

Active classification with comparison queries

Daniel M. Kane* Shachar Lovett† Shay Moran‡ Jiapeng Zhang§

April 19, 2017

Abstract

We study an extension of active learning in which the learning algorithm may ask the annotator to compare the distances of two examples from the boundary of their label-class. For example, in a recommendation system application (say for restaurants), the annotator may be asked whether she liked or disliked a specific restaurant (a label query); or which one of two restaurants did she like more (a comparison query).

We focus on the class of half spaces, and show that under natural assumptions, such as large margin or bounded bit-description of the input examples, it is possible to reveal all the labels of a sample of size n using approximately $O(\log n)$ queries. This implies an exponential improvement over classical active learning, where only label queries are allowed. We complement these results by showing that if any of these assumptions is removed then, in the worst case, $\Omega(n)$ queries are required.

Our results follow from a new general framework of active learning with additional queries. We identify a combinatorial dimension, called the *inference dimension*, that captures the query complexity when each additional query is determined by $O(1)$ examples (such as comparison queries, each of which is determined by the two compared examples). Our results for half spaces follow by bounding the inference dimension in the cases discussed above.

*Department of Computer Science and Engineering/Department of Mathematics, University of California, San Diego. dakane@ucsd.edu Supported by NSF Career Award ID 1553288.

†Department of Computer Science and Engineering, University of California, San Diego. slovett@cs.ucsd.edu. Research supported by NSF CAREER award 1350481.

‡Department of Computer Science and Engineering, University of California, San Diego, Simons Institute for the Theory of Computing, Berkeley, and Max Planck Institute for Informatics, Saarbrücken, Germany. shaymoran1@gmail.com.

§Department of Computer Science and Engineering, University of California, San Diego. jpeng.zhang@gmail.com. Research supported by NSF CAREER award 1350481.

1 Introduction

A central goal of *interactive learning* is understanding what type of interaction between a learner and a domain expert enhances the learning process, compared to the classical *passive learning* from labeled examples.

A basic model that was studied in this context is *pool-based active learning* (McCallum and Nigam, 1998). Here, the algorithm has an access to a large pool of *unlabeled examples* from which it can pick examples and query their labels. The goal is to make as few queries as possible while achieving generalization-guarantees which are comparable with these of a passive algorithm with an access to all of the labels.

A canonical example that demonstrates an advantage of active learning is the class of threshold functions¹ over the real line. Indeed, let c denote the learned threshold function, and let $x_1 < x_2 < \dots < x_n$ in \mathbb{R} be the given pool of unlabeled examples. It is possible to infer the labels of all n points by making at most $\log n + 2$ queries: query the labels of the extreme points $c(x_1), c(x_n)$; if $c(x_1) = c(x_n)$ then the remaining points must be labeled the same; otherwise, continue in a binary search fashion, by labeling the middle point of the interval whose extreme points have opposite labels. After at most $\log n$ such queries, the labels of all n points are revealed.

Unfortunately, this exponential improvement in the query complexity breaks for more general concept classes. In fact, even for the class of 2 dimensional threshold functions², namely the class of half-planes, the (worst-case) query complexity of active learning equals that of passive learning (see e.g. Dasgupta (2004)). Consequently, much of the literature was dedicated to developing theory that takes into consideration further properties of the unknown underlying distribution or the target concept (Balcan et al., 2007; Hanneke, 2007; Dasgupta et al., 2008; Balcan et al., 2009; Beygelzimer et al., 2010; Balcan et al., 2010; Koltchinskii, 2010; Balcan and Hanneke, 2011; Hanneke and Yang, 2012; El-Yaniv and Wiener, 2012; Balcan and Long, 2013; Gonen et al., 2013; Urner et al., 2013; Zhang and Chaudhuri, 2014; Wiener et al., 2015; Berlind and Urner, 2015).

We consider another approach by allowing the learning algorithm to further interact with the domain expert by asking additional queries. This poses a question:

Which additional queries can the algorithm use?

Allowing arbitrary queries will result in a very strong learner (indeed, by halving the set of potential hypothesis in every query, the number of queries can be made logarithmic). However, arbitrary queries are useless in practice: as experimental work by Lang and Baum (1992) shows, algorithms that use set of queries that are too rich may result in a poor practical performance. This is not surprising if we keep in mind that the human annotator who answers the queries is restricted (computationally and in other ways). Thus, a crucial factor in choosing the additional queries is compatibility with the annotator who answers them. A popular type of queries, which is used in applications involving human annotators, is *relative queries*. These queries poll relative information between two or more data points.

This work focuses on a basic kind of relative queries — *comparison queries*. Using such queries is sensible in settings in which there is some natural ordering of the instances with respect to the learned concept. As a toy example, consider the goal of classifying films according to whether a certain individual is likely to enjoy them or not (e.g. for recommending new films for this person). In this context, the input sample consists of films watched by the individual, a label query asks whether the person liked a film, and a comparison query asks which of two given films did the individual prefer. As we will see, comparison queries may significantly help.

¹ These are " $\mathbb{R} \rightarrow \{\pm 1\}$ " functions of the form $c(x) = \text{sign}(a \cdot x - b)$, where $a, b \in \mathbb{R}$.

² These are " $\mathbb{R}^2 \rightarrow \{\pm 1\}$ " functions of the form $c(x) = \text{sign}(\langle a, x \rangle - b)$, where $a \in \mathbb{R}^2, b \in \mathbb{R}$.

Another aspect implied by the restricted nature of the human annotator is that the learned concept resides in the class of concepts that can be computed by the annotator, which presumably has low capacity. Thus, realizability assumptions about the generating data distributions are plausible in this context.

1.1 Active learning with additional comparison queries

Consider a learned concept of the form $c(x) = \text{sign}(f(x))$, where f is a real valued function (e.g. half spaces, neural nets), and consider two instances x_1, x_2 such that, say $f(x_1) = 10$, and $f(x_2) = 1000$. Both $c(x_1)$ and $c(x_2)$ equal $+1$, however that $f(x_2) \gg f(x_1)$ suggests that x_2 is a “more positive instance” than x_1 . In the setting of film classification this is naturally interpreted as that the person likes the film x_2 more than the film x_1 . We call the query “ $f(x_2) \geq f(x_1)$?” a *comparison query*.

1.1.1 Example: learning half-planes with comparison and label queries

To be concrete, consider the class of half-planes in \mathbb{R}^2 . Here, a comparison query is equivalent to asking which of two sample points x_1, x_2 lies closer to the boundary line. Do such queries improve the query complexity over standard active learning? It is known that without such queries, in the worst-case, the learner essentially has to query all labels (Dasgupta, 2004).

The following algorithm demonstrates an exponential improvement when comparison queries are allowed. Later, we will present more general results showing that this can be generalized to higher dimensions under some natural restrictions, and that such restrictions are indeed necessary.

Interactive learning algorithm with comparison queries for half planes in \mathbb{R}^2

(see Figure 1 for a graphical illustration)

Setting: an unlabeled input sample of n points $x_1, \dots, x_n \in \mathbb{R}^2$ with hidden labels according to a half-plane $c(x) = \text{sign}(f(x))$, where $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is affine.

Repeat until all points are labeled:

1. Sample uniformly a subsample S of 30 points.
2. Query the labels of the points in S , and denote by:
 - \mathcal{P} — points in S labeled by $+1$,
 - \mathcal{N} — points in S labeled by -1 .
3. Use comparison queries to find:
 - (i) q — the closest point in \mathcal{P} to the boundary line,
 - (ii) v — the closest point in \mathcal{N} to the boundary line.
4. Denote by:
 - $[p, q], [q, r]$ — the two edges of the convex hull of \mathcal{P} that are incident to q ,
 - $[u, v], [v, w]$ — the two edges of the convex hull of \mathcal{N} that are incident to v .
5. Infer that:
 - all points inside the cone $\angle pqr$ are labeled $+1$,
 - all points inside the cone $\angle uvw$ are labeled -1 .
6. Repeat on the remaining unlabeled points.

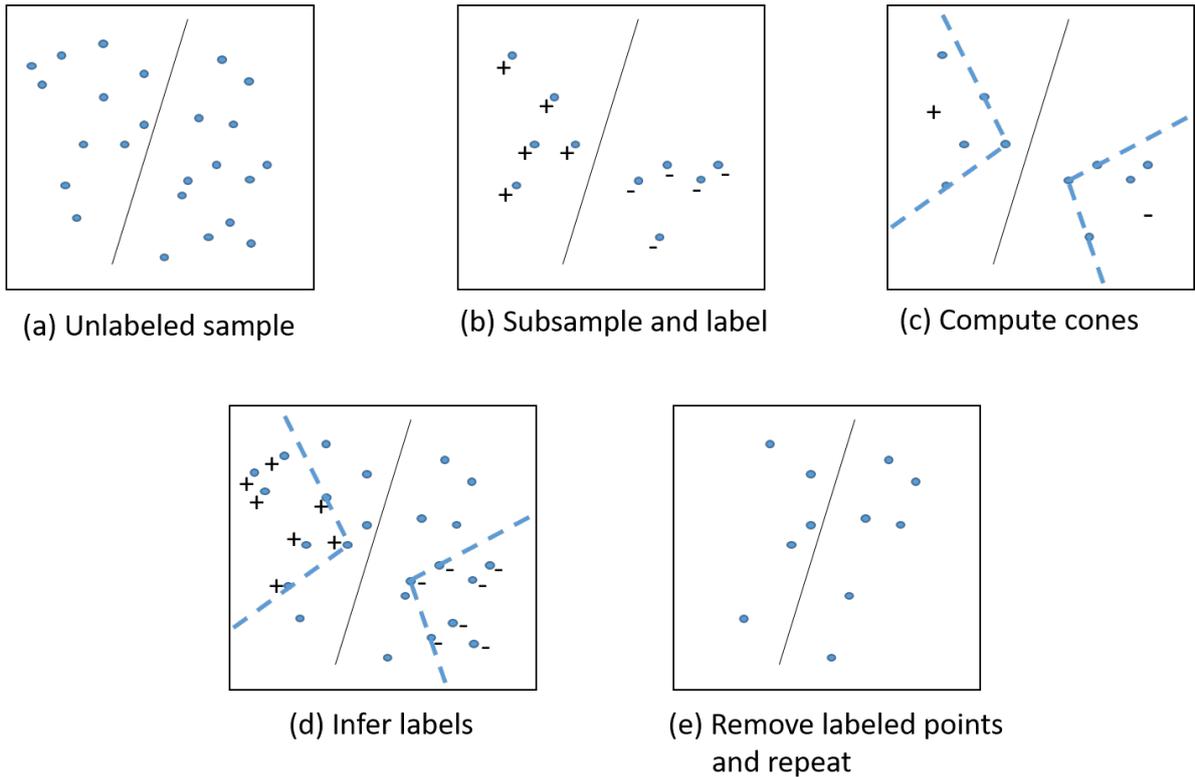


Figure 1: An illustration of a single iteration.

The algorithm proceeds by iterations: it repeats steps 1-5 until all points are labeled. At each iteration at most 60 queries are performed: 30 label queries in step 2 and at most 30 comparison queries in step 3, for finding the points q, v of minimal distance. In each iteration, the algorithm infers the labels of all points in the union of the angles $\angle pqr, \angle uvw$. We will refer to this region as “the confident region”.

We claim that after an expected number of $O(\log n)$ iterations (and therefore using only $O(\log n)$ queries) the algorithm infers the labels of all input points. Establishing this statement boils down, via a boosting argument (Theorem 3.2) to the following two important properties, which are easy to verify:

- **Confidence:** every point which is labeled by the algorithm is labeled correctly.
- **Inference from a small subsample:** the confident region is determined by a subsample of 6 labeled points of S (i.e. p, q, r and u, v, w).

These properties imply that at each iteration, with probability at least $1/2$, half of the remaining unlabeled points are labeled correctly (see Lemma 3.3). Thus, the expected number of queries is $O(\log n)$.

