# Distribution Free Learning with Local Queries* 

Galit Bary-Weisberg ${ }^{\dagger} \quad$ Amit Daniely ${ }^{\ddagger} \quad$ Shai Shalev-Shwartz ${ }^{\S}$

March 8, 2016


#### Abstract

The model of learning with local membership queries interpolates between the PAC model and the membership queries model by allowing the learner to query the label of any example that is similar to an example in the training set. This model, recently proposed and studied by Awasthi et al. [5], aims to facilitate practical use of membership queries.

We continue this line of work, proving both positive and negative results in the distribution free setting. We restrict to the boolean cube $\{-1,1\}^{n}$, and say that a query is $q$-local if it is of a hamming distance $\leq q$ from some training example. On the positive side, we show that 1-local queries already give an additional strength, and allow to learn a certain type of DNF formulas. On the negative side, we show that even $\left(n^{0.99}\right)$-local queries cannot help to learn various classes including Automata, DNFs and more. Likewise, $q$-local queries for any constant $q$ cannot help to learn Juntas, Decision Trees, Sparse Polynomials and more. Moreover, for these classes, an algorithm that uses $\left(\log ^{0.99}(n)\right)$-local queries would lead to a breakthrough in the best known running times.


[^0]
## 1 Introduction

A child typically learns to recognize a cat based on two types of input. The first is given by her parents, pointing at a cat and saying "Look, a cat!". The second is given as a response to the child's frequent question "What is that?". These two types of input were the basis for the learning model originally suggested by Valiant [21]. Indeed, in Valiant's model, the learner can randomly sample labelled examples from "nature", but it can also make a membership query ( $M Q$ ) for the label of any unseen example. Today, the acronym PAC stands for the restricted model in which MQ are forbidden, while the full model is called PAC+MQ. Much work has been done investigating the limits and the strengths of MQ. In particular, membership queries were proven stronger than the vanilla PAC model $[1,10,12,16]$. Yet, MQ are rarely used in practice. This is commonly attributed to the fact that MQ algorithms query very artificial examples, that are uninterpretable by humans (e.g., [8]).

Awasthi et al. [5] suggested a solution to the problem of unnatural examples. They considered a mid-way model that allows algorithms to make only local queries, i.e., query examples that are close to examples from the sample set. Hopefully, examples which are similar to natural examples will appear natural to humans. Awasti et al. considered the case where the instance space is $\{-1,1\}^{n}$ and the distance between examples is the Hamming distance. They proved positive results on learning sparse polynomials with $O(\log (n))$-local queries under what they defined as locally smooth distributions ${ }^{1}$. They also proposed an algorithm that learns DNF formulas under the uniform distribution in quasi-polynomial time using $O(\log (n))$-local queries.

Our work follows Awasthi et al. and investigates local queries in the distribution free setting, in which no explicit assumptions are made on the underlying distribution, but only on the learned hypothesis. We prove both positive and negative results in this context:

- One of the strongest and most beautiful results in the MQ model shows that automata are learnable with membership queries [1]. We show that unfortunately, this is probably not the case with local queries. Concretely, we show that even $\left(n^{0.99}\right)$-local queries cannot help to learn automata. Namely, such an algorithm will imply a standard PAC algorithm for automata. As learning automata is hard under several assumptions [17, 13], our result suggests that it is hard to learn automata even with $\left(n^{0.99}\right)$-local queries.
- We prove a similar result for several additional classes. Namely, we show that $\left(n^{0.99}\right)$ local queries cannot help to learn DNFs, intersection of halfspaces, decision lists, depth$d$ circuits for any $d \geq 2$, and depth- $d$ threshold circuits for any $d \geq 2$. Likewise, for any constant $q, q$-local queries cannot help to learn Juntas, Decision Trees, Sparse polynomials, and Sparse polynomial threshold functions. In fact, we show that even $\left(\log ^{0.99}(n)\right)$-local queries are unlikely to lead to polynomial time algorithms. Namely, any algorithm that uses $\left(\log ^{0.99}(n)\right)$-local queries will result with a PAC algorithm whose running time significantly improves the state of the art in these well studied problems.
- On the positive side we show that already 1-local queries are probably stronger than the vanilla PAC model. Concretely, we show that a certain simplistic but natural

[^1]learning problem, which we term learning DNFs with evident examples, is learnable with 1-local queries. Furthermore, we show that without queries, this problem is at least as hard as learning decision trees, that are conjectured to be hard to learn.

## 2 Previous Work

Membership Queries Several concept classes are known to be learnable only if membership queries are allowed: Deterministic Finite Automatons [1], k-term DNF for $k=$ $\frac{\log (n)}{\log (\log (n))}$ [10], decision trees and k -almost monotone-DNF formulas [12], intersections of khalfspaces [7] and DNF formulas under the uniform distribution [16]. The last result builds on Freund's boosting algorithm [15] and the Fourier-based technique for learning using membership queries due to [18]. We note that there are cases in which MQ do not help. E.g., in the case of learning DNF and CNF formulas [2], assuming that one way functions exist, and in the case of distribution free agnostic learning [14].

Local Membership Queries Awasthi et al. focused on learning with $O(\log (n))$-local queries. They showed that $t$-sparse polynomials are learnable under locally smooth distributions using $O(\log (n)+\log (t))$-local queries, and that DNF formulas are learnable under the uniform distribution in quasi-polynomial time $\left.\left(n^{O(\log \log n}\right)\right)$ using $O(\log (n))$-local queries. They also presented some results regarding the strength of local MQ. They proved that under standard cryptographic assumptions, $(r+1)$-local queries are more powerful than $r$-local queries (for every $1 \leq r \leq n-1$ ). They also showed that local queries do not always help. They showed that if a concept class is agnostically learnable under the uniform distribution using $k$-local queries (for constant $k$ ) then it is also agnostically learnable (under the uniform distribution) in the PAC model.

We remark that besides local queries, there were additional suggestions to solve the problem of unnatural examples. For example, to allow "I don't know" answers, or to be tolerant to some incorrect answers $[3,4,11,20,9]$.

## 3 Setting

We consider binary classification where the instance space is $\mathcal{X}=\mathcal{X}_{n}=\{-1,1\}^{n}$ and the label space is $\mathcal{Y}=\{0,1\}$. A learning problem is defined by a hypothesis class $\mathcal{H} \subset\{0,1\}^{\mathcal{X}}$. The learner receives a training set

$$
S=\left\{\left(\mathbf{x}_{1}, h^{\star}\left(\mathbf{x}_{1}\right)\right),\left(\mathbf{x}_{2}, h^{\star}\left(\mathbf{x}_{2}\right)\right), \ldots,\left(\mathbf{x}_{m}, h^{\star}\left(\mathbf{x}_{m}\right)\right)\right\} \in(\mathcal{X} \times \mathcal{Y})^{m}
$$

where the $\mathbf{x}_{i}$ 's are sampled i.i.d. from some unknown distribution $\mathcal{D}$ on $\mathcal{X}$ and $h^{\star}: \mathcal{X} \rightarrow \mathcal{Y}$ is some unknown hypothesis. The learner is also allowed to make membership queries. Namely, to call an oracle, which receives as input some $\mathbf{x} \in \mathcal{X}$ and returns $h^{\star}(\mathbf{x})$. We say that a membership query for $\mathbf{x} \in \mathcal{X}$ is $q$-local if there exists a training example $\mathbf{x}_{i}$ whose Hamming distance from $\mathbf{x}$ is at most $q$.

