all $\lambda > 0$,*

$$\mathbb{P}\left[|S \cap \mathbf{H}| \leq (\tfrac{1}{2} - \lambda)|S|\right] \leq \frac{1}{4\lambda^2 |S|}.$$

Proof. Let $\mathbf{Z} := |S \cap \mathbf{H}|$. We have $\mathbb{E}[\mathbf{Z}] = |S|/2$ and

$$\begin{aligned}\mathbb{E}[\mathbf{Z}^2] &= \sum_{x \in S} \mathbb{P}[x \in \mathbf{H}] + \sum_{x,y \in S\,:\,x \neq y} \mathbb{P}[x, y \in \mathbf{H}] \\ &= \frac{|S|}{2} + |S|(|S|-1)\frac{2^{n-1}-1}{2(2^n-1)} \;\leq\; \frac{1}{4}(|S| + |S|^2).\end{aligned}$$

By Chebyshev's inequality,

$$\mathbb{P}\left[\mathbf{Z} \leq (\tfrac{1}{2} - \lambda)|S|\right] \leq \mathbb{P}\left[|\mathbf{Z} - \mathbb{E}[\mathbf{Z}]| \geq \lambda |S|\right] \leq \frac{\mathsf{Var}(\mathbf{Z})}{\lambda^2 |S|^2} = \frac{\mathbb{E}[\mathbf{Z}^2] - \mathbb{E}[\mathbf{Z}]^2}{\lambda^2 |S|^2} \leq \frac{1}{4\lambda^2 |S|}.$$

This concludes the proof. □

Lemma 4.4. *Let $S \subseteq \{0,1\}^n$, let $\mathbf{B} = \mathbf{A}_{n-j}$ be a uniform random affine subspace of $\{0,1\}^n$ of co-dimension j, and let $b = 2^{-1/4}$. Then*

$$\mathbb{P}[\mathbf{B} \cap S = \emptyset] \leq \frac{2^{j+4(1+b+b^2+\cdots+b^j)}}{|S|}.$$

Proof. We argue by induction on j. In base case $j = 0$ (where $\mathbf{B} = \{0,1\}^n$), the lemma holds since $\mathbb{P}[\mathbf{B} \cap S = \emptyset] = 0$.

For induction step, let $j \geq 1$ and assume the lemma holds for $j-1$. By the induction hypothesis, for every affine hyperplane H,

$$\mathbb{P}[\mathbf{B} \cap S = \emptyset \mid \mathbf{B} \subseteq H] \leq \frac{2^{j-1+4(1+b+b^2+\cdots+b^{j-1})}}{|S \cap H|}.$$

Let $\mathbf{H}$ be a uniform random affine hyperplane. Note that $\mathbf{H}$ is independent of the event that $\mathbf{B} \subseteq \mathbf{H}$.

Let $\lambda := b^j/4$. We have

$$\begin{aligned}
\mathbb{P}[\mathbf{B} \cap S = \emptyset] &= \mathbb{P}[\mathbf{B} \cap S = \emptyset \mid \mathbf{B} \subseteq \mathbf{H}] \\
&\le \mathbb{P}\left[\mathbf{B} \cap S = \emptyset \text{ or } |S \cap \mathbf{H}| < (\tfrac{1}{2} - \lambda)|S| \;\middle|\; \mathbf{B} \subseteq \mathbf{H}\right] \\
&\le \mathbb{P}\left[|S \cap \mathbf{H}| < (\tfrac{1}{2} - \lambda)|S|\right] \\
&\quad + \mathbb{P}\left[\mathbf{B} \cap S = \emptyset \;\middle|\; \mathbf{B} \subseteq \mathbf{H} \text{ and } |S \cap \mathbf{H}| \ge (\tfrac{1}{2} - \lambda)|S|\right] \\
&\le \frac{1}{4\lambda^2 |S|} + \frac{2^{j-1+4(1+b+b^2+\cdots+b^{j-1})}}{(\frac{1}{2} - \lambda)|S|} \quad \text{(Lemma 4.3 and ind. hyp.)} \\
&= \left(2^{(j+4)/2} + \frac{2^{j+4(1+b+b^2+\cdots+b^{j-1})}}{1 - (b^j/2)}\right)\frac{1}{|S|}.
\end{aligned}$$