Paper organization. In Section 1.2 we present and discuss our results, later, in Section 1.3 we survey previous related works, and in Section 1.4 we ceil the introduction with suggested directions of future research and open problems. Sections 2, 3, and 4 contain the technical definitions and proofs.

1.2 Results

We next present and discuss our main results.

1. Subsection 1.2.1 is dedicated to our results concerning half spaces,
2. Subsection 1.2.2 is dedicated to the inference dimension, and how it captures the query complexity of active learning with additional comparison queries, and
3. Subsection 1.2.3 focuses on the general framework of active learning with additional queries.

1.2.1 Interactive learning of half spaces with comparison queries

We start by discussing our results for interactive learning of half spaces in \mathbb{R}^d when both label queries and comparison queries are allowed. We show that a general algorithm, as the one we described for \mathbb{R}^2 , cannot exist for $d \geq 3$. However, we identify two useful properties that allow for such a learning algorithm: bounded bit complexity and margin.

Exact recovery of labels. We first describe our results in the context of exact recovery. Here, the labels of all n sample points needs to be revealed, using as few queries as possible.

We first show that in the worst case, in \mathbb{R}^d this requires $\Omega(n)$ queries for any $d \geq 3$ (recall that $O(\log n)$ queries suffice in \mathbb{R}^2).

Theorem 1.1 (Theorem 4.11, informal version). *There are n points in \mathbb{R}^3 that require $\Omega(n)$ label and comparison queries for revealing all labels.*

Our first positive result shows that efficient exact recovery of labels is possible if the points have low bit complexity.

Theorem 1.2 (Theorem 4.1, informal version). *Consider an arbitrary realizable sample of n points in \mathbb{R}^d whose individual bit complexity is B . The labels of all sample points can be learned using $\tilde{O}(B \log n)$ label and comparison queries.*

As an example, consider the sample consisting of the 2^N point in the boolean hypercube $\{0, 1\}^N$. The bit complexity of every point is N . The above theorem implies that given an unknown threshold function on $\{0, 1\}^N$, it is possible to reveal all 2^N labels using $\tilde{O}(N^2)$ comparison and label queries. This should be compared to the situation where only label queries are allowed, where all 2^N queries are necessary. We discuss this example in more detail in Section 1.4.4, and leave as an open question whether revealing the 2^N labels can be done efficiently (the above $\tilde{O}(N^2)$ bound applies only to the query complexity, and not to the total running time).

Our second positive result shows that a similar algorithm exists if the margin is large.

Theorem 1.3 (Theorem 4.7, informal version). *Assume a sample of n points in \mathbb{R}^d with maximal ℓ_2 norm ρ and margin at least γ . The labels of all points can be recovered using $\tilde{O}(d \log(\rho/\gamma) \log n)$ label and comparison queries.*

Our bound is in fact stronger: it is

$$\tilde{O}(d \log(1/\eta) \log n),$$

where η is the *minimal-ratio* of the input sample, defined by $\frac{|\min_i f(x_i)|}{|\max_i f(x_i)|}$, where x_1, \dots, x_n are the sample points and $\text{sign}(f(x))$ is the learned concept. In Section 4.1.2 we show that the maximal ratio is lower bounded by the margin (and therefore yields a stronger statement).

Note that the above bound depends on ρ/γ logarithmically, and therefore applies even in settings when the margin is exponentially small. Similar dependence of on ρ/γ is obtained by the Ellipsoid method (Karmarkar, 1984), and Vaidya Cutting Plane method (Vaidya, 1996),

that use $O(d^2 \log(\rho/\gamma))$ and $O(d \log(d\rho/\gamma))$ iterations respectively, when used to find a linear classifier that is consistent with a realizable sample with margin ρ/γ in \mathbb{R}^d .

The upper bound in the above theorem depends on the dimension, which is often avoided in bounds that depend on the margin. It is therefore natural to ask whether there is a bound that depends only on the margin. We show that it is impossible:

Theorem 1.4 (Theorem 4.17, informal version). *There is a sample of n unit vectors in \mathbb{R}^{n+1} with $\Omega(1)$ margin that require $\Omega(n)$ label and comparison queries to recover all the labels.*

Statistical learning. Using standard arguments, we translate the results above to the statistical setting and get bounds on the sample and query complexity: the algorithmic results above directly give a learning algorithm for realizable distributions with sample complexity $n(\epsilon, \delta) = O\left(\frac{d + \log(1/\delta)}{\epsilon}\right)$ and query complexity approximately logarithmic in $n(\epsilon, \delta)$, where ϵ is the error and $1 - \delta$ is the confidence of the algorithm. (See Section 4.1.1 for the bit complexity and Section 4.1.2 for the margin.) Our lower bounds translate analogously to realizable distributions that require $\Omega(1/\epsilon)$ queries for achieving error ϵ and constant confidence (say $1 - \delta = 5/6$). See Section 4.2 for details.

1.2.2 Inference dimension

Our results for learning half spaces rely on a common combinatorial property that we describe next. Let X be a set and H a concept class where each concept is of the form $\text{sign}(f(x))$ for $f : X \rightarrow \mathbb{R}$. For example, X may be a finite set $X = \{x_1, \dots, x_n\} \subseteq \mathbb{R}^d$ and H the class of all half spaces with margin at least $1/100$ with respect to X ; or $X = \{0, 1\}^d$ and H is all half spaces; or $X = \mathbb{R}^d$ and H a class of (signs of) low degree polynomials; etcetera.

Let $S \subseteq X$ be an unlabeled sample. An S -query is either a label query regarding some $x \in S$, or a comparison query regarding $x_1, x_2 \in S$. Namely, the allowed queries are “ $f(x) \geq 0$?” (label query) or “ $f(x_1) \geq f(x_2)$ ” (comparison query). For $x \in X$ and $c \in H$, let

$$S \xrightarrow{f} x$$

denote the statement that the comparison and label queries on S determine the label of x , when the learned concept is $c = \text{sign}(f(x))$.

The *inference-dimension* of (X, H) is the smallest number k such that for every $c = \text{sign}(f(x)) \in H$, and every $S \subset X$ of size at least k , there exists $x \in S$ such that

$$S \setminus \{x\} \xrightarrow{f} x.$$

In other words, if the inference dimension is k then in every sample of size k or more, there is a point whose label can be inferred from the label and comparison queries on the remaining points.

For example the inference dimension of (X, H) , where $X = \mathbb{R}$, and H is the class of threshold functions is 3: indeed, to see it is at most 3, note that if all 3 points have the same label, then the label of the midpoint can be inferred from the other two labels, and if not all points have the same label then the midpoint and, say the right point have the same label, and so the label of the right point can be inferred from the other 2 labels. In this example comparison queries are not required. Another example, which requires comparison queries, is where $X = \mathbb{R}^2$ and H is the class of half-planes. Here the inference dimension³ is at most 7: indeed, in any sample of at least 7 points there are 4 points with the same label, and the label of one of these points can be inferred from the other 3 (see Figure 1 and Section 1.1.1).

³One can show the inference dimension here is 5.

The next Theorem shows that inference dimension captures the query-complexity in active learning with comparison queries. It is worth noting that in the classical setting of active learning, when only label queries are allowed, the inference dimension specializes to the *star dimension* (Hanneke and Yang, 2015), which similarly captures the (worst-case) query complexity in this setting.

Theorem 1.5. *Let k denote the inference dimension of (X, H) . Then:*

1. *There is an algorithm that reveals the labels of any realizable sample of size n using at most $O(k \log k \log n)$ queries.*
2. *Any algorithm that reveals the labels of any realizable sample of size k must use $\Omega(k)$ queries in the worst-case.*

The upper bound (the first item) is a corollary of Theorem 3.2, and the lower bound is a corollary of Theorem 3.5. Both Theorems are discussed in the next subsection. While the lower bound is relatively straight forward, our derivation of the upper bound requires several steps, which we summarize next.

- (i) **Low inference dimension \implies weak confident learner:** we first show that if the inference dimension is at most k then there is a *weak confident learner* for (X, H) with query complexity $O(k \log k)$ (see Lemma 3.3). A confident learner is a learning algorithm that may abstain from predicting on some points $x \in X$, but must be correct on every point where it does not abstain. A weak confident learner is a confident learner that with constant probability does not abstain on a constant fraction of X (see Section 3.1 for a formal definition).
- (ii) **Boosting the weak confident learner:** once a weak confident learner is derived, we transform it into the desired learning algorithm using a simple boosting argument.

While the boosting part is rather standard, showing that low inference dimension implies a weak confident learner relies on a symmetrization argument. This symmetrization argument can be replaced by a more standard sample compression argument, however this would result in a suboptimal query complexity bound.

We get the following corollary of Theorem 1.5 in the statistical setting:

Corollary 1.6. *Let k denote the inference dimension of (X, H) .*

1. *Let $n(\epsilon, \delta)$ denote the passive sample complexity of learning (X, H) with error ϵ and confidence $1 - \delta$. There is an algorithm that learns (X, H) with sample complexity $n(\epsilon, \delta)$ and query complexity $O(k \log k \log(n(\epsilon, \delta)))$.*
2. *Any algorithm that learns (X, H) with error at most $\epsilon = 1/k$ and confidence at least $5/6$ must use at least $\Omega(1/\epsilon)$ queries for some realizable distribution.*

This Corollary follows from Corollaries 1.8 and 3.6.

A dichotomy. The passive sample complexity of (X, H) is $n(\epsilon, \delta) = \Theta(\frac{d + \log(1/\delta)}{\epsilon})$, where d is the VC-dimension of H (Hanneke, 2015). Thus the above corollary implies a dichotomy: if the inference dimension is finite then the query complexity is logarithmic in $1/\epsilon$, and if it is infinite then the query complexity is $\Omega(1/\epsilon)$.

The results presented in the previous section regarding learning half spaces with comparison and label queries are derived by analyzing the inference-dimension of the relevant classes, which we sketch next.

Sketch of upper bounding the inference dimension of half spaces. Our upper bounds in terms of margin and bit-complexity follow a similar outline. We next give a rough sketch of the arguments in order to convey their flavor.

Consider the case where every instance $x \in X$ has bounded bit-complexity, say $X \subset [N]^d$ for some bounded $N \in \mathbb{N}$. We wish to show that for a sufficiently large $Y \subseteq X$, and every half space $c = \text{sign}(f)$ there is some $x \in Y$ such that

$$Y \setminus \{x\} \xrightarrow{f} x.$$

By removing at most half of the elements of Y , we may assume that c is constant on Y , without loss of generality assume that $c(x) = +1$, for every $x \in Y$. Let x_1, x_2, \dots be an ordering of the elements of Y such that

$$f(x_1) \leq f(x_2) \leq \dots$$

The first observation is that it suffices to show that there exists i_0 such that $x_{i_0} - x_1$ is a nonnegative linear combination of the $x_i - x_j$, where $i > j$ (see Claim 4.5).

The existence of such an i_0 is achieved by a pigeon hole argument showing that if Y is sufficiently large then there are two distinct linear combinations of the $x_{i+1} - x_i$'s with coefficients from $\{0, 1\}$ that yield the same vector:

$$\sum_{i=1}^{k/2} \beta_i (x_{i+1} - x_i) = \sum_{i=1}^{k/2} \gamma_i (x_{i+1} - x_i),$$

Then a short calculation yields that the maximal j such that $\beta_j \neq \gamma_j$ serves as the desired i_0 .

The upper bound in terms of margin is technically more involved. Instead of nonnegative linear combinations, one can consider linear combinations such that every coefficient is at least $-\gamma$, where γ is the margin, and the pigeon hole argument is replaced by a volume argument.

Sketch of lower bounding the inference dimension of half spaces. Our lower bounds are based on embedding the class $\{\emptyset, \{i\} : 1 \leq i \leq n\}$ as half spaces in a way that the $n + 1$ half spaces induce the same ordering on the n points in terms of distance from the boundary. Comparison queries are useless for such an embedding, which implies that the inference dimension is at least n . Indeed, consider the half space c_\emptyset that corresponds to the empty-set; then for any subset Y of at most $n - 1$ points c_\emptyset is indistinguishable from $c_{\{i\}}$, the half space corresponding to $\{i\}$, where $i \notin Y$.