The learner returns (a description of) a hypothesis $\hat{h}: \mathcal{X} \rightarrow \mathcal{Y}$. The goal is to approximate $h^{\star}$, namely to find $\hat{h}: \mathcal{X} \rightarrow \mathcal{Y}$ with loss as small as possible, where the loss is defined as
$L_{\mathcal{D}, h^{\star}}(\hat{h})=\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(\hat{h}(\mathbf{x}) \neq h^{\star}(\mathbf{x})\right)$. We will focus on the so-called realizable case where $h^{\star}$ is assumed to be in $\mathcal{H}$, and will require algorithms to return a hypothesis with loss $<\epsilon$ in time that is polynomial in $n$ and $\frac{1}{\epsilon}$. Concretely,

Definition 3.1 (Membership-Query Learning Algorithm) We say that a learning algorithm $\mathcal{A}$ learns $\mathcal{H}$ with $q$-local membership queries if

- There exists a function $m_{\mathcal{A}}(n, \epsilon) \leq$ poly $\left(n, \frac{1}{\epsilon}\right)$, such that for every distribution $\mathcal{D}$ over $\mathcal{X}$, every $h^{\star} \in \mathcal{H}$ and every $\epsilon>0$, if $\mathcal{A}$ is given access to q-local membership queries, and a training sequence

$$
S=\left\{\left(\mathbf{x}_{1}, h^{\star}\left(\mathbf{x}_{1}\right)\right),\left(\mathbf{x}_{2}, h^{\star}\left(\mathbf{x}_{2}\right)\right), \ldots,\left(\mathbf{x}_{m}, h^{\star}\left(\mathbf{x}_{m}\right)\right)\right\}
$$

where the $\mathbf{x}_{i}$ 's are sampled i.i.d. from $\mathcal{D}$ and $m \geq m_{\mathcal{A}}(n, \epsilon)$, then with probability of at least $t^{2} \frac{3}{4}$ (over the choice of $S$ ), the output $\hat{h}$ of $\mathcal{A}$ satisfies $L_{\mathcal{D}, h^{\star}}(\hat{h})<\epsilon$.

- Given a training set of size $m$
- $\mathcal{A}$ asks at most $\operatorname{poly}(m, n)$ membership queries.
- $\mathcal{A}$ runs in time $\operatorname{poly}(m, n)$.
- The hypothesis returned by $\mathcal{A}$ can be evaluated in time $\operatorname{poly}(m, n)$.

We remark that learning with 0-local queries is equivalent to PAC learning, while learning with $n$-local queries is equivalent to $\mathrm{PAC}+\mathrm{MQ}$ learning.

## 4 Learning DNFs with Evident Examples

Intuitively, when evaluating a DNF formula on a given example, we check a few conditions corresponding to the formula's terms, and deem the example positive if one of them holds. We will consider the case that for each of these conditions, there is a chance to see a "prototype example", that satisfies it in a strong or evident way. In the sequel, we denote by $h_{F}$ : $\{-1,1\}^{n} \rightarrow\{0,1\}$ the function induced by a DNF formula $F$ over $n$ variables.

Definition 4.1 Let $F=T_{1} \vee T_{2} \vee \ldots \vee T_{d}$ be a DNF formula. We say that an example $\mathbf{x} \in\{-1,1\}^{n}$ satisfies a term $T_{i}$ (with respect to the formula $F$ ) evidently and denote $T_{i}(\mathbf{x}) \equiv 1$ if :

- It satisfies $T_{i}$. (In particular, $h_{F}(\mathbf{x})=1$ )
- It does not satisfy any other term $T_{k}($ for $k \neq i)$ from $F$.
- No coordinate change will turn $T_{i}$ False and another term $T_{k}$ True. Concretely, if for $j \in[n]$ we denote $\mathbf{x}^{\oplus j}=\left(x_{1}, \ldots, x_{j-1},-x_{j}, x_{j+1}, \ldots, x_{n}\right)$, then for every coordinate $j \in[n]$, if $\mathbf{x}^{\oplus j}$ satisfies $F$ (i.e. if $h_{F}\left(\mathbf{x}^{\oplus j}\right)=1$ ) then $\mathbf{x}^{\oplus j}$ satisfies $T_{i}$ and only $T_{i}$.

[^2]Definition 4.2 Let $F=T_{1} \vee T_{2} \vee \ldots \vee T_{d}$ be a DNF formula. We say that $h^{\star}:\{-1,1\}^{n} \rightarrow$ $\{0,1\}$ is realized by $F$ with evident examples w.r.t. a distribution $\mathcal{D}$ if $h^{\star}=h_{F}$ and for every term ${ }^{3} T_{i}, \operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(T_{i}(\mathbf{x}) \equiv 1 \mid T_{i}(\mathbf{x})=1\right) \geq \frac{1}{n}$.

For example, the definition holds whenever $h^{\star}$ is realized by a DNF in which any pair of different terms contains two opposite literals. Also, functions that are realized by a decision tree can be also realized by a DNF in which every positive example satisfies a single term, corresponding to root-to-leaf path that ends with a leaf labelled by 1 . Hence, the assumption holds for decision trees provided that for every such path, there is a chance to see an example satisfying it evidently, in the sense that flipping each variable in the path will turn the example negative.

### 4.1 An algorithm

```
Algorithm 1
Input: \(\quad S_{1}, S_{2} \in\left(\{-1,1\}^{n} \times\{0,1\}\right)^{m}\)
Output: A DNF formula \(H\)
    Start with an empty DNF formula \(H\)
    for all positive examples \((\mathbf{x}, y) \in S_{1}\) do
        Define \(T=x_{1} \wedge \overline{x_{1}} \wedge x_{2} \wedge \overline{x_{2}} \wedge \ldots \wedge x_{n} \wedge \overline{x_{n}}\)
        for \(1 \leq j \leq n\) do
            Query \(\mathbf{x}^{\oplus j}\) to get \(h^{\star}\left(\mathbf{x}^{\oplus j}\right)\)
            if \(h^{\star}\left(\mathbf{x}^{\oplus j}\right)=1\) then
                Remove \(x_{j}\) and \(\overline{x_{j}}\) from \(T\)
            if \(h^{\star}\left(\mathbf{x}^{\oplus j}\right)=0\) then
                if \(x_{j}=1\) then
                    Remove \(\overline{x_{j}}\) from \(T\)
                if \(x_{j}=0\) then
                    Remove \(x_{j}\) from \(T\)
        \(H=H \vee T\)
    for all \(T\) in \(H\) do
        if \(T(\mathbf{x})=1\) but \(y=0\) for some \((\mathbf{x}, y) \in S_{2}\) then
            Remove \(T\) from \(H\)
    Return H
```

Theorem 4.3 Algorithm 1 learns with 1-local-queries poly-sized DNFs with evident examples.