Noting that $1 - (b^j/2) \ge 2^{-b^j}$, we have

$$\begin{aligned}
2^{(j+4)/2} + \frac{2^{j+4(1+b+b^2+\cdots+b^{j-1})}}{1 - (b^j/2)} &\le 2^{(j+4)/2} + 2^{j+4(1+b+b^2+\cdots+b^{j-1})+b^j} \\
&\le 2^{j+4(1+b+b^2+\cdots+b^{j-1})+b^j}(1 + 2^{-(j+4)/2}) \\
&\le 2^{j+4(1+b+b^2+\cdots+b^{j-1})+b^j} e^{2^{-(j+4)/2}} \\
&\le 2^{j+4(1+b+b^2+\cdots+b^{j-1}+b^j)},
\end{aligned}$$

where the last inequality uses $e^{2^{-(j+4)/2}} < 2^{3b^j}$ (as we have $\log(e)2^{-(j+4)/2} < 2^{-j/2} < 2^{-j/4} = b^j < 3b^j$.) The proof is completed by combining the above inequalities. □

We now prove Theorem 4.5 which gives a positive answer to Question 4.1 in case of intersection queries.

Theorem 4.5. *Let $S \subseteq \{0,1\}^n$. For all $0 \le k \le n$, let $p_k := \mathbb{P}[\mathbf{A}_k \cap S \ne \emptyset]$. Let $\tau := n - \log |S|$. Then $\min\{p_k, 1 - p_k\} \le 2^{-|\tau - k| + O(1)}$ for all k.*

(Note that $|\theta - \tau| = O(1)$ for θ as in Question 4.1.)

Proof. The case where $k \le \tau$ follows from a simple union bound. Let $\mathbf{a}_1, \ldots, \mathbf{a}_{2^k}$ enumerate the elements of $\mathbf{A}_k$ in any order. Then

$$p_k \;=\; \mathbb{P}[\mathbf{A}_k \cap S \ne \emptyset] \;\le\; \sum_{i=1}^{2^k} \mathbb{P}[\mathbf{a}_i \in S] \;=\; \sum_{i=1}^{2^k} \frac{|S|}{2^n} \;=\; 2^{-(\tau - k)}.$$

In the case $k > \tau$, we invoke Lemma 4.4 (with $j = n - k$):

$$1 - p_k = \mathbb{P}[\mathbf{A}_k \cap S = \emptyset] \le \frac{2^{n-k+4(1+b+\cdots+b^{n-k})}}{|S|} \le 2^{\tau - k + 4\sum_{j=0}^{\infty} b^j} \le 2^{-(k-\tau)+26}.$$

Therefore, $\min\{p_k, 1 - p_k\} \le 2^{-|\tau - k| + O(1)}$ completing the proof. □

Our proof of Theorem 1.7 ($m_{\text{affine}}(\mathsf{Intersection}) = \Omega(n^2)$) uses the same argument as our proof of Theorem 1.3 ($m(\mathsf{Monotone}) = \Omega(n^2)$). Formally, we require the following analogues of Lemmas 2.3, 2.4 and 2.5, in which $\mathbf{W}$ is replaced by $\mathbf{A}$ and we restrict attention to the case where Q is an intersection query (rather than an arbitrary monotone query).

Lemma 4.6. *$\mathbb{H}(Q(\mathbf{A}) \mid \mathbf{k}) = O(1/n)$ for every intersection query Q.*

Proof. (The following argument is identical to Lemma 2.3, with Theorem 4.5 crucially playing the role of Lemma 2.2.) If Q is identically $\bot$ or $\top$, then the statement is trivial (as $\mathbb{H}(Q(\mathbf{A}) \mid \mathbf{k}) = 0$). So assume Q is a non-trivial intersection query. Let $p_k := \mathbb{P}[Q(\mathbf{A}_k) = \top]$ for all $k \in \{0, \dots, n\}$, and let θ be the unique threshold such that $p_\theta \leq 1/2 < p_{\theta+1}$. By Theorem 4.5, we have

$$\begin{aligned} \mathbb{H}(Q(\mathbf{A}) \mid \mathbf{k}) &= \sum_{k=0}^{n/2} \mathbb{P}[\mathbf{k} = k] \cdot \mathbb{H}(Q(\mathbf{A}_k)) \\ &= \frac{2}{n} \sum_{k=1}^{n/2} H(p_k) \ \leq \ \frac{2}{n} \sum_{k=1}^{n/2} \frac{|\theta - k|}{2^{|\theta-k|-26}} = O(1/n). \end{aligned}$$ □