1.2.3 General framework

We next describe how the notion of inference-dimension, as well as Theorem 1.5 extend to settings where the additional queries are not necessarily comparison queries.

Consider an interactive model, where the learning algorithm is allowed to use additional queries from a prescribed set of queries \mathcal{Q} . More formally, let S be the unlabeled input sample. In pool-based active learning, the algorithm may query the label of any point in S . Here, the algorithm is allowed to use additional queries from a set $\mathcal{Q} = \mathcal{Q}(S)$ (we stress that the allowable queries depend on the input sample). For example, in the setting discussed in the previous section, $\mathcal{Q}(S)$ contains all comparison queries among points in S . Another example, which is used by crowd-sourcing algorithms (Tamuz et al., 2011) involve 3-wise queries of the form “Is x_2 more similar to x_1 than to x_3 ?”.

Upper bound. The next “boosting result” generalizes the upper bound from Theorem 1.5 and shows that if there are not too many queries in $\mathcal{Q}(S)$, or alternatively if there is an algorithm that infers the answers to the queries in $\mathcal{Q}(S)$ using a few queries, then it is possible to reveal all labels using a logarithmic number of queries:

Theorem 1.7 (Restatement of Theorem 3.2). *Let k denote the inference dimension of (X, H) . Assume that there is an algorithm that, given a realizable sample S of size n as input, uses at most $q(n)$ queries and infers the answers to all queries in $\mathcal{Q}(S)$ and all labels in S . Then there is a randomized algorithm that infers all the labels in S using at most*

$$2q(4k) \log n$$

queries in expectation.

For example, in the setting of comparison queries, $q(n) = n + n \log n$ queries suffice to infer all comparison queries and label queries in $\mathcal{Q}(S)$, and thus, the upper bound in Theorem 1.5 follows from the above theorem.

An interactive algorithm that infers all labels using a few queries can be combined with any passive learner by first using the interactive learner to reveal all labels of the input sample, and then apply the passive learner on the labeled sample. Thus, we get:

Corollary 1.8. *Let (X, H) be as in Theorem 3.2, and let $n(\epsilon, \delta)$ denote the (passive) sample complexity of learning (X, H) with error ϵ and confidence $1 - \delta$. Then there exists an algorithm that learns H with sample complexity $n(\epsilon, \delta)$, and query complexity $O(q(4k) \log n(\epsilon, \delta))$.*

Lower bound. The next lower bound on the query complexity generalizes the one from Theorem 1.5. It demonstrates a property of comparison and label queries which suffices for the lower bound in Theorem 1.5 to hold. Call an additional query *t-local*, if its answer is determined by $f(x_1), \dots, f(x_t)$ for some $x_1, \dots, x_t \in X$. For example, every label-query is 1-local and every comparison query is 2-local. The following Theorem extends Theorem 1.5 to this setting:

Theorem 1.9 (Restatement of Theorem 3.5). *Assume that the inference dimension of (X, H) is $> k$. Assume that every additional query is *t-local*, and that for every sample S the set of allowable queries $\mathcal{Q}(S)$ is the set of all queries that are determined by subsets (of size at most t) of S . Then any algorithm that reveals the labels of any realizable sample of size k must use $\Omega(k/t)$ queries in the worst-case.*

In the statistical setting, we get the following Corollary:

Corollary 1.10 (Restatement of Corollary 3.6). *Set $\epsilon = \frac{1}{k}$, $\delta = \frac{1}{6}$. Then any algorithm that learns (X, H) with error ϵ and confidence $1 - \delta$ must use $\Omega(1/t\epsilon)$ queries.*

1.3 Related work

Studying statistical learning where the learner has access to additional queries was considered by various works. A partial list includes: Angluin (1987); Baum (1991); Lang and Baum (1992); Turán (1993); Jackson (1997); Kwek and Pitt (1998); Blum et al. (1998); Bshouty et al. (2004); Feldman (2009a,b); Nowak (2011); Chen et al. (2016). Many of these focus on the case where the additional queries are *membership queries*. We discuss some of these results in Section 1.4.1.

In the context of active learning, which is considered in this paper, Balcan and Hanneke (2011) considered additional queries of two types: *Class conditional queries* and *Mistake queries*, in the first type the learner provides the annotator with a list of examples, and a label and asks her to point out an example in this list with the given label. In the second type, the learner gives the oracle a list of examples with proposed labels and she replies whether it is correct or points out a mistake. Note that these queries may have more than two answers, which is different than binary queries, which are considered in this paper.

Wauthier et al. (2012) give an active learning algorithm for clustering using pairwise similarity queries, and Ashtiani et al. (2016) gives a clustering algorithm that uses queries that ask whether two elements belong to the same cluster. Vikram and Dasgupta (2016) consider

clustering in an interactive setting where the algorithm may present the annotator a clustering of a $O(1)$ size subset of the domain, and the annotator replies whether the target-clustering agrees with this partition, and points out a difference in case it does not.

Jamieson and Nowak (2011) give an active ranking algorithm from pairwise comparisons. They consider a setting that bears resemblance with ours: their goal is to find the ranking among a sample of points in \mathbb{R}^d , where the ranking is determined according to the euclidean distance from a fixed, unknown reference point. They present an algorithm that use an expected number of $O(d \log n)$ comparisons when the ranking is chosen uniformly at random.

More recently, ? considered a similar setting to ours. They also study active classification with additional comparison queries, but they focus on minimizing the total number of label queries, while the number of comparison queries can be large (roughly $(1/\epsilon)^2$, where ϵ is the error). They also extend their study to noisy settings.

1.4 Discussion and suggested future research

We next offer few potential directions for future research.

1.4.1 Other types of additional queries

Relative queries It is natural to study other types of additional queries, and when do they enhance the learning process. Comparison queries belong to a type of queries, called relative queries, which is popular in many applications involving human annotators. A popular example of relative queries, which are used in practice, take 3 data points x_1, x_2, x_3 and queries “is x_2 more similar to x_1 than to x_3 ?”. Such queries were studied in various works, including Xing et al. (2002); Schultz and Joachims (2003); Agarwal et al. (2007); McFee and Lanckriet (2010); Huang et al. (2011); Tamuz et al. (2011); Qian et al. (2013).

Membership queries from generative models A natural and well-studied type of queries is *membership queries*; thus, the learner is allowed to query the labels of any point in the domain, including points outside the input sample. Using membership queries to enhance learning was considered in the PAC model (Baum, 1991; Turán, 1993; Jackson, 1997; Kwek and Pitt, 1998; Bshouty et al., 2004; Feldman, 2009a,b; Chen et al., 2016). However, as an experimental work by Lang and Baum (1992) suggested, querying the labels of arbitrary points may result in a poor practical performance. A possible explanation is that some domain points represent noise; for example, if the data distribution is supported on some low dimensional manifold, then querying outside it may make no sense to the human annotator. Consequently, various restrictions on this model were studied (Blum et al., 1998; Sloan and Turán, 1997; Awasthi et al., 2013; Bary-Weisberg et al., 2016). One option, studied by Awasthi et al. (2013); Bary-Weisberg et al. (2016)) is to restrict the type of membership queries by considering only local queries, within a small neighborhood of the input sample. Another possible restriction is to consider only membership queries of points that are sums and differences of a few sample points (e.g. a comparison query is a membership query on the difference of the two compared points in the case of half spaces).

An alternative direction for dealing with querying noise is to use a generative model for constructing data outside of the input sample: imagine we have an access to a generative model (e.g. Generative Adversarial Networks (Goodfellow et al., 2014)) g that gets as an input any point $s \in [0, 1]^d$, and outputs a (non-noisy) unlabeled example $g(s) \in X$. This gives an access to a large pool membership queries the learning algorithm can use.

From a theoretical perspective, this motivates the study of active learning with additional membership queries, where the learned concept is some $c : [0, 1]^d \rightarrow \{\pm 1\}$ of low complexity (say c is the sign of a low degree polynomial or a small network, etc.). It is worth noting that the case where c is a half space was studied in the discrete geometry community in the

context of linear decision trees (Meyer auf der Heide, 1984; Meiser, 1993; Liu, 2004; Cardinal et al., 2016; Ezra and Sharir, 2016). The currently best known upper bound are due to Ezra and Sharir (2016), who give an algorithm that reveals all labels of a realizable sample of size n using $\tilde{O}(d^2 \log d \log n)$ membership queries (the best known lower bound is the information theoretical $\Omega(d \log n)$, which applies more generally for any type of additional yes/no queries).

1.4.2 Noisy comparisons

It is natural to extend our algorithms by making them robust in the presence of *noisy comparisons*. Imagine, for example that $f(x_1) = 1.001$, and $f(x_2) = 1.0001$. Our algorithms are based on the assumption that the answer to the query “ $f(x_1) \geq f(x_2)$?” will always be “yes”. It is plausible to consider a noise model that may give the wrong answer in such cases where $f(x_1) \approx f(x_2)$. One such model is suggested by Bradley and Terry (1952). According to this model, when⁴ $f(x_1), f(x_2) > 0$ then the answer to the query “ $f(x_1) \geq f(x_2)$?”, is “yes” with probability $\frac{f(x_1)}{f(x_1)+f(x_2)}$.

1.4.3 Streaming-based interactive learning with comparison queries

This paper focused on the pool-based model. Another natural and well studied model is the streaming-based model. Here, the algorithm gets to observe the unlabeled examples one-by-one in an online fashion, and has to decide whether to query the current example. Consider a setting in which the algorithm has a bounded memory in which it may store M past examples. Upon receiving a new example it may query its label and/or compare it with any of the saved examples. In the realizable setting, a possible goal is to minimize the number of queries and the memory size while maintaining a (partial) hypothesis that is correct on all past examples.

Consider for example the class of threshold function over \mathbb{R} . Given an infinite sequence of independent unlabeled examples x_1, x_2, \dots drawn from an unknown distribution over \mathbb{R} . A simple online algorithm with memory $M = 2$ saves the 2 examples x, y where the sign-change occurs. Upon receiving a new example x_n , it checks whether x_n is between x and y ; if it is then it queries its label and replaces it with x or y (according to the label of x_n), otherwise the label of x_n can be inferred and the algorithm moves to the next example. One can show that this algorithm makes an expected number of $\log n$ queries after observing the first n examples. One can show that a slightly more complicated algorithm, which uses comparison queries and memory $M \leq 10$, has a similar behaviour for half-planes in \mathbb{R}^2 . We leave as an open question whether similar algorithms exist for arbitrary classes (X, H) with low inference dimension.

1.4.4 Exactly learning threshold functions

Consider the class of thresholds functions over the N -dimensional hypercube $\{0, 1\}^N$. By Lemma 4.2, the inference dimension of this class is $O(N \log N)$. Therefore, Theorem 4.1 implies that the 2^n labels of any threshold function can be revealed using at most $O(N^2 \log^2 N)$ queries in expectation. This demonstrates the additional power of comparison queries over label queries: indeed, standard arguments show that if only label queries are allowed then at least $2^{\Omega(N)}$ queries are needed.

While just $O(N^2 \log^2 N)$ queries suffice to reveal all 2^N labels, it remains open whether this can be done efficiently, in a $\text{poly}(N)$ time. A naive implementation of the algorithm from Theorem 4.1 would check at each step for each point in $\{0, 1\}^N$ whose label is still no known, whether its be inferred by the queries performed in this step. This takes at least 2^N such checks. We leave the question, of efficiently learning all 2^N labels of an unknown threshold function as an open question for future work.

⁴if $f(x_1), f(x_2) < 0$ then it is “yes” with probability $\frac{f(x_2)}{f(x_1)+f(x_2)}$.