Idea The algorithm is based on the following claim that follows easily from definition 4.1.

[^3]Claim 1 Let $F=T_{1} \vee T_{2} \vee \ldots \vee T_{d}$ be a DNF formula over $\{-1,1\}^{n}$. For every $\mathbf{x} \in\{-1,1\}^{n}$ that satisfies a term $T_{i}$ evidently (with respect to $F$ ), and for every $j \in[n]$ it holds that:

$$
h_{F}\left(\mathbf{x}^{\oplus j}\right)=1 \Longleftrightarrow \text { the term } T_{i} \text { does not contain the variable } x_{j}
$$

By this claim, if $\mathbf{x}$ is an evident example for a certain term $T$, one can easily reconstruct $T$. Indeed, by flipping the value of each variable and checking if the label changes, one can infer which variables appear in $T$. Furthermore, the sign of these variables can be inferred from their sign in $\mathbf{x}$. Hence, after seeing an evident example for all terms, one can have a list of terms containing all terms in the DNF. This list might have terms that are not part of the DNF. Yet, such terms can be thrown away later by testing if they make wrong predictions.

Proof (of Theorem 4.3) We will show that algorithm 1 learns with 1-local-queries any function that is realized by a DNF of size $\leq n^{2}$ with evident examples. Adapting the proof to hold with $n^{c}$ instead of $n^{2}$, for any $c>0$, is straight-forward. First, it is easy to see that this algorithm is efficient. Now, fix a distribution $\mathcal{D}$ and let $h^{\star}:\{-1,1\}^{n} \rightarrow\{0,1\}$ be a hypothesis that is realized, w.r.t. $\mathcal{D}$, by a DNF formula $F=T_{1} \vee T_{2} \vee \ldots \vee T_{d}$, $d \leq n^{2}$ with evident examples. Let $\epsilon>0$, and suppose we run the algorithm on two indepent samples from $\mathcal{D}$, denoted $S_{1}=\left\{\left(\mathbf{x}_{i}, h^{\star}\left(\mathbf{x}_{i}\right)\right\}_{i=1}^{m_{1}}\right.$ and $S_{2}=\left\{\left(\mathbf{x}_{i}^{\prime}, h^{\star}\left(\mathbf{x}_{i}^{\prime}\right)\right\}_{i=1}^{m_{2}}\right.$. We will show that if $m_{1} \geq \frac{32 n^{3}}{\epsilon} \log \frac{32 n^{2}}{\epsilon} \geq \frac{32 n d}{\epsilon} \log \frac{32 d}{\epsilon}$ and $m_{2} \geq \frac{32 m_{1}}{\epsilon} \log \frac{32 m_{1}}{\epsilon}$ then with probability of at least $\frac{3}{4}$, the algorithm will return a function $\hat{h}$ with $L_{\mathcal{D}, h^{\star}}(\hat{h})<\epsilon$. Let $H$ be the DNF formula returned by the algorithm, and let $\hat{h}$ be the function induced by $H$. We have that

$$
\begin{aligned}
L_{\mathcal{D}, h^{\star}}(\hat{h}) & =\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(h^{\star}(\mathbf{x}) \neq \hat{h}(\mathbf{x})\right) \\
& =\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(h^{\star}(\mathbf{x})=1 \text { and } \hat{h}(\mathbf{x})=0\right)+\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(h^{\star}(\mathbf{x})=0 \text { and } \hat{h}(\mathbf{x})=1\right)
\end{aligned}
$$

The proof will now follow from claims 2 and 3 .
Claim 2 With probability at least $\frac{7}{8}$ over the choice of $S_{1}, S_{2}$ we have

$$
\begin{equation*}
\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(h^{\star}(\mathbf{x})=1 \text { and } \hat{h}(\mathbf{x})=0\right) \leq \frac{\epsilon}{2} \tag{1}
\end{equation*}
$$

Proof We first note that if there is an evident example $\mathbf{x}$ for a term $T_{i}$ in $S_{1}$, then $T_{i}$ will be in the output formula. Indeed, in the for-loop that go over the examples in $S_{1}$, when processing the example $\left(\mathbf{x}, h^{\star}(\mathbf{x})\right)$, it is not hard to see that $T_{i}$ will be added. We furthermore claim that the term $T_{i}$ won't be removed at the for loop that tests the terms collected in the first loop. Indeed, if for some $\left(\mathbf{x}, h^{\star}(\mathbf{x})\right) \in S_{2}$ we have $T_{i}(\mathbf{x})=1$, it must be the case that $h^{\star}(\mathbf{x})=1$. Now, we say that the term $T_{i}$ is revealed if we see an evident example for this term in $S_{1}$. We also denote $p_{i}=\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(T_{i}(\mathbf{x})=1\right)$. We have

$$
\begin{aligned}
\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(h^{\star}(\mathbf{x})=1 \text { and } \hat{h}(\mathbf{x})=0\right) & \leq \sum_{i=1}^{d} \operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(T_{i}(\mathbf{x})=1 \text { and } \hat{h}(\mathbf{x})=0\right) \\
& \leq \sum_{i: T_{i} \text { is not revealed }} p_{i}
\end{aligned}
$$

Now, by the assumption that $h^{\star}$ is realized with evident queries, the probability (over the choice of $S_{1}, S_{2}$ ) that $T_{i}$ is not revealed is at most $\left(1-\frac{p_{i}}{n}\right)^{m_{1}}$. Hence, if we denote by $A_{i}$ the event that $T_{i}$ is not revealed, we have

$$
\begin{aligned}
\underset{S_{1} \sim \mathcal{D}^{m_{1}}}{\mathbb{E}}\left[\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(h^{\star}(\mathbf{x})=1 \text { and } \hat{h}(\mathbf{x})=0\right)\right] & \leq \mathbb{E}_{S_{1} \sim \mathcal{D}^{m_{1}}}\left[\sum_{i=1}^{d} p_{i} \cdot 1_{A_{i}}\right] \\
& =\sum_{i=1}^{d} p_{i} \mathbb{E}_{S_{1} \sim \mathcal{D}^{m_{1}}}\left[1_{A_{i}}\right] \\
& =\sum_{i=1}^{d} p_{i} \operatorname{Pr}_{S_{1} \sim \mathcal{D}^{m_{1}}}\left[A_{i}\right] \\
& =\sum_{i=1}^{d} p_{i}\left(1-\frac{p_{i}}{n}\right)^{m_{1}} \\
& =\sum_{i \left\lvert\, p_{i}<\frac{\epsilon}{32 d}\right.} p_{i}\left(1-\frac{p_{i}}{n}\right)^{m_{1}}+\sum_{i \left\lvert\, p_{i} \geq \frac{\epsilon}{32 d}\right.} p_{i}\left(1-\frac{p_{i}}{n}\right)^{m_{1}} \\
& \leq \sum_{i \left\lvert\, p_{i}<\frac{\epsilon}{32 d}\right.} \frac{\epsilon}{32 d}+\sum_{i \left\lvert\, p_{i} \geq \frac{\epsilon}{32 d}\right.}\left(1-\frac{p_{i}}{n}\right)^{m_{1}} \\
& \leq d \cdot \frac{\epsilon}{32 d}+\sum_{i \left\lvert\, p_{i} \geq \frac{\epsilon}{32 d}\right.} e^{-\frac{m_{1} p_{i}}{n}} \\
& \leq \frac{\epsilon}{32}+d \cdot e^{-\frac{m_{1} \epsilon}{32 d n}}
\end{aligned}
$$