Lemma 4.7. *For every random variable* $\mathbf{z}$ *on* $\{0,1\}^n$ *(not necessarily independent of* $\mathbf{A}$*),*

$$\mathbb{P}[\mathbf{z} \in \mathbf{A}] \leq \frac{4}{n}\mathbb{H}(\mathbf{z}) + \frac{1}{2^{n/4}}.$$

Proof. The only property of $\mathbf{W}$ used in the proof of Lemma 2.4 is that $\mathbb{P}[x \in \mathbf{W}] < 2^{-n/2}$ for every $x \in \{0,1\}^n$. This property also holds of $\mathbf{A}$ (since $\mathbf{A}$ is a random affine set of dimension $\mathbf{k}$ and $\mathbb{P}[\mathbf{k} \leq n/2] = 1$). □

Lemma 4.8. *For all intersection queries* $Q_1, \dots, Q_m$ *and every function* $f : \{\top, \bot\}^m \to \{0,1\}^n$*,*

$$\mathbb{P}[f(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A})) \in \mathbf{A}] \leq O(m/n^2) + o(1).$$

Proof. (The following argument is identical to the proof of Lemma 2.5, with Lemmas 4.6 and 4.7 playing the role of Lemmas 2.3 and 2.4.) By standard entropy inequalities,

$$\begin{aligned} \mathbb{H}(f(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A}))) &\leq \mathbb{H}(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A})) \\ &\leq \mathbb{H}(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A}), \mathbf{k}) \\ &= \mathbb{H}(\mathbf{A}) + \mathbb{H}(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A}) \mid \mathbf{k}) \\ &\leq \mathbb{H}(\mathbf{k}) + \mathbb{H}(Q_1(\mathbf{A}) \mid \mathbf{k}) + \cdots + \mathbb{H}(Q_m(\mathbf{A}) \mid \mathbf{k}). \end{aligned}$$

Since $\mathbb{H}(\mathbf{A}) = \log(n/2)$ and $\mathbb{H}(Q_i(\mathbf{A}) \mid \mathbf{k}) = O(1/n)$ for all i by Lemma 4.6, we have

$$\mathbb{H}(f(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A}))) \leq O(m/n) + \log n.$$

Since $f(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A}))$ is a random variable on $\{0,1\}^n$, we can apply Lemma 4.7 to get

$$\begin{aligned} \mathbb{P}[f(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A})) \in \mathbf{A}] &\leq \frac{4}{n}\mathbb{H}(f(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A}))) + \frac{1}{2^{n/4}} \\ &\leq O(m/n^2) + \frac{4 \log n}{n} + \frac{1}{2^{n/4}} \\ &= O(m/n^2) + o(1). \end{aligned}$$ □

Finally, we prove our lower bound on the intersection query complexity of affine witness finding.

Theorem 4.7. (restated) *The intersection query complexity of affine witness finding,* $m_{\text{affine}}(\mathsf{Intersection})$*, is* $\Omega(n^2)$*.*

Proof. (The following argument is identical to the proof of Theorem 1.3, with Lemma 4.8 playing the role of Lemma 2.5.) Let $m = m_{\text{affine}}(\mathsf{Intersection})$. By Lemma 2.1, there exist intersection queries $Q_1, \dots, Q_m$ and a function $f : \{\top, \bot\}^m \to \{0,1\}^n$ such that

$$\mathbb{P}[f(Q_1(\mathbf{A}), \dots, Q_m(\mathbf{A})) \in \mathbf{A}] > 1/2.$$

(Note that $\mathbf{A}$ is never empty (in contrast to $\mathbf{W}$), so there is no need to condition on the event $\{\mathbf{A} \neq \emptyset\}$.) By Lemma 4.8,

$$\begin{aligned}\mathbb{P}[f(Q_1(\mathbf{A}),\dots,Q_m(\mathbf{A})) \in \mathbf{A}] &= \mathbb{P}[f(Q_1(\mathbf{A}),\dots,Q_m(\mathbf{A})) \in \mathbf{A}] \\ &\le O(m/n^2) + o(1).\end{aligned}$$