Another comment that seems in place is that the $O(N^2 \log^2 N)$ upper bound is tight up to the $\log^2 N$ factor: indeed, it is known that there are some $2^{\Theta(N^2)}$ threshold functions (Goto and Takahasi, 1962), and therefore an information theoretic lower bound of $\Omega(N^2)$ queries holds even if the algorithm is allowed to use arbitrary yes/no queries. While it remains open whether the information theoretic barrier can be achieved using comparison queries, it is possible shave one log factor from the upper bound and achieve $O(N^2 \log N)$ queries in expectation, as we sketch next. This follows since: (i) the number of linear orders over a set $\{x_1, \dots, x_m\} \subseteq \mathbb{R}^N$ that are induced by a linear function is at most some⁵ $m^{O(2N)}$, and therefore (ii) Fredman’s algorithm (Fredman, 1976) finds such an ordering using $q(m) = O(m + N \log m)$ comparisons (which is better than the standard $O(m \log m)$ when $m \gg N$). Thus, in Theorem 3.2 we can plug $q(4k) = O(k + N \log k)$, instead of $q(4k) = O(k \log k)$ (recall that $k = O(N \log N)$, which yields $O(N^2 \log N)$ queries in expectation.

2 Preliminaries

Basic definitions. A hypothesis class is a pair (X, H) , where X is a set, and H is a class of functions $h : X \rightarrow \{\pm 1\}$. Each function $h : X \rightarrow \{\pm 1\}$ is called a *hypothesis*, or a *concept*. In this paper we study classes $H = H_{\mathcal{F}}$ of the form

$$H = \{\text{sign}(f) : f \in \mathcal{F}\},$$

where $\mathcal{F} = \{f : X \rightarrow \mathbb{R}\}$ is a class of real-valued functions. For example, when \mathcal{F} is the class of $\mathbb{R}^d \rightarrow \mathbb{R}$ affine functions then H is the class of d -dimensional half spaces. Other examples include neural-nets, low degree polynomials, and more. The class \mathcal{F} will be later used to define comparison queries. An *example* is a pair $(x, y) \in X \times \{\pm 1\}$. A *labeled sample* \bar{S} is a finite sequence of examples. We denote by S the *unlabeled sample* underlying \bar{S} that is obtained by removing the labels from the examples in \bar{S} . We will sometimes abuse notation and treat S as a subset of X (formally it is a sequence). Given a distribution D on $X \times \{\pm 1\}$, the (expected) *loss* of a hypothesis h is defined by:

$$L_D(h) \triangleq \mathbb{E}_{(x,y) \sim D} [1_{h(x) \neq y}].$$

Given a labeled sample \bar{S} , the *empirical loss* of h is defined by:

$$L_{\bar{S}}(h) \triangleq \frac{1}{|\bar{S}|} \sum_{(x,y) \in \bar{S}} 1_{h(x) \neq y}.$$

A distribution D on $X \times \{\pm 1\}$ is *realizable* by H if there exists $c \in H$ with $L_D(c) = 0$. We will sometimes refer to such a c as the “learned concept”. A labeled sample \bar{S} is *realizable* by H if there exists $c \in H$ with $L_{\bar{S}}(c) = 0$.

Passive learning (PAC learning). A learning algorithm is an (efficiently computable) mapping that gets a sample as an input and outputs a hypothesis. An algorithm A learns H if there exists a sample complexity bound $n(\epsilon, \delta)$, such that for every realizable distribution D , given a labeled sample \bar{S} of size $n \geq n(\epsilon, \delta)$, A outputs $h = A(\bar{S})$ such that:

$$\Pr_{\bar{S} \sim D^n} [L_D(h) > \epsilon] \leq \delta.$$

The parameter ϵ is called the error of the algorithm, and $1 - \delta$ is called the confidence of the algorithm. We will assume throughout that a learning algorithm for H also receives ϵ, δ as part of the input and can compute $n(\epsilon, \delta)$.

⁵if $f_1, f_2 : \mathbb{R}^N \rightarrow \mathbb{R}$ are linear and induce a different linear orders on $\{x_1, \dots, x_m\}$, then $\text{sign}(f_1(x_i - x_j)) \neq \text{sign}(f_2(x_i - x_j))$ for some i, j , and so the number of distinct orders is at most the number of sign patterns $(\text{sign}(f(x_i - x_j)))_{1 \leq i < j \leq m}$, which is at most $O(m^{2N})$.

2.1 Active learning

It is helpful to recall the framework of active learning before extending it by allowing additional queries. A (pool-based) active learning algorithm has an access to the unlabeled sample S underlying the input labeled sample \bar{S} . It queries the labels of a subsample of it, and outputs a hypothesis. The choice of which subsample A queries may be adaptive. Each active learning algorithm is associated with two complexity measures: (i) the sample-complexity $n(\epsilon, \delta)$, is the number of examples required to achieve error at most ϵ with confidence at least $1 - \delta$ (like in the passive setting), and (ii) the query-complexity (also called label complexity) $q(\epsilon, \delta)$, is the number of queries it makes.

In the process of active learning, it is natural to distinguish between points whose label can be inferred and points for which there is uncertainty concerning their label. It is therefore convenient to consider partial hypotheses. A *partial hypothesis* is a partially labeled hypothesis $h : X \rightarrow (Y \cup \{?\})$, where if $h(x) = "?"$ it means that h abstains from labeling x . We extend the 0/1 loss function to "?", such that abstaining is always treated as a mistake. The *coverage* of h with respect to a distribution D is defined as

$$C_D(h) \triangleq \Pr_{x \sim D} [h(x) \neq ?],$$

and the *empirical-coverage* of h with respect to a sample S is defined as

$$C_S(h) \triangleq \frac{|\{x \in S_X : h(x) \neq ?\}|}{|S|}.$$

Since abstaining counts as error it follows that $C_D(h) \leq 1 - L_D(h)$ for every partial hypothesis h , and every distribution D .

Confident algorithms. A learning algorithm is *confident* if it outputs a partial hypothesis that is correct on all points where it does not abstain:

Definition 2.1 (Confident learning algorithm). *A learning algorithm A is confident with respect to a class (X, H) if it satisfies the following additional requirement. For every realizable labeled sample \bar{S} that is consistent with a learned concept $c \in H$, the output hypothesis $h \triangleq A(\bar{S})$ satisfies that whenever $h(x) \neq ?$ then $h(x) = c(x)$.*

Thus, if A is confident then $C_D(h) = 1 - L_D(h)$ for every distribution D , where $h = A(\bar{S})$. Therefore, in the context of confident learners we will only discuss their coverage, and omit explicit reference to their error. For example a class (X, H) is learned by a confident learner A if there exists a sample complexity $n(\epsilon, \delta)$ such that for every realizable distribution D , given a labeled sample \bar{S} of size $n \geq n(\epsilon, \delta)$, A outputs $h = A(\bar{S})$ such that

$$\Pr_{\bar{S} \sim D^n} [C_D(h) < 1 - \epsilon] \leq \delta.$$

2.2 Interactive learning with additional queries

Consider an extension of the active learning setting, by allowing the learning algorithm to use additional queries from a prescribed set of queries \mathcal{Q} . An additional query is modeled as a boolean function $q : \mathcal{F} \rightarrow \{True, False\}$. We stress that the query may depend on the function f that underlies the learned concept $c = \text{sign}(f)$ (consider a comparison query as an example). In this setting, given an input sample \bar{S} , the algorithm is given access to S , the unlabeled sample underlying \bar{S} , and is allowed to:

- query the label of any of point in S ,

- query an additional query q from a prescribed set of queries $\mathcal{Q}(S)$.

We stress that the set of allowable queries $\mathcal{Q}(S)$ depend on the input sample \bar{S} . For example, a comparison query on x_1, x_2 is the query “ $f(x_1) \leq f(x_2)$?”, where $c = \text{sign}(f)$ is the learned concept. An *answered query* in a pair (q, b) , where q is a comparison/label-query, and b is a possible answer to q . For example, (“ $f(x_1) \leq f(x_2)$?”, $True$) is an answered comparison query, and (“ $c(x_1) = ?$ ”, -1) is an answered label query. The standard notions of *version space* and *agreement set* are naturally extended to this context. Let

$$\bar{Q} = ((q_1, b_1), (q_2, b_2), \dots, (q_m, b_m))$$

be a sequence of answered queries. Define the *version space*, denoted by $V(\bar{Q})$, as the set of hypotheses in H that are consistent with \bar{Q} :

$$V(\bar{Q}) \triangleq \{h \in H : q_i(h) = b_i, i = 1, \dots, m\},$$

and the confidence region, $\text{Conf}(\bar{Q})$, as the agreement set of $V(\bar{Q})$:

$$\text{Conf}(\bar{Q}) \triangleq \{x \in X : \text{all hypotheses in } V(\bar{Q}) \text{ agree on } x\}.$$

3 Inference Dimension

Let (X, H) be a hypothesis class, where $H = H_{\mathcal{F}}$ for some class of real functions \mathcal{F} , and let \mathcal{Q} be a set of additional queries. For $S \subseteq X$, $x \in X$, $f : X \rightarrow \mathbb{R}$ where $f \in \mathcal{F}$ and $c = \text{sign}(f) \in H$, let

$$S \xrightarrow[f]{} x$$

denote the statement that there exists a sequence \bar{Q} of answered label queries of S and/or additional queries from $\mathcal{Q}(S)$ that determine the label of x , when the learned concept is c . Namely, that $x \in \text{Conf}(\bar{Q})$.

Definition 3.1 (Inference dimension). *The inference dimension of (X, H) is the minimal number k such that for every $S \subseteq X$ of size k , and every $c \in H$ there exists $x \in S$ such that*

$$S \setminus \{x\} \xrightarrow[f]{} x.$$

If no such k exists then the inference dimension of (X, H) is defined as ∞ .

3.1 Upper bound

Theorem 3.2 (Boosting). *Let k denote the inference dimension of (X, H) . Assume that there is an algorithm that, given a realizable sample \bar{S} of size n as input, uses at most $q(n)$ queries and infers the answers to all queries in $\mathcal{Q}(S)$ and all labels in \bar{S} . Then there is a randomized algorithm that infers all the labels in \bar{S} using at most*

$$2q(4k) \log n$$

queries in expectation.

We prove Theorem 3.2 in two steps: (i) First, Lemma 3.3 shows that (X, H) has a weak confident learner A that, given a realizable input sample of size $4k$, uses at most $q(4k)$ queries, and outputs a partial hypothesis h with coverage $1/2$. (ii) Then, we show that the labels of a given sample \bar{S} of size n are revealed after applying A on roughly $\log n$ random subsamples of \bar{S} .

Lemma 3.3 (Weak confident-learning). *Let k denote the inference dimension of (X, H) . Then there exists a confident learner for (X, H) that is defined on input samples of length $4k$, makes at most $q(4k)$ queries, and has coverage $\geq 1/2$ with probability $\geq 1/2$. That is, for any distribution D over X ,*

$$\Pr_{S \sim D^{4k}} [C_D(h) \geq 1/2] \geq 1/2,$$

where h is the output hypothesis of the algorithm.

Proof. The learner is defined as follows. Given a realizable input sample $S = ((x_i, c(x_i)))_{i=1}^{4k}$, from D^{4k} , the algorithm infers the answers to all queries in $\mathcal{Q}(S)$ and all labels in \bar{S} (by assumption, this can be done with $q(4k)$ queries). It outputs the partial hypothesis h , which labels any x whose label can be inferred from the queries. Namely:

$$h(x) \triangleq \begin{cases} c(x) & S \xrightarrow{f} x \\ ? & \text{otherwise} \end{cases}$$

We next claim that the expected coverage of h , $C_D(h)$, is at least $3/4$. This implies that $\Pr_S[C_D(h) \geq 1/2] \geq 1/2$, which shows that the learning algorithm is a weak confident learner.

To this end we use the following observation.