Since $m_{1} \geq \frac{32 d n}{\epsilon} \log \frac{32 d}{\epsilon}$ the last expression is bounded by $\frac{\epsilon}{16}$. By Markov's inequality we conclude that the probabilty over the coice of $S_{1}, S_{2}$ that (1) does not lold is less than $\frac{1}{8}$.

Claim 3 With probability at least $\frac{7}{8}$ over the choice of $S_{1}, S_{2}$ we have

$$
\begin{equation*}
\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(h^{\star}(\mathbf{x})=0 \text { and } \hat{h}(\mathbf{x})=1\right) \leq \frac{\epsilon}{2} \tag{2}
\end{equation*}
$$

Proof Let $\hat{T}_{1}, \ldots, \hat{T}_{r}$ be the terms that were added to $H$ at the first for-loop. Denote $q_{i}=\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(\hat{T}_{i}(\mathbf{x})=1\right.$ and $\left.h^{\star}(\mathbf{x})=0\right)$. We have

$$
\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left(h^{\star}(\mathbf{x})=0 \text { and } \hat{h}(\mathbf{x})=1\right) \leq \sum_{i: \hat{T}_{i} \text { is not removed }} q_{i}
$$

Now, the probability that $\hat{T}_{i}$ is not removed is $\left(1-q_{i}\right)^{m_{2}}$. Hence, using an argument similar to the one used in the proof of claim 2 , and since $m_{2} \geq \frac{32 m_{1}}{\epsilon} \log \frac{32 m_{1}}{\epsilon} \geq \frac{32 r}{\epsilon} \log \frac{32 r}{\epsilon}$, the claim follows.

### 4.2 A matching Lower Bound

In this section we provide evidence that the use of queries in our algorithm is crucial. We will show that the problem of learning poly-sized decision trees can be reduced to the problem of learning DNFs with evident examples. As learning decision trees is conjectured to be intractable, this reduction serves as an indication that learning DNFs with strongly evident examples is hard without membership queries. In fact, we will show that learning decision trees can be even reduced to the case that all positive examples are evident.

Theorem 4.4 Learning poly-sized DNFs with evident examples is as hard as learning polysized decision trees.

We denote by $h_{T}$ the function induced by a decision tree $T$. The proof will use the following claim:

Claim 4 There exists a mapping (a reduction) $\varphi:\{-1,1\}^{n} \rightarrow\{-1,1\}^{2 n}$, that can be evaluated in poly $(n)$ time so that for every decision tree $\mathcal{T}$ over $\{-1,1\}^{n}$ there exists a DNF formula $F$ over $\{-1,1\}^{2 n}$ such that the following holds:

1. The number of terms in $F$ is upper bounded by the number of leaves in $\mathcal{T}$
2. $h_{\mathcal{T}}=h_{F} \circ \varphi$
3. $\forall \mathbf{x}$ such that $h_{\mathcal{T}}(\mathbf{x})=1, \varphi(\mathbf{x})$ satisfies some term in $F$ evidently.

Proof Define $\varphi$ as follows:

$$
\forall \mathbf{x}=\left(x_{1}, x_{2}, \ldots, x_{n}\right) \in \mathcal{X}_{n} \quad \varphi\left(x_{1}, x_{2}, \ldots, x_{n}\right)=\left(x_{1}, x_{1}, x_{2}, x_{2}, \ldots, x_{n}, x_{n}\right)
$$

Now, for every tree $\mathcal{T}$, we will build the desired DNF formula $F$ as follows: First we build a DNF formula $F^{\prime}$ over $\{-1,1\}^{n}$. Every leaf labeled ' 1 ' in $\mathcal{T}$ will define the following termtake the path from the root to that leaf and form the logical AND of the literals describing the path. $F^{\prime}$ will be a disjunction of these terms. Now, for every term $T$ in $F^{\prime}$ we will define a term $\phi(T)$ over $\mathcal{X}_{2 n}$ in the following way: Let $P_{T}=\left\{i \in[n]: x_{i}\right.$ appear in $\left.T\right\}$ and $N_{T}=\left\{i \in[n]: \overline{x_{i}}\right.$ appear in $\left.T\right\}$. So

$$
T=\bigwedge_{j \in P_{T}} x_{j} \bigwedge_{j \in N_{T}} \overline{x_{j}}
$$

Define

$$
\phi(T)=\bigwedge_{j \in P_{T}} x_{2 j-1} \bigwedge_{j \in P_{T}} x_{2 j} \bigwedge_{j \in N_{T}} \overline{x_{2 j-1}} \bigwedge_{j \in N_{T}} \overline{x_{2 j}}
$$

Finally, define $F$ to be the DNF formula over $\mathcal{X}_{2 n}$ given by

$$
F=\bigvee_{T \in F^{\prime}} \phi(T)
$$

We will now prove that $\varphi$ and $F$ satisfy the required conditions. First, $\varphi$ can be evaluated in linear time in $n$. Second, it is easy to see that $h_{\mathcal{T}}=h_{F} \circ \varphi$, and as every term in $F$ matches one of $\mathcal{T}$ 's leaves, the number of terms in $F$ cannot exceed the number of leaves in $\mathcal{T}$. It is left to show that the third requirement holds. Let there be an $\mathbf{x}$ such that $h_{\mathcal{T}}(\mathbf{x})=1$, then $\mathbf{x}$ is matched to one and only one path from $\mathcal{T}$ 's root to a leaf labeled '1'. From the construction of $F$, x satisfies one and only one term in $F^{\prime}$. Regarding the last requirement, that no coordinate change will make one term from $F$ False and another one True, we made sure this will not happen by "doubling" each variable. By this construction, in order to change a term from False to True at least two coordinate must change their value.
We are now ready to prove theorem 4.4.
Proof [of theorem 4.4] Suppose that $\mathcal{A}$ learns size- $n$ DNFs with evident examples. Using the reduction from claim 4 we will build an efficient algorithm $\mathcal{B}$ that learns size- $n$ decision trees. For every training set

$$
S=\left\{\left(\mathbf{x}_{1}, h^{\star}\left(\mathbf{x}_{1}\right)\right),\left(\mathbf{x}_{2}, h^{\star}\left(\mathbf{x}_{2}\right)\right), \ldots,\left(\mathbf{x}_{m}, h^{\star}\left(\mathbf{x}_{m}\right)\right)\right\} \in\left(\mathcal{X}_{n} \times\{0,1\}\right)^{m}
$$

we define

$$
\left.\varphi(S):=\left\{\left(\varphi\left(\mathbf{x}_{1}\right), h^{\star}\left(\mathbf{x}_{1}\right)\right),\left(\varphi\left(\mathbf{x}_{2}\right)\right), h^{\star}\left(\mathbf{x}_{2}\right)\right), \ldots,\left(\varphi\left(\mathbf{x}_{m}\right), h^{\star}\left(\mathbf{x}_{m}\right)\right)\right\} \in\left(\mathcal{X}_{2 n} \times\{0,1\}\right)^{m}
$$