It follows that $1/2 < O(m/n^2) + o(1)$ and hence $m = \Omega(n^2)$. □

5 Conclusion

We initiated the study of the information-theoretic witness finding problem. For three natural classes of queries (intersection queries, monotone queries, NP queries), we proved lower bounds of $\Omega(n^2)$ on the query complexity of witness finding over arbitrary subsets of $\{0,1\}^n$. These lower bounds match upper bounds coming from classic results of Valiant and Vazirani [8] and Ben-David et al. [3]. In addition, we considered the setting where witness sets are affine subspaces of $\{0,1\}^n$ and proved a tight lower bound of $\Omega(n^2)$ for intersection queries. (All of our lower bounds hold even under the strong interpretation of Ω, i.e., for all but finitely many n.) Our investigation of affine witness finding led to an interesting and apparently new question about the threshold behavior of monotone properties in the affine lattice (Question 4.1). Other questions left open by this work are to resolve the monotone and NP query complexity of affine witness finding (i.e. $m_{\text{affine}}(\mathsf{Monotone})$ and $m_{\text{affine}}(\mathsf{NP})$). Finally, we wonder whether the idea in §3 of approximating NP queries by monotone queries might have other applications in complexity theory.

Acknowledgements

We thank Oded Goldreich for feedback on an earlier manuscript. We are also grateful to the anonymous reviewers for their detailed and extremely helpful comments.

A Proof of Lemma 2.1

In order to apply Yao's minimax principle [9], we express $m(\mathcal{W},\mathcal{Q})$ in terms of a particular matrix M. Let $\mathcal{F}$ be the set of functions $\{\top,\bot\}^m \to \{0,1\}^n$. Let $\mathcal{A} := \mathcal{Q}^m \times \mathcal{F}$ (representing the set of deterministic witness finding algorithms). Let $\mathcal{W}_0 := \mathcal{W} \setminus \{\emptyset\}$. Finally, let M be the $\mathcal{A} \times \mathcal{W}_0$-matrix defined by

$$M_{(Q_1,\dots,Q_m;f),W} := \begin{cases} 1 & \text{if } f(Q_1(W),\dots,Q_m(W)) \in W, \\ 0 & \text{otherwise.}\end{cases}$$

In this context, Yao's minimax principle states that for all random variables $\mathbf{W}$ on $\mathcal{W}_0$ and $(\mathbf{Q}_1,\dots,\mathbf{Q}_m;\mathbf{f})$ on $\mathcal{A}$,

$$\min_{(Q_1,\dots,Q_m;f)\in\mathcal{A}} \mathbb{E}[M_{(Q_1,\dots,Q_m;f),\mathbf{W}}] \le \max_{W\in\mathcal{W}_0} \mathbb{E}[M_{(\mathbf{Q}_1,\dots,\mathbf{Q}_m;\mathbf{f}),W}].$$

It follows that, if $\mathbb{P}[f(Q_1(\mathbf{W}),\dots,Q_m(\mathbf{W})) \in \mathbf{W}] \le 1/2$ for all $Q_1,\dots,Q_m \in \mathcal{Q}$ and every function $f : \{\top,\bot\}^m \to \{0,1\}^n$, then for all $(\mathbf{Q}_1,\dots,\mathbf{Q}_m;\mathbf{f}) \in \mathcal{A}$ (including the special case where $\mathbf{f}$ is deterministic, as in the definition of witness finding procedures), there exists $W \in \mathcal{W}_0$ such that $\mathbb{P}[\mathbf{f}(\mathbf{Q}_1(W),\dots,\mathbf{Q}_m(W)) \in W] \le 1/2$. Therefore, the $\mathcal{Q}$-query complexity of $\mathcal{W}$-witness finding is greater than m.

B Proof of Lemma 2.2

For inequality (1), let $\mathbf{Y}_1,\dots,\mathbf{Y}_{2^i}$ be independent copies of $\mathbf{W}_{\theta-i}$. For each $x \in \{0,1\}^n$, we have

$$\mathbb{P}[x \in (\mathbf{Y}_1 \cup \dots \cup \mathbf{Y}_{2^i})] = 1 - (1 - 2^{\theta-i-n})^{2^i} < 2^{\theta-n} = \mathbb{P}[x \in \mathbf{W}_\theta].$$