Observation 3.4. *For any set Y of size $4k + 1$, there are $x_{i_1}, \dots, x_{i_{3k+1}} \in Y$ such that for all $1 \leq j \leq 3k + 1$ it holds that*

$$Y \setminus \{x_{i_j}\} \xrightarrow{f} x_{i_j}.$$

Proof. This follows since the inference dimension of (X, H) is k . Assume we already constructed $x_{i_1}, \dots, x_{i_{j-1}}$ for $j \leq 3k + 1$. Let $Y' = Y \setminus \{x_{i_1}, \dots, x_{i_{j-1}}\}$. As $|Y'| \geq k$, there exists $x_{i_j} \in Y'$ such that $Y' \setminus \{x_{i_j}\} \xrightarrow{f} x_{i_j}$. But then also $Y \setminus \{x_{i_j}\} \xrightarrow{f} x_{i_j}$. \square

Next, we show that this observation implies that $\mathbb{E}_S[C_D(h)] \geq 3/4$. Clearly, we have that

$$\mathbb{E}_S[C_D(h)] = \Pr_{(x_1, x_2, \dots, x_{4k+1}) \sim D^{4k+1}} [\{x_1, \dots, x_{4k}\} \xrightarrow{f} x_{4k+1}].$$

Letting $T = \{x_1, x_2, \dots, x_{4k+1}\}$, this is the probability that $T \setminus \{x_{4k+1}\} \xrightarrow{f} x_{4k+1}$. However, by symmetry, this is the same as the probability that $T \setminus \{x_i\} \xrightarrow{f} x_i$ for any $1 \leq i \leq 4k + 1$.

Taking the average over i , we have

$$\mathbb{E}_S[C_D(h)] = \mathbb{E}_T \left[\frac{1}{4k+1} |\{i : T \setminus \{x_i\} \xrightarrow{f} x_i\}| \right] \geq \frac{3k+1}{4k+1} \geq \frac{3}{4}.$$

\square

Proof of Theorem 3.2. Let A denote the weak confident learner from Lemma 3.3. Let \bar{S} be a realizable sample, corresponding to an unknown concept $c \in H$. Our goal is to fully recover the labels of \bar{S} .

The algorithm proceed in iterations $t = 1, 2, \dots$. At each iteration it applies A on a sub-sample of size $4k$ of \bar{S} . Let h_t denote the output hypothesis of A on iteration t , and let

$$DIS_t = \{x : h_s(x) = ? \text{ for all } s \leq t\}.$$

Since A is confident it follows for any point $x \notin DIS_t$, the label $c(x)$ is equal to $h_s(x)$ for any $h_s(x)$ such that $h_s(x) \neq ?$. As long as $DIS_t \cap \bar{S} \neq \emptyset$, perform the following update step.

Update step at time t :

- (1) Let D_t be the uniform distribution over $DIS_t \cap S$. Sample $\bar{R}_t \sim (D_t)^{4k}$.
- (2) Apply A to R_t , the unlabeled sample underlying \bar{R}_t .
- (3) Let $h_t = A(\bar{R}_t)$ be the confident partial hypothesis that A outputs on \bar{R}_t .
- (4) Compute $e_t = C_{D_t}[h_t] = \Pr_{x \sim D_t}[h_t(x) \neq ?]$.
- (5) If $e_t < 1/2$ then go back to step (1). Otherwise set $t \leftarrow t + 1$ and continue.

Since A is confident, it follows that once $DIS_t \cap S = \emptyset$ then all the labels of \bar{S} are revealed.

Query-complexity. In order to analyze the query-complexity of the algorithm, first observe that since $\Pr[e_t \geq 1/2] \geq 1/2$ then in expectation, we proceed to the next iteration after at most two samples of \bar{R}_t . Next, if $e_t \geq 1/2$ then by definition $|DIS_{t+1} \cap S_X| \leq |DIS_t \cap S_X|/2$. Thus, we only apply the update step at most $t_{max} \leq 2 \log n$ many times. It follows that the expected query-complexity is at most $2q(4k) \log n$. \square

Computational complexity. The algorithm derived in Theorem 3.2 has expected running time of $O(T_{\text{update}} \log n)$, where T_{update} is the running time of the update step. In every update step the algorithm makes $q(4k)$ queries and determines e_t , by checking for each unlabeled point, whether its label can be inferred by the queries performed in this step. Assume that testing this for each point takes T_{infer} . So,

$$T_{\text{update}} = O(q(4k) + n \cdot T_{\text{infer}})$$

and the total running time is

$$T_{\text{total}} = O((q(4k) + n \cdot T_{\text{infer}}) \log n).$$

For example, when the hypothesis class is half spaces in \mathbb{R}^d , and the the set of additional queries is comparisons, the total running time is polynomial in n . This is since $q(4k) = O(k \log k)$ (by sorting), and since checking if the label of a point is inferred by a set of label and comparison queries can be phrased as a linear program and solved in polynomial time (see Claim 4.5).

3.2 Lower bound

Next, we show that if the inference dimension is large then many queries are needed to infer all the labels. We further assume that every query is t -local, in the sense that it depends on $f(x_1), \dots, f(x_t)$ for some $x_1, \dots, x_t \in X$. We set of allowable queries $\mathcal{Q}(S)$ to be all queries that are associated with subsets of S of size t .

Let (X, H) be a hypothesis class with inference dimension $> k$, for some $k \geq 3$. This means that there exists $Z \subseteq X$ of size k and a concept $c \in H$ such that for every $z \in Z$ there is $c_z \in C$ with $c_z(z) \neq c(z)$, but $c_z(x) = c(x)$ for all $x \in Z \setminus \{z\}$, and moreover, $q(c) = q(c_z)$ for every query $q \in \mathcal{Q}(Z \setminus \{z\})$.

Theorem 3.5. *Any algorithm that reveals the labels of any realizable sample of size k must use at least k/t queries in the worst-case.*

Proof. Consider the realizable sample $\bar{S} = (x, c(x))_{x \in Z}$, and assume that less than k/t queries were used on it. This means that there exist $z \in S$ that is not associated with any of the queries that were used, and hence only queries from $\mathcal{Q}(Z \setminus \{z\})$ were used. However, this implies that both c_z and c are consistent with the queries that were used, and $c_z(z) \neq c(z)$. Thus, the label of z can not be inferred by the queries that were used. \square

Corollary 3.6. *Let $\epsilon = \frac{1}{k}, \delta = \frac{1}{6}$, and let D be the uniform distribution over Z . Then any learning algorithm that makes less than $\frac{1}{t\epsilon}$ queries suffers a loss of ϵ , with probability at least δ .*

Proof. Let A be a learning algorithm. Consider an adversary that picks the secret-concept $c = c_z$ where $z \in Z$ is chosen uniformly. Define

$$E \triangleq \{c_x : \text{a query associated with } x \text{ was queried by } A\}.$$

If A makes less than $\frac{k}{2t}$ queries (i.e. $|E| < \frac{k}{2}$) then:

1. $\Pr[c_z \notin E] \geq \frac{1}{2}$, and
2. $\Pr\left[L(A) \geq \frac{1}{k} \mid c_z \notin E\right] \geq 1 - \frac{2}{k} \geq \frac{1}{3}$,

where

$$L(A) = \frac{|\{x \in Z : h(x) \neq c_z(x)\}|}{k}$$

is the loss of the hypothesis h that A outputs. To see why the second item holds, note that since the answers to the queries are the same for every $c_x \notin E$, it follows that the distribution of c_z conditioned on $c_z \notin E$ is uniform over E^c whose size is at least $\frac{k}{2}$, and so the probability that A outputs c_i is at most $\frac{2}{k}$. In any other case, the loss of h is at least $\frac{1}{k}$.

Combining the above two items together yields that

$$\Pr\left[L(A) \geq \frac{1}{k}\right] = \Pr[c_z \notin E] \cdot \Pr\left[L(A) \geq \frac{1}{k} \mid c_z \notin E\right] \geq \frac{1}{6}.$$

\square

4 Interactive learning of half spaces with comparison-queries

In this section we restrict our attention to the class $H_d = \{\text{sign}(f) : f : \mathbb{R}^d \rightarrow \mathbb{R}\}$ of half spaces in \mathbb{R}^d , where for simplicity of exposition we consider linear functions f (these correspond to homogeneous half spaces). Our results extend to the non-homogeneous case, as non-homogeneous half spaces in dimension d can be embedded as homogeneous half spaces in dimension $d+1$. The additional queries allowed are comparison queries. That is, a label query returns the answer to $\text{sign}(f(x))$ and a comparison query returns the answer to $f(x_1) \geq f(x_2)$.

In Subsection 4.1 we present our upper bounds on the query complexity, under two natural conditions: small bit complexity, or large margin. In Subsection 4.2 we present lower bounds showing that these conditions are indeed necessary for obtaining query complexity sub-linear in the sample complexity.

4.1 Upper bounds

4.1.1 Bit-complexity

Here we show that if the examples can be represented using a bounded number of bits then comparison-queries can reduce the query-complexity. We formalize bounded bit-complexity by assuming that $X = [N]^d$, where $[N] = \{0, \dots, N\}$. Note that each example can be represented by $B = d \log N$ bits. We provide a bound on the query-complexity that depends efficiently on d and $\log N$. Variants of the arguments we use apply to other standard ways of quantifying bounded bit-complexity.

Theorem 4.1. Consider the class $([N]^d, H_d)$. There exists an algorithm that reveals the labels of any realizable input sample of size n using at most $O(k \log k \log n)$ label/comparison-queries in expectation, where $k = O(d \log(Nd))$.

As a consequence it follows that the hypothesis class $([N]^d, H_d)$ is learnable with

$$\text{sample complexity } \tilde{O}(d/\epsilon) \text{ and query-complexity } \tilde{O}(d \log(N) \log(1/\epsilon)),$$

where the \tilde{O} notation suppresses lower order terms and the usual $\log(1/\delta)$ dependence.

In order to prove the above theorem, we use Theorem 1.5 that reduces it to the following lemma.

Lemma 4.2. Let k such that $2^{k/2} > 2(kN + 1)^d$. Then the inference dimension of the class $([N]^d, H_d)$ is at most k . In particular, it is at most $16d \log(4Nd)$.

Proof. Let $c = \text{sign}(f) \in H_d$. We will use the following claims.

Observation 4.3. Let $x \in \mathbb{R}^d$, and $y \in \{\pm 1\}$. Then

$$yf(x) > 0 \implies c(x) = y.$$

Observation 4.4. Let $x_1, x_2 \in \mathbb{R}^d$ such that $c(x_1) = c(x_2) = y$. Then

$$yf(x_2 - x_1) > 0 \implies |f(x_2)| > |f(x_1)|.$$

Claim 4.5. Let $x_1, \dots, x_m \in \mathbb{R}^d$, $y \in \{\pm 1\}$ such that:

- (i) $c(x_i) = y$ for all i , and
- (ii) $|f(x_1)| < \dots < |f(x_m)|$.

Then, if $x \in \mathbb{R}^d$ satisfies $x - x_1 = \sum_{i=1}^{m-1} \alpha_i (x_{i+1} - x_i)$ where $\alpha_i \geq 0$, then $c(x) = y$. In particular $\{x_1, \dots, x_m\} \xrightarrow{f} x$.

Proof. If $x = x_1$ then clearly $c(x) = y$. Thus, assume $x \neq x_1$. Therefore, at least one of the α_i 's is positive, and by linearity of f , and Observation 4.4:

$$yf(x - x_1) = \sum_i \alpha_i yf(x_{i+1} - x_i) > 0.$$

Therefore $yf(x) > yf(x_1) \geq 0$, which implies that $c(x) = y$ (by Observation 4.3). \square

Let k be as in the formulation of the lemma, and let $Y \subseteq [N]^d$ of size k . For simplicity of exposition assume k is even, and let $m = k/2$. We need to show that there exists $x \in Y$ such that $Y \setminus \{x\} \xrightarrow{f} x$. Without loss of generality, assume that at least m of the points in Y have a 1-label. Let $x_1, x_2, \dots, x_m \in Y$ be an ordering of the 1-labeled points such that $|f(x_1)| \leq \dots \leq |f(x_m)|$. By Claim 4.5 it suffices to show that there exists $2 \leq i^* \leq m$ such that

$$x_{i^*} - x_1 = \sum_{i=1}^{i^*-2} \alpha_i (x_{i+1} - x_i) \text{ where } \alpha_i \geq 0. \quad (1)$$

This follows by the following pigeon hole argument: consider all possible boolean combinations of the form $\sum_{i=1}^{m-1} \beta_i (x_{i+1} - x_i)$ where $\beta_i \in \{0, 1\}$. There are 2^{m-1} boolean combinations, each of which yields a vector in $\{-mN, \dots, mN\}^d$. Therefore, by our assumptions on k we have that

$$2^{m-1} > (2mN + 1)^d$$

there exists two boolean combinations:

$$\sum_{i=1}^{m-1} \beta_i(x_{i+1} - x_i) = \sum_{i=1}^{m-1} \gamma_i(x_{i+1} - x_i),$$

such that $\beta_i \neq \gamma_i$ for at least one i . Let i_0 be the maximal such i . We verify that $i^* = i_0 + 1$ satisfies Equation (1) above.