The algorithm $\mathcal{B}$ will work as follows: Given a training set $S, \mathcal{B}$ will run $\mathcal{A}$ on $\varphi(S)$, and will return $\hat{h} \circ \varphi$, where $\hat{h}$ is the hypothesis returned by $\mathcal{A}$. Since $\varphi$ can be evaluated in poly $(n)$ time and $\mathcal{A}$ is efficient, $\mathcal{B}$ is also efficient. We will prove that $\mathcal{B}$ learns size- $n$ trees. Since $\mathcal{A}$ learns size- $n$ DNFs with evident examples, there exists a function $m_{\mathcal{A}}(n, \epsilon) \leq \operatorname{poly}\left(n, \frac{1}{\epsilon}\right)$, such that if $\mathcal{A}$ is given a training sequence

$$
S=\left\{\left(\mathbf{x}_{1}, h^{\star}\left(\mathbf{x}_{1}\right)\right),\left(\mathbf{x}_{2}, h^{\star}\left(\mathbf{x}_{2}\right)\right), \ldots,\left(\mathbf{x}_{m}, h^{\star}\left(\mathbf{x}_{m}\right)\right)\right\} \in\left(\mathcal{X}_{n} \times\{0,1\}\right)^{m}
$$

where the $\mathbf{x}_{i}$ 's are sampled i.i.d. from a distribution $\mathcal{D}, h^{\star}$ is realized by a poly-sized DNF with evident examples, and $m \geq m_{\mathcal{A}}(n, \epsilon)$, then with probability of at least $\frac{3}{4}$ (over the choice of $S$ ), the output $\hat{h}$ of $\mathcal{A}$ satisfies $L_{\mathcal{D}, h^{\star}}(\hat{h}) \leq \epsilon$. Let $\mathcal{D}$ be a distribution on $\mathcal{X}_{n}$ and let $h_{\mathcal{T}}$ be a hypothesis that can be realized by a tree with $\leq n$ leafs. Define a distribution $\tilde{\mathcal{D}}$ on $\mathcal{X}_{2 n}$ by,

$$
(\tilde{\mathcal{D}})(\mathbf{z})= \begin{cases}\mathcal{D}(\mathbf{x}) & \text { if } \exists \mathbf{x} \in \mathcal{X}_{n} \text { such that } \mathbf{z}=\varphi(\mathbf{x}) \\ 0 & \text { otherwise }\end{cases}
$$

Now, since $h_{\mathcal{T}}$ is realized by $\mathcal{T}$, from the conditions that $\varphi$ satisfies, we get that $h_{\mathcal{T}}=h \circ \varphi$, where $h$ is realized by a DNF of size $\leq n$ with evident examples w.r.t. $\tilde{\mathcal{D}}$. Now if $S \in$ $\left(\mathcal{X}_{n} \times\{0,1\}\right)^{m}$ is an i.i.d. sample with $m \geq m_{\mathcal{A}}(2 n, \epsilon)$ we have that with probability of at least $\frac{3}{4}$ it holds that

$$
\begin{aligned}
L_{\mathcal{D}, h_{\mathcal{T}}}(\mathcal{B}(S)) & =L_{\mathcal{D}, h_{\mathcal{T}}}(\hat{h} \circ \varphi) \\
& =\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}\left[h_{\mathcal{T}}(\mathbf{x}) \neq \hat{h} \circ \varphi(\mathbf{x})\right] \\
& =\operatorname{Pr}_{\mathbf{x} \sim \mathcal{D}}[h \circ \varphi(\mathbf{x}) \neq \hat{h} \circ \varphi(\mathbf{x})] \\
& =\operatorname{Pr}_{\mathbf{z} \sim \tilde{\mathcal{D}}}[h(\mathbf{z}) \neq \hat{h}(\mathbf{z})] \\
& =L_{\tilde{\mathcal{D}}, h}(\hat{h}) \\
& =L_{\tilde{\mathcal{D}}, h}(\mathcal{A}(\varphi(S)))<\epsilon
\end{aligned}
$$

## 5 Lower Bounds

We first present a general technique to prove lower bounds on learning with local queries. For $A \subset \mathcal{X}_{n}$ and $q>0$ denote

$$
B(A, q)=\left\{\mathbf{x} \in \mathcal{X}_{n} \mid \exists \mathbf{a} \in A, d(\mathbf{x}, \mathbf{a}) \leq q\right\}
$$

We say that a mapping $\varphi:\{-1,1\}^{n} \rightarrow\{-1,1\}^{n^{\prime}}$ is a $q$-reduction of type $A$ from a class $\mathcal{H}$ to a class $\mathcal{H}^{\prime}$ if the following holds:

1. $\varphi$ is efficiently computable.
2. For every $h \in \mathcal{H}$ there is $h^{\prime} \in \mathcal{H}^{\prime}$ such that
(a) $h=h^{\prime} \circ \varphi$
(b) The restriction of $h^{\prime}$ to $B\left(\varphi\left(\mathcal{X}_{n}\right), q\right) \backslash \varphi\left(\mathcal{X}_{n}\right)$ is the constant function 1.

We say that a mapping $\varphi:\{-1,1\}^{n} \rightarrow\{-1,1\}^{n^{\prime}}$ is a $q$-reduction of type $B$ if the same requirements hold, except that (2b) is replaced by the following condition: For every $\mathbf{z} \in$ $B\left(\varphi\left(\mathcal{X}_{n}\right), q\right)$ there is a unique $\mathbf{x} \in \mathcal{X}_{n}$ satisfying $d(\mathbf{z}, \varphi(\mathbf{x})) \leq q$, and furthermore, $h^{\prime}(\mathbf{z})=$ $h(\mathbf{x})$.

Lemma 5.1 Suppose that there is a q-reduction $\varphi$ from $\mathcal{H}$ to $\mathcal{H}^{\prime}$. Then, learning $\mathcal{H}^{\prime}$ with $q$-local queries is as hard as PAC learning $\mathcal{H}$.