Moreover, events $\{x \in \mathbf{Y}_1 \cup \cdots \cup \mathbf{Y}_{2^i}\}$ are independent over $x \in \{0,1\}^n$. Thus, $\mathbf{Y}_1 \cup \cdots \cup \mathbf{Y}_{2^i}$ and $\mathbf{W}_\theta$ are both product distributions on the lattice of subsets of $\{0,1\}^n$ with biases $p = 1-(1-2^{\theta-i-n})^{2^i}$ and $q = 2^{\theta_n}$ respectively. Since $p < q$, $\mathbf{Y}_1 \cup \cdots \cup \mathbf{Y}_{2^i}$ is stochastically dominated by $\mathbf{W}_\theta$. (That is, $\mathbb{E}[f(\mathbf{Y}_1 \cup \cdots \cup \mathbf{Y}_{2^i})] \leq \mathbb{E}[f(\mathbf{W}_\theta)]$ for every monotone increasing function f of subsets of $\{0,1\}^n$.) Since Q is a monotone query, we have

$$\mathbb{P}[Q(\mathbf{Y}_1) \vee \cdots \vee Q(\mathbf{Y}_{2^i})] \leq \mathbb{P}[Q(\mathbf{Y}_1 \cup \cdots \cup \mathbf{Y}_{2^i})] \leq \mathbb{P}[Q(\mathbf{W}_\theta)].$$

By independence of $\mathbf{Y}_1, \ldots, \mathbf{Y}_{2^i}$, it follows that

$$1/2 \geq \mathbb{P}[Q(\mathbf{W}_\theta)] \geq \mathbb{P}[\textstyle\bigvee_{j=1}^{2^i} Q(\mathbf{Y}_j)] = 1 - \mathbb{P}[\neg Q(\mathbf{W}_{\theta-i})]^{2^i} = 1-(1-p_{\theta-i})^{2^i}.$$

Therefore, $p_{\theta-i} \leq 1-(1/2)^{1/2^i} < (\ln 2)/2^i$.

For inequality (2), let $\mathbf{Z}_1, \ldots, \mathbf{Z}_{2^i}$ be independent copies of $\mathbf{W}_{\theta+1}$. By a similar argument, we have

$$p_{\theta+i+1} = \mathbb{P}[Q(\mathbf{W}_{\theta+i+1})] \geq \mathbb{P}[\textstyle\bigvee_{j=1}^{2^i} Q(\mathbf{Z}_j)] = 1 - \mathbb{P}[\neg Q(\mathbf{W}_{\theta+1})]^{2^i} > 1 - \dfrac{1}{2^{2^i}}.$$

Finally, for inequality (3), note that for all $p, q \in [0,1]$,

$$0 \leq \min(p, 1-p) \leq q \leq 1/2 \implies H(p) \leq H(q) \leq 2q\log(1/q).$$

By this observation, together with (1) and (2), we have

$$H(p_{\theta-i-1}) \leq 2\frac{\ln 2}{2^{i+1}}\log(\frac{2^{i+1}}{\ln 2}) < \frac{i+2}{2^i}, \quad H(p_{\theta+i+1}) \leq 2\frac{1}{2^{2^i}}\log(2^{2^i}) = \frac{1}{2^{2^i-i-1}}.$$

From these two inequalities, it follows that $H(p_k) \leq (|\theta-k|+1)/2^{|\theta-k|-1}$.

References

[1] N. Alon and J. Spencer, *The Probablistic Method* (3rd edition), Wiley 2008.

[2] M. Bellare and S. Goldwasser, The complexity of decision versus search, *SIAM Journal on Computing*, 23:97–119, 1994.

[3] S. Ben-David, B. Chor, O. Goldreich, and M. Luby, On the theory of average-case complexity, *Journal of Computer and System Sciences*, 44(2):193–219, 1992.

[4] B. Bollobás and A.G. Thomason, Threshold functions, *Combinatorica*, 7(1):35–38, 1987.

[5] A. Chowdhury and B. Patkos, Shadows and intersections in vector spaces, *J. of Combinatorial Theory*, Ser. A 117, 1095–1106, 2010.

[6] T. Cover and J. Thomas, *Elements of Information Theory*, Wiley-Interscience New York, NY, 1991.

[7] H. Dell, V. Kabanets, D. van Melkebeek, and O. Watanabe, Is the Valiant-Vazirani isolation lemma improvable?, in *Proc. 27th Conference on Computational Complexity*, 10–20, 2012.

[8] L. Valiant and V. Vazirani, NP is as easy as detecting unique solutions, *Theoretical Computer Science*, 47:85–93, 1986.

[9] A.C. Yao, Probabilistic computations: toward a unified measure of complexity, *Proc. of the 18th IEEE Sympos. on Foundations of Comput. Sci.*, IEEE, 222–227, 1977.