Assume without loss of generality that $\beta_{i_0} = 0, \gamma_{i_0} = 1$. Let $\alpha_i = \beta_i - \gamma_i$. Thus $\alpha_{i_0} = -1, \alpha_i = 0$ for all $i > i_0$ and $\alpha_i \in \{0, \pm 1\}$ for $i < i_0$. So

$$\sum_{i=1}^{i_0} \alpha_i(x_{i+1} - x_i) = 0.$$

Next, add $x_{i_0+1} - x_1 = \sum_{i=1}^{i_0} (x_{i+1} - x_i)$ to both sides. This yields

$$x_{i_0+1} - x_1 = \sum_{i=1}^{i_0} (\alpha_i + 1)(x_{i+1} - x_i).$$

This concludes the proof as $\alpha_{i_0} + 1 = 0$ and $\alpha_i + 1 \geq 0$ for all $i < i_0$. \square

4.1.2 Minimal-ratio and margin

Let $X \subseteq \mathbb{R}^d$ and let $c = \text{sign}(f) \in H_d$. The *minimal-ratio* of X with respect to c is defined by

$$\eta = \eta(c, X) \triangleq \frac{\min_{x \in X} |f(x)|}{\max_{x \in X} |f(x)|}.$$

Here we show that it is possible to reveal all labels using at most $\tilde{O}(d \log(1/\eta) \log n)$, where the minimal-ratio is η . Note that the minimal-ratio is invariant under scaling and that it is upper bounded by the margin:

Claim 4.6. *Let η be the minimal-ratio of X with respect to c , let $\rho = \max_{x \in X} \|x\|_2$, and let γ be the margin of X with respect to c . Then*

$$\frac{\gamma}{\rho} \leq \eta.$$

Proof. Let w such that $f(x) = \langle w, x \rangle$, and set $w' = \frac{w}{\|w\|}$. Thus,

$$\frac{\gamma}{\rho} = \frac{\min_x |\langle w', x \rangle|}{\max_x \|x\|_2} \leq \frac{\min_x |\langle w', x \rangle|}{\max_x |\langle w', x \rangle|} = \frac{\min_{x \in X} |f(x)|}{\max_{x \in X} |f(x)|} = \eta.$$

\square

Thus, the upper bound in Theorem 4.7 below applies when the minimal-ratio is replaced by the standard margin parameter γ/ρ .

Note that there are cases where $\eta \gg \gamma/\rho$. For example, assume $X = \{e_1, \dots, e_d\}$ is the standard basis and $c_w \in H_d$ is determined by the normal $w = \frac{1}{\sqrt{d}}(+1, -1, +1, -1, \dots)$. In this case, $\gamma/\rho = 1/\sqrt{d} \ll 1 = \eta$.

We next state and prove the upper bound. Let $X \subseteq \mathbb{R}^d$, and let $H_{d,\eta} \subseteq H_d$ be the set of all half spaces with minimal-ratio at least η with respect to X .

Theorem 4.7. *Consider the class $(X, H_{d,\eta})$. There exists an algorithm that reveals the labels of any realizable input sample of size n using at most $O(k \log k \log n)$ label/comparison-queries in expectation, where $k = O(d \log(d) \log(1/\eta))$.*

As a consequence it follows that the hypothesis class $(X, H_{d,\eta})$ is learnable with

$$\text{sample complexity } \tilde{O}(d/\epsilon) \text{ and query-complexity } \tilde{O}(d \log(1/\eta) \log(1/\epsilon)),$$

As before, the \tilde{O} notation suppresses lower order terms and the usual $\log(1/\delta)$ dependence.

The above theorem is a corollary of Theorem 1.5 via the following lemma, which upper bounds the inference dimension of the class $(X, H_{d,\eta})$.

Lemma 4.8. *Let k such that $(k/2 + 1)^d < 2^{k/2}(\eta/6)^d$. Then the inference dimension of the class $(X, H_{d,\eta})$ is at most k . In particular, it is at most $10d \log(d + 1) \log(2/\eta)$.*

Proof. Let $c = \text{sign}(f) \in H_{d,\eta}$. We will use the following claims (note the analogy with the claims in the proof of Lemma 4.2).

Observation 4.9. *Let $x \in X$, and $y \in \{\pm 1\}$. Then*

$$\frac{yf(x)}{\max_{x' \in X} |f(x')|} \geq -\eta \implies c(x) = y.$$

Claim 4.10. *Let $x_1, \dots, x_m \in X$, $y \in \{\pm 1\}$ such that:*

(i) $c(x_i) = y$ for all i , and

(ii) $|f(x_1)| < \dots < |f(x_m)|$.

Then, if $x \in X$ satisfies $x - x_1 = \sum_{i=1}^{m-1} \alpha_i(x_{i+1} - x_i)$ where $\alpha_i \geq -\eta$, then $c(x) = y$. In particular $\{x_1, \dots, x_m\} \xrightarrow{f} x$.

Proof. Assume towards contradiction that $c(x) \neq y$, namely that $yf(x) \leq 0$. By assumption, x satisfies:

$$yf(x - x_1) = \sum_{i=1}^{m-1} \alpha_i yf(x_{i+1} - x_i) \geq -\eta y \sum_{i=1}^{m-1} f(x_{i+1} - x_i) = -\eta yf(x_m - x_1),$$

where the equality follows by linearity of f and our assumptions. Therefore,

$$yf(x) \geq yf(x_1) - \eta yf(x_m) + \eta yf(x_1) \geq -\eta yf(x_m),$$

which implies that

$$0 \geq \frac{yf(x)}{\max_{x' \in X} |f(x')|} \geq \frac{yf(x)}{|f(x_m)|} = \frac{yf(x)}{yf(x_m)} \geq -\eta.$$

By Observation 4.9 this implies that $c(x) = y$, which contradicts the assumption that $c(x) \neq y$. \square

Let k be as in the formulation of the lemma, and let $Y \subseteq X$ of size k . For simplicity of exposition assume k is even, and let $m = k/2$. We need to show that there exists $x \in Y$ such that $Y \setminus \{x\} \xrightarrow{f} x$. Without loss of generality, assume that at least half of the points in Y have a 1-label. Let $x_1, x_2, \dots, x_m \in Y$ be an ordering of the 1-labeled points such that $|f(x_1)| \leq \dots \leq |f(x_m)|$. By Claim 4.10 it suffices to show that there exists $2 \leq i^* \leq m$ such that

$$x_{i^*} - x_1 = \sum_{i=1}^{m-1} \alpha_i(x_{i+1} - x_i) \text{ where } \alpha_i \geq -\eta \text{ and } \alpha_{i^*-1} = \alpha_{i^*}. \quad (2)$$

This follows by a volume argument: let \mathcal{C} be the convex hull of $\{x_2 - x_1, x_3 - x_1, \dots, x_m - x_1\}$. For a set $A \subseteq [m]$, let⁶

$$\mathcal{C}_A \triangleq \left(\sum_{j \in A} x_{j+1} - x_j \right) + \frac{\eta}{6} \mathcal{C}.$$

We claim that there are $A \neq B$ such that $\mathcal{C}_A \cap \mathcal{C}_B \neq \emptyset$. Indeed, assume towards contradiction that the \mathcal{C}_A are all mutually disjoint. Note that for all A : (i) $\mathcal{C}_A \subseteq (m + \eta/6)\mathcal{C}$, and (ii) $\text{vol}(\mathcal{C}_A) = (\eta/6)^d \text{vol}(\mathcal{C})$. So, if all the \mathcal{C}_A 's are mutually disjoint then

$$(m + 1)^d \text{vol}(\mathcal{C}) \geq 2^m (\eta/6)^d \text{vol}(\mathcal{C}),$$

which means that $(m + 1)^d > 2^m (\eta/6)^d$, which contradicts the property k is assumed to satisfy.

Thus, there exist two combinations:

$$\sum_{i=1}^m \beta_i (x_{i+1} - x_i) = \sum_{i=1}^m \gamma_i (x_{i+1} - x_i),$$

such that $\beta_i \in [b_i, b_i + \gamma/6]$ and $\gamma_i \in [c_i, c_i + \gamma/6]$ for some $b_i, c_i \in \{0, 1\}$, not all the same. Let i_0 be maximal such that $b_{i_0} \neq c_{i_0}$. We will prove that Equation (2) holds for $i^* = i_0 + 1$.

Assume without loss of generality that $b_{i_0} = 0, c_{i_0} = 1$. Define $\alpha_i = \beta_i - \gamma_i$. Thus $\alpha_{i_0} \in [-1 - \eta/6, -1 + \eta/6]$, and $\alpha_i \geq -\eta/6$ for all $i > i_0$. Now, adding

$$(\alpha_{i_0+1} - \alpha_{i_0})(x_{i_0+1} - x_1) = \sum_{j=1}^{i_0} (\alpha_{i_0+1} - \alpha_{i_0})(x_{i+1} - x_i)$$

to both sides of $\sum_{j=1}^m \alpha_i (x_{i+1} - x_i) = 0$ gives:

$$\begin{aligned} & (\alpha_{i_0+1} - \alpha_{i_0})(x_{i_0+1} - x_1) = \\ & \sum_{j=1}^{i_0-1} (\alpha_i + \alpha_{i_0+1} - \alpha_{i_0})(x_{i+1} - x_i) + \alpha_{i_0+1}(x_{i_0+2} - x_{i_0}) + \sum_{j=i_0+2}^m \alpha_i (x_{i+1} - x_i). \end{aligned}$$

We conclude the proof by verifying that Equation (2) is satisfied after after dividing both sides of the above formula by

$$\alpha_{i_0+1} - \alpha_{i_0} \geq 1 - \eta/3 \geq 1/2.$$

Indeed, the coefficients in front of $(x_{i_0} - x_{i_0-1})$ and $(x_{i_0+1} - x_{i_0})$ are both equal to α_{i_0+1} ; if $i \geq i_0 + 2$ then as $\alpha_i \geq -\eta/6$ we have $\alpha_i / (\alpha_{i_0+1} - \alpha_{i_0}) \geq (-\eta/6) / (1/2) \geq -\eta$; and if $i \leq i_0 - 1$ then as $\alpha_i + \alpha_{i_0+1} - \alpha_{i_0} \geq -(3\eta)/6 = -\eta/2$, we have $(\alpha_i + \alpha_{i_0+1} - \alpha_{i_0}) / (\alpha_{i_0+1} - \alpha_{i_0}) \geq -\eta$. \square

4.2 Lower bounds

In this Section we show that (in the worst-case) comparison-queries yield no advantage when the bit complexity is large in dimension $d \geq 3$, or when the dimension is large even if the margin is large.

4.2.1 Dimension $d \geq 3$

We show that (in the worst-case) comparison-queries do not yield a significant saving in query complexity for learning half spaces, already in \mathbb{R}^3 . This is tight since, as discussed in the introduction, in \mathbb{R}^2 comparison-queries yield an exponential saving.

⁶We use here the standard notation of $u + \alpha A = \{u + \alpha a : a \in A\}$, where $u \in \mathbb{R}^d, \alpha \in \mathbb{R}$, and $A \subseteq \mathbb{R}^d$.