Proof (sketch) Suppose that $\mathcal{A}^{\prime}$ learns $\mathcal{H}^{\prime}$ with $q$-local queries. We need to show that there is an algorithm that learns $\mathcal{H}$. Indeed, by an argument similar to the one in theorem 4.4, it is not hard to verify that the following algorithm learns $\mathcal{H}$. Given a sample $\left\{\left(\mathbf{x}_{1}, y_{1}\right), \ldots,\left(\mathbf{x}_{m}, y_{m}\right)\right\} \subset\{-1,1\}^{n} \times\{0,1\}$, run $\mathcal{A}^{\prime}$ on the sample $\left\{\left(\varphi\left(\mathbf{x}_{1}\right), y_{1}\right), \ldots,\left(\varphi\left(\mathbf{x}_{m}\right), y_{m}\right)\right\} \subset\{-1,1\}^{n^{\prime}} \times\{0,1\}$. Whenever $\mathcal{A}^{\prime}$ makes a $q$-local query for $\mathbf{z} \in\{-1,1\}^{n^{\prime}}$, respond 1 if $\varphi$ is of type A. If $\varphi$ is of type B , respond $y_{i}$, where $\mathbf{x}_{i}$ is the unique training sample satisfying $d\left(\mathbf{z}, \varphi\left(\mathbf{x}_{i}\right)\right) \leq q$. Finally, if $\mathcal{A}^{\prime}$ returned the hypothesis $h^{\prime}$, return $h^{\prime} \circ \varphi$.

We next use Lemma 5.1 to prove that for several classes, local queries are not useful. Namely, if the class is learnable with local queries then it is also learnable without queries. We will use the following terminology. We say that $q$-local queries cannot help to learn a class $\mathcal{H}$ if $\mathcal{H}$ is learnable if and only if it is learnable with $q$-local queries.

Corollary 5.2 For every $\epsilon_{0}>0$, $\left(n^{1-\epsilon_{0}}\right)$-local queries cannot help to learn poly-sized DNFs, intersection of halfspaces, decision lists, depth-d circuits for any $d=d(n) \geq 2$, and depth-d threshold circuits for any $d=d(n) \geq 2$.

Proof We will only prove the corollary for DNFs. The proofs for the remaining classes are similar. Also, for simplicity, we will assume that $\epsilon_{0}=\frac{1}{3}$. Consider the mapping $\varphi$ : $\{-1,1\}^{n} \rightarrow\{-1,1\}^{n^{3}}$ that replicates each coordinate $n^{2}$ times. To establish the corollary, we show that $\varphi$ is an $\left(\left(n^{\prime}\right)^{\frac{2}{3}}-1\right)$-reduction of type A from poly-sized DNF to poly-sized DNF.

Indeed, let $F=T_{1} \vee \ldots \vee T_{d}$ be a DNF formula on $n$ variables, consider the following formula on the $n^{3}$ variables $\left\{x_{i, j}\right\}_{1 \leq i \leq n, 1 \leq j \leq n^{2}}$. Let,

$$
\begin{aligned}
T_{t}^{\prime}\left(\left\{x_{i, j}\right\}_{i, j}\right) & =T_{t}\left(x_{1,1}, \ldots, x_{n, 1}\right) \\
G^{\prime}\left(\left\{x_{i, j}\right\}_{i, j}\right) & =\vee_{i=1}^{n} \vee_{j=1}^{n^{2}-1}\left(x_{i, j} \wedge \neg x_{i, j+1}\right) \vee\left(x_{i, j+1} \wedge \neg x_{i, j}\right) \\
F^{\prime} & =\left(T_{1}^{\prime} \vee \ldots \vee T_{d}^{\prime}\right) \vee G^{\prime}
\end{aligned}
$$

It is not hard to verify that $h_{F}=h_{F^{\prime}} \circ \varphi$. Moreover, if $\mathbf{x}=\left\{x_{i, j}\right\}_{i, j} \in B\left(\varphi\left(\mathcal{X}_{n}\right), n^{2}-1\right) \backslash \varphi\left(\mathcal{X}_{n}\right)$ then $x_{i, j} \neq x_{i, j+1}$ for some $i, j$, meaning that $h_{G}(x)=1$ and therefore also $h_{F}(x)=1$.

Corollary 5.3 For every $\epsilon_{0}>0$, $\left(n^{1-\epsilon_{0}}\right)$-local queries cannot help to learn poly-sized automata.

Proof Again, for simplicity, we assume that $\epsilon_{0}=\frac{1}{3}$, and consider the mapping $\varphi$ : $\{-1,1\}^{n} \rightarrow\{-1,1\}^{n^{3}}$ that replicates each coordinate $n^{2}$ times. To establish the corollary, we show that $\varphi$ is an $\left(\left(n^{\prime}\right)^{\frac{2}{3}}-1\right)$-reduction of type A from poly-sized automata to poly-sized automata.

Indeed, let $A$ be an automaton on $n$ variables. It is enough to show that there is an automaton $A^{\prime}$ on the $n^{3}$ variables $\left\{x_{i, j}\right\}_{1 \leq i \leq n, 1 \leq j \leq n^{2}}$ such that (i) the size of $A^{\prime}$ is polynomial, (ii) $A^{\prime}$ accepts any string in $B\left(\varphi\left(\mathcal{X}_{n}\right), n^{2}-1\right) \backslash \varphi\left(\mathcal{X}_{n}\right)$, and (iii) $A^{\prime}$ accepts $\varphi(x)$ if and only if $A$ accepts $x$. Now, by the product construction of automatons [19], if $A_{1}^{\prime}, A_{2}^{\prime}$ are automata that induce the functions $h_{A_{1}^{\prime}}, h_{A_{2}^{\prime}}:\{-1,1\}^{n^{3}} \rightarrow\{0,1\}$, then the function $h_{A_{1}^{\prime}} \vee h_{A_{2}^{\prime}}$ can be induced by an automaton of size $\left|A_{1}^{\prime}\right| \cdot\left|A_{2}^{\prime}\right|$. Hence, it is enough to show that there are poly-sized automata $A_{1}^{\prime}, A_{2}^{\prime}$ that satisfies (ii) and (iii) respectively.

A construction of a size $O\left(n^{2}\right)$ automaton that satisfies (ii) is a simple exercise. We next explain how to construct a poly-sized automaton $A_{2}^{\prime}$ satisfying (iii). The states of $A_{2}^{\prime}$ will be the $S(A) \times\left[n^{2}\right]$ (here, $S(A)$ denotes the set of states of $A$ ). The start state of $A_{2}^{\prime}$ will be ( $\alpha_{0}, 1$ ), where $\alpha_{0}$ is the start state of $A$, and the accept states of $A_{2}^{\prime}$ will be the cartesian product of the accept states of $A$ with $\left[n^{2}\right]$. Finally if $\delta: S(A) \times\{-1,1\} \rightarrow S(A)$ is the transition function of $A$, then the transition function of $A_{2}^{\prime}$ is defined by

$$
\delta^{\prime}((\alpha, i), b)= \begin{cases}(\alpha, i+1) & 1 \leq i<n^{2} \\ (\delta(\alpha, b), 1) & i=n^{2}\end{cases}
$$

It is not hard to verify that $A_{2}^{\prime}$ satisfies (iii).
In the next corollary, a Junta of size $t$ is a function $h:\{-1,1\}^{n} \rightarrow\{0,1\}$ that depends on $\leq \log (t)$ variables. The rational behind this definition of size, is the fact that in order to describe a general function that depends on $K$ variables, at least $2^{K}$ bits are required. Likewise, the sample complexity of learning Juntas is proportional to $t$ rather than $\log (t)$

Corollary 5.4 For every constant $q_{0}>0$, $q_{0}$-local queries cannot help to learn poly-sized Juntas.