Theorem 4.11. Consider the class (\mathbb{R}^3, H_3) of half spaces in \mathbb{R}^3 , Any algorithm that reveals the labels of any realizable sample of size n must use $\Omega(n)$ comparison/label queries in the worst-case.

In the statistical setting, we get that

Corollary 4.12. Let $\epsilon > 0$. Then any algorithm that learns (\mathbb{R}^3, H_3) with error ϵ and confidence at least $5/6$ must use $\Omega(1/\epsilon)$ comparison/label queries on some realizable distributions.

We derive these statements by showing that the inference dimension of (\mathbb{R}^3, H_3) is ∞ . Then, Theorem 4.11 and Corollary 4.12 follow by plugging $t = 2$ in Theorem 3.5 and Corollary 3.6 respectively. (Note that comparison queries are 2-local, and thus $t = 2$).

Theorem 4.13. The inference dimension of (\mathbb{R}^3, H_3) is ∞ .

Proof. We need to show that for every n , there is $c = \text{sign}(f) \in H_3$, and a set $X_n \subseteq \mathbb{R}^3$ of n points such that for every $x \in X_n$, it is **not** the case that $X_n \setminus \{x\} \xrightarrow{f} x$. We use the following lemma.

Lemma 4.14. There exists a 3 dimensional linear space V_n of functions $g : \{1, \dots, n\} \rightarrow \mathbb{R}$, and $n + 1$ functions $g_0, g_1, \dots, g_n \in V_n$ with the following properties:

- $g_i(j) > 0$ if and only if $i \neq j$ (in particular, $g_0(j) > 0$ for all j).
- $|g_i(1)| < |g_i(2)| < \dots < |g_i(n)|$ for all $0 \leq i \leq n$.

We first use Lemma 4.14 to prove Theorem 4.13. Let $v_1, v_2, v_3 \in V_n$ be a basis for the linear space from Lemma 4.14. Define

$$x_j \triangleq (v_1(j), v_2(j), v_3(j)) \in \mathbb{R}^3$$

and set $X_n = \{x_1, \dots, x_n\}$. Note that each function g_i can be represented by a linear combination $\alpha_{i,1}v_1 + \alpha_{i,2}v_2 + \alpha_{i,3}v_3$. Define $f_i : \mathbb{R}^3 \rightarrow \mathbb{R}$ by $f_i = \langle \alpha_i, x \rangle$, and set $c_i = \text{sign}(f_i)$. Thus, for all i, j ,

$$f_i(x_j) = g_i(j). \tag{3}$$

Consider the set $X_n \setminus \{x_i\}$. Equation (3) implies that there is no comparison nor label query that distinguishes $c_0 = \text{sign}(f_0)$ from $c_i = \text{sign}(f_i)$ on this set. This means that it is **not** the case that $X_n \setminus \{x_i\} \xrightarrow{f_0} x_i$, which finishes the proof of Theorem 4.13.

It remains to prove Lemma 4.14.

Proof of Lemma 4.14. For a sufficiently large M , let V_n denote the space of functions $g : \{1, \dots, n\} \rightarrow \mathbb{R}$ that is spanned by the three functions $M^x, x \cdot M^x, x^2 \cdot M^x$. Define

$$g_i(x) = M^x(1 - 2(x - i)^2) \in V_n.$$

It is easy to check that the first item in the conclusion is satisfied by the g_i 's, and that if M is sufficiently large, then so does the second item. □

□

4.2.2 Margin

We show here that, in the worst-case, comparison-queries do not yield a significant saving in query complexity for learning half spaces, even if it is guaranteed that the margin is large, say at least $1/8$.

Theorem 4.15. *For every n there is a class (X, H) , where $X \subseteq \mathbb{R}^{n+1}$, and $H \subseteq H_{n+1}$ contains all the half spaces with margin at least $1/8$ such that the following holds: any algorithm that reveals the labels of any realizable sample of size n must use $\Omega(n)$ comparison/label queries in the worst-case.*

In the statistical setting, we get that

Corollary 4.16. *For every $\epsilon > 0$, there is n and a class (X, H) , where $X \subseteq \mathbb{R}^{n+1}$, and $H \subseteq H_{n+1}$ contains all the half spaces with margin at least $1/8$ such that the following holds: any algorithm that learns (X, H) with error ϵ and confidence at least $5/6$ must use $\Omega(1/\epsilon)$ comparison/label queries on some realizable distributions.*

We derive these statements by establishing the existence of classes with large margin and large inference dimension. Then, Theorem 4.15 and Corollary 4.16 follow by plugging $t = 2$ in Theorem 3.5 and Corollary 3.6 respectively.

Theorem 4.17. *For every n , there is a set of n unit vectors $X = \{x_1, x_2, \dots, x_n\} \subset \mathbb{R}^{n+1}$ such that the class (X, H) has inference dimension at least n , where H contains all half spaces with margin at least $1/6$ with respect to X .*

Proof. Define $x_i \triangleq \frac{e_i + e_{n+1}}{\sqrt{2}}$ where the e_i 's are the vectors in the standard basis, and define $w_i \in \mathbb{R}^{n+1}$ for $i = 1, \dots, n$ by

$$w_i(j) \triangleq \begin{cases} 1 + \frac{j}{10n^2} & j = i \\ -\frac{1}{2} & j = n + 1 \\ -\frac{j}{10n^2} & \text{otherwise} \end{cases}$$

and $w_0 \in \mathbb{R}^{n+1}$ by

$$w_0(j) \triangleq \begin{cases} -\frac{1}{2} & j = n + 1 \\ -\frac{j}{10n^2} & \text{otherwise} \end{cases}$$

Let $f_i(x) = \langle w_i, x \rangle$, and define $c_i = \text{sign}(f_i)$, for $i = 0, \dots, n$. A short calculation shows that the margin $\gamma(c_i) = \min_j \frac{\langle w_i, x_j \rangle}{\|w_i\|_2} \geq 1/8$, and therefore, $c_i \in H$ for all $0 \leq i \leq n$. Another calculation shows that $c_i(x_j) = +1$ if and only if $i = j$, and that $|f_i(x_j)| = \frac{1/2 + j/10n^2}{\sqrt{2}}$. In particular note that $|f_i(x_j)|$ does not depend on i .

Consider the set $X_n \setminus \{x_i\}$. The previous paragraph implies that there is no comparison nor label query that distinguishes $c_0 = \text{sign}(f_0)$ from $c_i = \text{sign}(f_i)$ on this set. This means that it is **not** the case that $X_n \setminus \{x_i\} \xrightarrow{f_0} x_i$, which finishes the proof of Theorem 4.17. \square

5 Acknowledgements

This work has benefited from various discussions during the special program on *Foundations of Machine Learning* that took place at the *Simons Institute for the Theory of Computing*, in Berkeley. In particular, the authors would like to thank Sanjoy Dasgupta for inspiring the focus on comparison queries.

References

- S. Agarwal, J. Wills, L. Cayton, G. R. G. Lanckriet, D. J. Kriegman, and S. J. Belongie. Generalized non-metric multidimensional scaling. In M. Meila and X. Shen, editors, *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics, AISTATS 2007, San Juan, Puerto Rico, March 21-24, 2007*, volume 2 of *JMLR Proceedings*, pages 11–18. JMLR.org, 2007. URL <http://www.jmlr.org/proceedings/papers/v2/agarwal07a.html>.
- D. Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1987. doi: 10.1007/BF00116828. URL <http://dx.doi.org/10.1007/BF00116828>.
- H. Ashtiani, S. Kushagra, and S. Ben-David. Clustering with same-cluster queries. In D. D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 3216–3224, 2016. URL <http://papers.nips.cc/paper/6449-clustering-with-same-cluster-queries>.
- P. Awasthi, V. Feldman, and V. Kanade. Learning using local membership queries. In Shalev-Shwartz and Steinwart (2013), pages 398–431. URL <http://jmlr.org/proceedings/papers/v30/Awasthi13.html>.
- M. Balcan and S. Hanneke. Robust interactive learning. *CoRR*, abs/1111.1422, 2011. URL <http://arxiv.org/abs/1111.1422>.
- M. Balcan and P. M. Long. Active and passive learning of linear separators under log-concave distributions. In Shalev-Shwartz and Steinwart (2013), pages 288–316. URL <http://jmlr.org/proceedings/papers/v30/Balcan13.html>.
- M. Balcan, A. Z. Broder, and T. Zhang. Margin based active learning. In N. H. Bshouty and C. Gentile, editors, *Learning Theory, 20th Annual Conference on Learning Theory, COLT 2007, San Diego, CA, USA, June 13-15, 2007, Proceedings*, volume 4539 of *Lecture Notes in Computer Science*, pages 35–50. Springer, 2007. ISBN 978-3-540-72925-9. doi: 10.1007/978-3-540-72927-3_5. URL http://dx.doi.org/10.1007/978-3-540-72927-3_5.
- M. Balcan, A. Beygelzimer, and J. Langford. Agnostic active learning. *J. Comput. Syst. Sci.*, 75(1):78–89, 2009. doi: 10.1016/j.jcss.2008.07.003. URL <http://dx.doi.org/10.1016/j.jcss.2008.07.003>.
- M. Balcan, S. Hanneke, and J. W. Vaughan. The true sample complexity of active learning. *Machine Learning*, 80(2-3):111–139, 2010. doi: 10.1007/s10994-010-5174-y. URL <http://dx.doi.org/10.1007/s10994-010-5174-y>.
- G. Bary-Weisberg, A. Daniely, and S. Shalev-Shwartz. Distribution free learning with local queries. *CoRR*, abs/1603.03714, 2016. URL <http://arxiv.org/abs/1603.03714>.
- E. B. Baum. Neural net algorithms that learn in polynomial time from examples and queries. *IEEE Trans. Neural Networks*, 2(1):5–19, 1991. doi: 10.1109/72.80287. URL <http://dx.doi.org/10.1109/72.80287>.
- C. Berlind and R. Urner. Active nearest neighbors in changing environments. In F. R. Bach and D. M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 1870–1879. JMLR.org, 2015. URL <http://jmlr.org/proceedings/papers/v37/berlind15.html>.

- A. Beygelzimer, D. J. Hsu, J. Langford, and T. Zhang. Agnostic active learning without constraints. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23: 24th Annual Conference on Neural Information Processing Systems 2010. Proceedings of a meeting held 6-9 December 2010, Vancouver, British Columbia, Canada.*, pages 199–207. Curran Associates, Inc., 2010. URL <http://papers.nips.cc/paper/4014-agnostic-active-learning-without-constraints>.
- A. Blum, P. Chalasani, S. A. Goldman, and D. K. Slonim. Learning with unreliable boundary queries. *J. Comput. Syst. Sci.*, 56(2):209–222, 1998. doi: 10.1006/jcss.1997.1559. URL <http://dx.doi.org/10.1006/jcss.1997.1559>.
- R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952. ISSN 00063444. URL <http://www.jstor.org/stable/2334029>.
- N. H. Bshouty, J. C. Jackson, and C. Tamon. More efficient pac-learning of DNF with membership queries under the uniform distribution. *J. Comput. Syst. Sci.*, 68(1):205–234, 2004. doi: 10.1016/j.jcss.2003.10.002. URL <http://dx.doi.org/10.1016/j.jcss.2003.10.002>.
- J. Cardinal, J. Iacono, and A. Ooms. Solving k-sum using few linear queries. In P. Sankowski and C. D. Zaroliagis, editors, *24th Annual European Symposium on Algorithms, ESA 2016, August 22-24, 2016, Aarhus, Denmark*, volume 57 of *LIPIcs*, pages 25:1–25:17. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016. ISBN 978-3-95977-015-6. doi: 10.4230/LIPIcs.ESA.2016.25. URL <http://dx.doi.org/10.4230/LIPIcs.ESA.2016.25>.
- L. Chen, S. H. Hassani, and A. Karbasi. Dimension coupling: Optimal active learning of halfspaces via query synthesis. *CoRR*, abs/1603.03515, 2016. URL <http://arxiv.org/abs/1603.03515>.
- S. Dasgupta. Analysis of a greedy active learning strategy. In *Advances in Neural Information Processing Systems 17 [Neural Information Processing Systems, NIPS 2004, December 13-18, 2004, Vancouver, British Columbia, Canada]*, pages 337–344, 2004. URL <http://papers.nips.cc/paper/2636-analysis-of-a-greedy-active-learning-strategy>.
- S. Dasgupta, D. J. Hsu, and C. Monteleoni. A general agnostic active learning algorithm. In *International Symposium on Artificial Intelligence and Mathematics, ISAIM 2008, Fort Lauderdale, Florida, USA, January 2-4, 2008*, 2008. URL <http://isaim2008.unl.edu/PAPERS/SS3-ActiveLearning/isaim08-sanjoy-dasgupta.pdf>.
- R. El-Yaniv and Y. Wiener. Active learning via perfect selective classification. *Journal of Machine Learning Research*, 13:255–279, 2012. URL <http://dl.acm.org/citation.cfm?id=2188394>.
- E. Ezra and M. Sharir. The decision tree complexity for k-sum is at most nearly quadratic. *CoRR*, abs/1607.04336, 2016. URL <http://arxiv.org/abs/1607.04336>.
- V. Feldman. On the power of membership queries in agnostic learning. *Journal of Machine Learning Research*, 10:163–182, 2009a. doi: 10.1145/1577069.1577076. URL <http://doi.acm.org/10.1145/1577069.1577076>.
- V. Feldman. Hardness of approximate two-level logic minimization and PAC learning with membership queries. *J. Comput. Syst. Sci.*, 75(1):13–26, 2009b. doi: 10.1016/j.jcss.2008.07.007. URL <http://dx.doi.org/10.1016/j.jcss.2008.07.007>.