Proof Consider the mapping $\varphi:\{-1,1\}^{n} \rightarrow\{-1,1\}^{\left(2 q_{0}+1\right) n}$ that replicates each coordinate $2 q_{0}+1$ times. To establish the corollary, we show that $\varphi$ is a $q_{0}$-reduction of type B from poly-sized Juntas to poly-sized Juntas. Indeed, let $h:\{-1,1\}^{n} \rightarrow\{0,1\}$ be a function that depends on $K$ variables. It is enough to show that there is a function $h^{\prime}$ on the variables $\left\{z_{i, j}\right\}_{i \in[n], j \in\left[2 q_{0}+1\right]}$ that satisfies (i) $h=h^{\prime} \circ \varphi$, (ii) for every $\mathbf{x} \in \mathcal{X}$, if $\mathbf{z} \in\{-1,1\}^{\left(2 q_{0}+1\right) n}$ is obtained from $\varphi(\mathbf{x})$ by modifying $\leq q_{0}$ coordinates, then $h^{\prime}(\mathbf{z})=h^{\prime}(\varphi(\mathbf{x}))$ and (iii) $h^{\prime}$ depends on $\left(2 q_{0}+1\right) K$ variables. It is not hard to check that the following function satisfies these requirements:

$$
h^{\prime}(\mathbf{z})=h\left(\operatorname{MAJ}\left(z_{1,1}, \ldots, z_{1,2 q_{0}+1}\right), \ldots, \operatorname{MAJ}\left(z_{n, 1}, \ldots, z_{n, 2 q_{0}+1}\right)\right)
$$

We remark that by taking $q_{0}=\log ^{1-\epsilon_{0}}(n)$ for some $\epsilon_{0}>0$, and using a similar argument, it can be shown that an efficient algorithm for learning poly-sized Juntas with $\left(\log ^{1-\epsilon_{0}}(n)\right)$-local queries would imply a PAC algorithm for poly-sized Juntas that runs in time $n^{O\left(\log ^{1-\epsilon_{0}}(n)\right)}$.

Corollary 5.5 For every constant $q_{0}>0$, $q_{0}$-local queries cannot help to learn poly-sized decision tress.

Proof As with Juntas, consider the mapping $\varphi:\{-1,1\}^{n} \rightarrow\{-1,1\}^{\left(2 q_{0}+1\right) n}$ that replicates each coordinate $2 q_{0}+1$ times. To establish the corollary, we show that $\varphi$ is a $q_{0}$-reduction of type B from poly-sized decision trees to poly-sized decision trees. Indeed, let $h:\{-1,1\}^{n} \rightarrow$ $\{0,1\}$ be a function that is realized by a decision tree $T$. It is enough to show that there is a function $h^{\prime}$ on the variables $\left\{z_{i, j}\right\}_{i \in[n], j \in\left[2 q_{0}+1\right]}$ that satisfies (i) $h=h^{\prime} \circ \varphi$, (ii) for every $\mathbf{x} \in \mathcal{X}$, if $\mathbf{z} \in\{-1,1\}^{\left(2 q_{0}+1\right) n}$ is obtained from $\varphi(\mathbf{x})$ by modifying $\leq q_{0}$ coordinates, then $h^{\prime}(\mathbf{z})=h^{\prime}(\varphi(\mathbf{x}))$ and (iii) $h^{\prime}$ can be realized by a tree of size $|T|^{2 q_{0}+1}$. As we explain, this will hold for the following function:

$$
h^{\prime}(\mathbf{z})=\operatorname{MAJ}\left(h\left(z_{1,1}, \ldots, z_{n, 1}\right), \ldots, h\left(z_{1,2 q_{0}+1}, \ldots, z_{n, 2 q_{0}+1}\right)\right)
$$

It is not hard to check that $h^{\prime}$ satisfies (i) and (ii). As for (iii), consider the following tree $T^{\prime}$. First replicate $T$ on the variables $z_{1,1}, \ldots, z_{n, 1}$. Then, on the obtained tree, replace each leaf by a replica of $T$ on the variables $z_{1,2}, \ldots, z_{n, 2}$. Then, again, replace each leaf by a replica of $T$ on the variables $z_{1,3}, \ldots, z_{n, 3}$. Continues doing so, until the leafs are replaced by a replica on the variables $z_{1,2 q_{0}+1}, \ldots, z_{n, 2 q_{0}+1}$. Now, each leaf in the resulted tree corresponds to $\left(2 q_{0}+1\right)$ root-to-leaf paths in the original tree $T$. The label of such leaf will be the majority of the labels of these paths.

As with Juntas, we remark that by taking $q_{0}=\log ^{1-\epsilon_{0}}(n)$ for some $\epsilon_{0}>0$, and using a similar argument, it can be shown that an efficient algorithm for learning poly-sized trees with $\left(\log ^{1-\epsilon_{0}}(n)\right)$-local queries would imply a PAC algorithm for poly-sized trees that runs in time $n^{O\left(\log ^{1-\epsilon_{0}}(n)\right)}$.

In the sequel, we say that a function $h:\{-1,1\}^{n} \rightarrow\{0,1\}$ is realized by a poly-sized polynomial, if it is the function induced by a polynomial with polynomially many nonzero coefficient and degree $O(\log (n))$. A similar convention applies to the term poly-sized polynomial threshold function. We remark that the following corollary and its proof remain correct also if in our definition, we replace the number of coefficients with the $\ell^{1}$ norm of the coefficients vector.

Corollary 5.6 For every constant $q_{0}>0$, $q_{0}$-local queries cannot help to learn poly-sized polynomials, as well as poly-sized polynomial threshold functions.