- M. L. Fredman. How good is the information theory bound in sorting? *Theor. Comput. Sci.*, 1(4):355–361, 1976. doi: 10.1016/0304-3975(76)90078-5. URL [http://dx.doi.org/10.1016/0304-3975\(76\)90078-5](http://dx.doi.org/10.1016/0304-3975(76)90078-5).
- Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors. *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, 2014. URL <http://papers.nips.cc/book/advances-in-neural-information-processing-systems-27-2014>.
- A. Gonen, S. Sabato, and S. Shalev-Shwartz. Efficient active learning of halfspaces: an aggressive approach. *Journal of Machine Learning Research*, 14(1):2583–2615, 2013. URL <http://dl.acm.org/citation.cfm?id=2567744>.
- I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio. Generative adversarial nets. In Ghahramani et al. (2014), pages 2672–2680. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets>.
- E. Goto and H. Takahasi. Some theorems useful in threshold logic for enumerating boolean functions. In *IFIP Congress*, pages 747–752, 1962.
- S. Hanneke. A bound on the label complexity of agnostic active learning. In Z. Ghahramani, editor, *Machine Learning, Proceedings of the Twenty-Fourth International Conference (ICML 2007), Corvallis, Oregon, USA, June 20-24, 2007*, volume 227 of *ACM International Conference Proceeding Series*, pages 353–360. ACM, 2007. ISBN 978-1-59593-793-3. doi: 10.1145/1273496.1273541. URL <http://doi.acm.org/10.1145/1273496.1273541>.
- S. Hanneke. The optimal sample complexity of PAC learning. *CoRR*, abs/1507.00473, 2015. URL <http://arxiv.org/abs/1507.00473>.
- S. Hanneke and L. Yang. Surrogate losses in passive and active learning. *CoRR*, abs/1207.3772, 2012. URL <http://arxiv.org/abs/1207.3772>.
- S. Hanneke and L. Yang. Minimax analysis of active learning. *Journal of Machine Learning Research*, 16:3487–3602, 2015. URL <http://jmlr.org/papers/v16/hanneke15a.html>.
- K. Huang, Y. Ying, and C. Campbell. Generalized sparse metric learning with relative comparisons. *Knowl. Inf. Syst.*, 28(1):25–45, 2011. doi: 10.1007/s10115-010-0313-0. URL <http://dx.doi.org/10.1007/s10115-010-0313-0>.
- J. C. Jackson. An efficient membership-query algorithm for learning {DNF} with respect to the uniform distribution. *Journal of Computer and System Sciences*, 55(3):414 – 440, 1997. ISSN 0022-0000. doi: <http://dx.doi.org/10.1006/jcss.1997.1533>. URL [//www.sciencedirect.com/science/article/pii/S0022000097915336](http://www.sciencedirect.com/science/article/pii/S0022000097915336).
- K. G. Jamieson and R. D. Nowak. Active ranking using pairwise comparisons. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. C. N. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain.*, pages 2240–2248, 2011. URL <http://papers.nips.cc/paper/4427-active-ranking-using-pairwise-comparisons>.
- N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica*, 4(4): 373–396, 1984. doi: 10.1007/BF02579150. URL <http://dx.doi.org/10.1007/BF02579150>.

- V. Koltchinskii. Rademacher complexities and bounding the excess risk in active learning. *Journal of Machine Learning Research*, 11:2457–2485, 2010. URL <http://portal.acm.org/citation.cfm?id=1953014>.
- S. Kwek and L. Pitt. PAC learning intersections of halfspaces with membership queries. *Algorithmica*, 22(1/2):53–75, 1998. doi: 10.1007/PL00013834. URL <http://dx.doi.org/10.1007/PL00013834>.
- K. Lang and E. Baum. Query learning can work poorly when a human oracle is used. In *Proceedings of the IEEE International Joint Conference on Neural Networks*, pages 335–340. IEEE Press, 1992. *Proceedings of the IEEE International Joint Conference on Neural Networks*, pages 335–340, 1992.
- D. Liu. A note on point location in arrangements of hyperplanes. *Inf. Process. Lett.*, 90(2):93–95, 2004. doi: 10.1016/j.ipl.2004.01.010. URL <http://dx.doi.org/10.1016/j.ipl.2004.01.010>.
- A. McCallum and K. Nigam. Employing EM and pool-based active learning for text classification. In J. W. Shavlik, editor, *Proceedings of the Fifteenth International Conference on Machine Learning (ICML 1998), Madison, Wisconsin, USA, July 24-27, 1998*, pages 350–358. Morgan Kaufmann, 1998. ISBN 1-55860-556-8.
- B. McFee and G. R. G. Lanckriet. Learning similarity in heterogeneous data. In J. Z. Wang, N. Boujemaa, N. O. Ramirez, and A. Natsev, editors, *Proceedings of the 11th ACM SIGMM International Conference on Multimedia Information Retrieval, MIR 2010, Philadelphia, Pennsylvania, USA, March 29-31, 2010*, pages 243–244. ACM, 2010. ISBN 978-1-60558-815-5. doi: 10.1145/1743384.1743427. URL <http://doi.acm.org/10.1145/1743384.1743427>.
- S. Meiser. Point location in arrangements of hyperplanes. *Inf. Comput.*, 106(2):286–303, 1993. doi: 10.1006/inco.1993.1057. URL <http://dx.doi.org/10.1006/inco.1993.1057>.
- F. Meyer auf der Heide. A polynomial linear search algorithm for the n-dimensional knapsack problem. *J. ACM*, 31(3):668–676, 1984. doi: 10.1145/828.322450. URL <http://doi.acm.org/10.1145/828.322450>.
- R. D. Nowak. The geometry of generalized binary search. *IEEE Trans. Information Theory*, 57(12):7893–7906, 2011. doi: 10.1109/TIT.2011.2169298. URL <http://dx.doi.org/10.1109/TIT.2011.2169298>.
- B. Qian, X. Wang, F. Wang, H. Li, J. Ye, and I. Davidson. Active learning from relative queries. In F. Rossi, editor, *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, pages 1614–1620. IJCAI/AAAI, 2013. ISBN 978-1-57735-633-2. URL <http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6320>.
- M. Schultz and T. Joachims. Learning a distance metric from relative comparisons. In S. Thrun, L. K. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems 16 [Neural Information Processing Systems, NIPS 2003, December 8-13, 2003, Vancouver and Whistler, British Columbia, Canada]*, pages 41–48. MIT Press, 2003. ISBN 0-262-20152-6. URL <http://papers.nips.cc/paper/2366-learning-a-distance-metric-from-relative-comparisons>.
- S. Shalev-Shwartz and I. Steinwart, editors. *COLT 2013 - The 26th Annual Conference on Learning Theory, June 12-14, 2013, Princeton University, NJ, USA*, volume 30 of *JMLR Workshop and Conference Proceedings*, 2013. JMLR.org. URL <http://jmlr.org/proceedings/papers/v30/>.

- R. H. Sloan and G. Turán. Learning from incomplete boundary queries using split graphs and hypergraphs. In S. Ben-David, editor, *Computational Learning Theory, Third European Conference, EuroCOLT '97, Jerusalem, Israel, March 17-19, 1997, Proceedings*, volume 1208 of *Lecture Notes in Computer Science*, pages 38–50. Springer, 1997. ISBN 3-540-62685-9. doi: 10.1007/3-540-62685-9_5. URL http://dx.doi.org/10.1007/3-540-62685-9_5.
- O. Tamuz, C. Liu, S. J. Belongie, O. Shamir, and A. Kalai. Adaptively learning the crowd kernel. In L. Getoor and T. Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning, ICML 2011, Bellevue, Washington, USA, June 28 - July 2, 2011*, pages 673–680. Omnipress, 2011.
- G. Turán. Lower bounds for PAC learning with queries. In L. Pitt, editor, *Proceedings of the Sixth Annual ACM Conference on Computational Learning Theory, COLT 1993, Santa Cruz, CA, USA, July 26-28, 1993.*, pages 384–391. ACM, 1993. ISBN 0-89791-611-5. doi: 10.1145/168304.168382. URL <http://doi.acm.org/10.1145/168304.168382>.
- R. Urner, S. Wulff, and S. Ben-David. PLAL: cluster-based active learning. In Shalev-Shwartz and Steinwart (2013), pages 376–397. URL <http://jmlr.org/proceedings/papers/v30/Urner13.html>.
- P. M. Vaidya. A new algorithm for minimizing convex functions over convex sets. *Math. Program.*, 73:291–341, 1996. doi: 10.1007/BF02592216. URL <http://dx.doi.org/10.1007/BF02592216>.
- S. Vikram and S. Dasgupta. Interactive bayesian hierarchical clustering. In M. Balcan and K. Q. Weinberger, editors, *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2081–2090. JMLR.org, 2016. URL <http://jmlr.org/proceedings/papers/v48/vikram16.html>.
- F. L. Wauthier, N. Jojic, and M. I. Jordan. Active spectral clustering via iterative uncertainty reduction. In Q. Yang, D. Agarwal, and J. Pei, editors, *The 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, Beijing, China, August 12-16, 2012*, pages 1339–1347. ACM, 2012. ISBN 978-1-4503-1462-6. doi: 10.1145/2339530.2339737. URL <http://doi.acm.org/10.1145/2339530.2339737>.
- Y. Wiener, S. Hanneke, and R. El-Yaniv. A compression technique for analyzing disagreement-based active learning. *Journal of Machine Learning Research*, 16:713–745, 2015. URL <http://dl.acm.org/citation.cfm?id=2831137>.
- E. P. Xing, A. Y. Ng, M. I. Jordan, and S. J. Russell. Distance metric learning with application to clustering with side-information. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15 [Neural Information Processing Systems, NIPS 2002, December 9-14, 2002, Vancouver, British Columbia, Canada]*, pages 505–512. MIT Press, 2002. ISBN 0-262-02550-7. URL <http://papers.nips.cc/paper/2164-distance-metric-learning-with-application-to-clustering-with-side-information>.
- C. Zhang and K. Chaudhuri. Beyond disagreement-based agnostic active learning. In Ghahramani et al. (2014), pages 442–450. URL <http://papers.nips.cc/paper/5435-beyond-disagreement-based-agnostic-active-learning>.