Proof We will prove the corollary for polynomials. The proof for polynomial threshold functions is similar. Consider the mapping $\varphi:\{-1,1\}^{n} \rightarrow\{-1,1\}^{\left(2 q_{0}+1\right) n}$ that replicates each coordinate $2 q_{0}+1$ times. To establish the corollary, we show that $\varphi$ is a $q_{0}$-reduction of type B from poly-sized polynomials to poly-sized polynomials. Indeed, let $h:\{-1,1\}^{n} \rightarrow$ $\{0,1\}$ be a poly-sized polynomial. It is enough to show that there is a poly-sized polynomial $h^{\prime}$ on the variables $\left\{z_{i, j}\right\}_{i \in[n], j \in\left[2 q_{0}+1\right]}$ that satisfies (i) $h=h^{\prime} \circ \varphi$, and (ii) for every $\mathbf{x} \in \mathcal{X}$, if $\mathbf{z} \in\{-1,1\}^{\left(2 q_{0}+1\right) n}$ is obtained from $\varphi(\mathbf{x})$ by modifying $\leq q_{0}$ coordinates, then $h^{\prime}(\mathbf{z})=$ $h^{\prime}(\varphi(\mathbf{x}))$. As we explain next, this will hold for the following polynomial:

$$
h^{\prime}(\mathbf{z})=h\left(\operatorname{MAJ}\left(z_{1,1}, \ldots, z_{1,2 q_{0}+1}\right), \ldots, \operatorname{MAJ}\left(z_{n, 1}, \ldots, z_{n, 2 q_{0}+1}\right)\right)
$$

It is not hard to check that $h^{\prime}$ satisfies (i) and (ii). It remains to show that $h^{\prime}$ is poly-sized. Indeed, the majority function on $2 q_{0}+1$ coordinates is a polynomial of degree $\leq 2 q_{0}+1$ with at most $2^{2 q_{0}+1}$ non zero coefficients (this is true for any function on $2 q_{0}+1$ coordinates). Hence, if we replace each variable in $h$ by a polynomial of the form $\operatorname{MAJ}\left(z_{i, 1}, \ldots, z_{i, 2 q_{0}+1}\right)$, the degree is multiplied by at most $\left(2 q_{0}+1\right)$, while the number of non zero coefficients is multiplied by at most $2^{2 q_{0}+1}$.
As with Juntas and decision trees, by taking $q_{0}=\log ^{1-\epsilon_{0}}(n)$ for some $\epsilon_{0}>0$, and using a similar argument, it can be shown that an efficient algorithm for learning poly-sized polynomials or polynomial threshold functions with $\left(\log ^{1-\epsilon_{0}}(n)\right)$-local queries would imply a PAC algorithm for the same problem that runs in time $n^{O\left(\log ^{1-\epsilon_{0}}(n)\right)}$.

## 6 Conclusion and Future Work

We have shown that in the distribution free setting, for many hypothesis classes, local queries are not useful. As our proofs show, this stems from the fact that learning these classes without queries can be reduced to a case where local queries are pointless, in the sense that the answer to them is either always 1 , or the label of the closest training example. On the other hand, the learning problem of DNFs with evident examples circumvents this property. Indeed, the underlying assumption enforces that local changes can change the label in a
non-trivial manner. While this assumption might be intuitive in some cases, it is certainly very restrictive. Therefore, a natural future research direction is to seek less restrictive assumptions, that still posses this property.

More concrete direction arising from our work concern classes for which we have shown that $\left(\log ^{0.99}(n)\right)$-local queries are unlikely to lead to efficient algorithms. We conjecture that for some $a>0$ even $\left(n^{a}\right)$-local queries won't lead to efficient distribution free algorithms for these classes.

## References

[1] Dana Angluin. Learning regular sets from queries and counterexamples. Information and computation, 75(2):87-106, 1987.
[2] Dana Angluin and Michael Kharitonov. When won t membership queries help? Journal of Computer and System Sciences, 50(2):336-355, 1995.
[3] Dana Angluin and Donna K Slonim. Randomly fallible teachers: Learning monotone dnf with an incomplete membership oracle. Machine Learning, 14(1):7-26, 1994.
[4] Dana Angluin, Mārtiņš Kriķis, Robert H Sloan, and György Turán. Malicious omissions and errors in answers to membership queries. Machine Learning, 28(2-3):211-255, 1997.
[5] Pranjal Awasthi, Vitaly Feldman, and Varun Kanade. Learning using local membership queries. arXiv preprint arXiv:1211.0996, 2012.
[6] Galit Bary. Learning using 1-local membership queries. arXiv preprint arXiv:1512.00165, 2015.
[7] Eric B Baum. Neural net algorithms that learn in polynomial time from examples and queries. Neural Networks, IEEE Transactions on, 2(1):5-19, 1991.
[8] Eric B Baum and Kenneth Lang. Query learning can work poorly when a human oracle is used. In International Joint Conference on Neural Networks, volume 8, 1992.
[9] Laurence Bisht, Nader H Bshouty, and Lawrance Khoury. Learning with errors in answers to membership queries. Journal of Computer and System Sciences, 74(1):2-15, 2008.
[10] Avrim Blum and Steven Rudich. Fast learning of k-term dnf formulas with queries. In Proceedings of the twenty-fourth annual ACM symposium on Theory of computing, pages 382-389. ACM, 1992.
[11] Avrim Blum, Prasad Chalasani, Sally A Goldman, and Donna K Slonim. Learning with unreliable boundary queries. In Proceedings of the eighth annual conference on Computational learning theory, pages 98-107. ACM, 1995.
[12] Nader H Bshouty. Exact learning boolean functions via the monotone theory. Information and Computation, 123(1):146-153, 1995.
[13] Amit Daniely and Shai Shalev-Shwatz. Complexity theoretic limitations on learning dnf's. arXiv preprint arXiv:1404.3378, 2014.
[14] Vitaly Feldman. On the power of membership queries in agnostic learning. The Journal of Machine Learning Research, 10:163-182, 2009.
[15] Yoav Freund. Boosting a weak learning algorithm by majority. Information and computation, 121(2):256-285, 1995.
[16] Jeffrey Jackson. An efficient membership-query algorithm for learning dnf with respect to the uniform distribution. In Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on, pages 42-53. IEEE, 1994.
[17] Michael Kearns and Leslie Valiant. Cryptographic limitations on learning boolean formulae and finite automata. Journal of the ACM (JACM), 41(1):67-95, 1994.
[18] Eyal Kushilevitz and Yishay Mansour. Learning decision trees using the fourier spectrum. SIAM Journal on Computing, 22(6):1331-1348, 1993.
[19] Michael Sipser. Introduction to the Theory of Computation, volume 2.
[20] Robert H Sloan and György Turán. Learning with queries but incomplete information. In Proceedings of the seventh annual conference on Computational learning theory, pages 237-245. ACM, 1994.
[21] Leslie G Valiant. A theory of the learnable. Communications of the ACM, 27(11): 1134-1142, 1984.


[^0]:    *This paper is based on the M.Sc. thesis [6] of the first author. The thesis offers a more elaborated discussion, as well as experiments.
    ${ }^{\dagger}$ Matific inc. Most work was done while the author was an M.Sc. student at the Hebrew University, Jerusalem, Israel
    ${ }^{\ddagger}$ Google inc. Most work was done while the author was a Ph.D. student at the Hebrew University, Jerusalem, Israel
    ${ }^{\S}$ School of Computer Science and Engineering, The Hebrew University, Jerusalem, Israel

[^1]:    ${ }^{1} \mathrm{~A}$ distribution is locally $\alpha$-smooth for $\alpha \geq 1$ if its density function is $\log (\alpha)$-Lipschitz.

[^2]:    ${ }^{2}$ The success probability can be amplified to $1-\delta$ by repetition.

[^3]:    ${ }^{3}$ We note that the quantity $\frac{1}{n}$ in the following equation is arbitrary, and can be replaced by $\frac{1}{n^{c}}$ for any $c>0$.

