Teaching via Best-Case Counterexamples in the Learning-with-Equivalence-Queries Paradigm

Akash Kumar[1,2]
akk002@ucsd.edu

Yuxin Chen[3]
chenyuxin@uchicago.edu

Adish Singla[1]
adishs@mpi-sws.org

[1]Max Planck Institute for Software Systems (MPI-SWS),
[2]UC San Diego,
[3]University of Chicago

Abstract

We study the sample complexity of teaching, termed as "teaching dimension" (TD) in the literature, for the learning-with-equivalence-queries (LwEQ) paradigm. More concretely, we consider a learner who asks equivalence queries (i.e., "is the queried hypothesis the target hypothesis?"), and a teacher responds either "yes" or "no" along with a counterexample to the queried hypothesis. This learning paradigm has been extensively studied when the learner receives worst-case or random counterexamples; in this paper, we consider the optimal teacher who picks best-case counterexamples to teach the target hypothesis within a hypothesis class. For this optimal teacher, we introduce LwEQ-TD, a notion of TD capturing the teaching complexity (i.e., the number of queries made) in this paradigm. We show that a significant reduction in queries can be achieved with best-case counterexamples, in contrast to worst-case or random counterexamples, for different hypothesis classes. Furthermore, we establish new connections of LwEQ-TD to the well-studied notions of TD in the learning-from-samples paradigm.

1 Introduction

Learning-with-queries paradigm involves a learner who asks structured queries to a teacher in order to locate a target hypothesis. This paradigm has been extensively studied in machine learning and formal methods research, including automata learning [1, 2, 3, 4], model checking [5], oracle-guided synthesis (OGIS) [6], model learning [7], among others. Classical literature involves different kinds of queries that a learner could ask [1, 8, 9], such as *membership* queries (i.e., "is the queried instance consistent with the target hypothesis?") and *equivalence* queries (i.e., "is the queried hypothesis the target hypothesis?"). In this paper, we consider the learning-with-equivalence-queries (LwEQ) paradigm where the learner is only asking equivalence queries and the teacher responds to a query either "yes" or "no" along with a counterexample on which the current hypothesis disagrees with the target hypothesis. LwEQ paradigm captures a variety of important problem settings, such as counterexample-guided synthesis (CEGIS) [10], data augmentation (CEGAR) [11], and learning regular languages from counterexamples [12]. The focus of our work is to understand the query complexity for the LwEQ paradigm, i.e., the number of equivalence queries needed by the learner for the exact identification of a target hypothesis [9].

The query complexity clearly depends on the learner model (i.e., the query function deciding the next hypothesis to query), as well as the informativeness of the counterexamples provided by the teacher. In the literature, the query complexity for the LwEQ paradigm has been extensively studied when the learner receives worst-case or random counterexamples [1, 9, 13, 14, 15]. To show worst-case bounds [1, 13], classical works have studied a teacher who responds with worst-case counterexamples

35th Conference on Neural Information Processing Systems (NeurIPS 2021).

to maximize the number of equivalence queries asked by the learner. Recent work [14] has studied this query complexity when facing a more benign teacher who provides counterexamples selected at random from a known probability distribution. In particular, [14] proposed a learning algorithm (i.e., a query function for the learner) that achieves $\mathcal{O}(\log n)$ query complexity on the expected number of random counterexamples for any hypothesis class of size n. When contrasting this bound with worst-case bounds, it shows that random counterexamples could lead to exponential improvement in the query complexity when compared with worst-case counterexamples [13, 14]. Along these lines, an important research question is to understand the query complexity where the counterexamples are picked by a more informed and helpful teacher, instead of random or worst-case counterexamples.

In this paper, we consider a more powerful teaching setting, where the optimal teacher picks best-case counterexamples to steer the learner towards a target hypothesis. Our goal is to characterize the query complexity for this optimal teacher (also, referred to as *teaching complexity* in the paper), for the LwEQ paradigm. This teaching complexity has been extensively studied for binary classification in the learning-from-samples (LfS) paradigm [16, 17] and is termed as "teaching dimension" (TD) [18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29]. Beyond the LfS paradigm, various notions of teaching complexity have also been studied in other learning paradigms (see Section 1.1 for further details). We introduce a new notion of teaching complexity for the LwEQ paradigm, namely "LwEQ teaching dimension" (LwEQ-TD), capturing the number of queries needed by the learner when best-case counterexamples are provided by the optimal teacher. Our study of LwEQ-TD for some prominent hypothesis classes reveals the power of teaching via best-case counterexamples—we show a significant reduction in query complexity when compared to that for worst-case and random counterexamples. Furthermore, we establish several new connections of LwEQ-TD to the existing notions of TD in the LfS paradigm.[1] Table 1 provides a summary of these different teaching settings and learning paradigms; our main results and contributions are summarized below:

I. We characterize the query complexity for the optimal teacher in the LwEQ paradigm, termed as *learning-with-equivalence-queries teaching dimension* (LwEQ-TD). (see Section 3)

II. We study the query complexity in the LwEQ paradigm under different teaching settings: worst-case, random, and best-case, distinguished by the informativeness of counterexamples. We showcase the power of best-case counterexamples picked by the optimal teacher, in contrast to worst-case or random counterexamples, for different hypothesis classes, including Axes-aligned hyperplanes, Monotone monomials, and Orthogonal rectangles. (see Section 4)

III. We establish new connections between LwEQ-TD and LfS-TD by studying LwEQ-TD for different learner models based on the richness of their query functions. We show that LwEQ-TD is the same as wc-TD [18], RTD [22, 24], and NCTD [27] for a hypothesis class when restricting query functions to specific families. In general, LwEQ-TD is weaker than LfS-TD, e.g., LwEQ-TD is lower-bounded by local-PBTD [26, 29] of the hypothesis class when the learner's next query depends on the previous query. (see Section 5)

Learning \ Teaching	Worst-case Teacher	Random-case Teacher	Best-case Teacher
Learning-with-equivalence-queries (LwEQ)	Worst-case counterexamples [1, 30, 9, 31, 13]	Random counterexamples [14, 32]	LwEQ-TD **This work**
Learning-from-samples (LfS)	Worst-case examples (i.e., least informative)	i.i.d learning [16, 17]	LfS-TD / classical TD [18, 22, 25, 27, 29]

Table 1: An overview of different teaching settings in the context of LwEQ and LfS paradigms.

1.1 Background and Related Work

Learning-with-queries paradigm and equivalence queries. Learning-with-queries paradigm was introduced in [1] which proposed L^* algorithm for exact identification of DFAs (deterministic finite automaton) when the learner is allowed to ask membership queries and equivalence queries. Classical work has studied different types of queries (subset, membership, equivalence, correction, among others) [1, 8, 9]. Among these works, learning with membership queries has been explored in a variety of problem settings, such as PAC learning [16], active learning [33, 34, 35], and agnostic

[1] We will collectively refer to these different notions of TD in the LfS paradigm as LfS-TD.

learning [36]. In the learning-with-equivalence-queries (LwEQ) paradigm, the learner can ask only equivalence queries; furthermore, in our work, we consider *proper* queries, i.e., the queried hypothesis is within the hypothesis class (see Section 6 for a discussion on *improper* queries). In this LwEQ paradigm, [13] studied query complexity when the teacher picks worst-case counterexamples for some key hypothesis classes (e.g., DFA, NFA, Context-free Grammars), and showed an exponential lower bound. [14] studied random counterexamples, and proposed a learning algorithm, namely Max-Min, which achieves a substantially improved bound on query complexity. We will investigate the query complexity of Max-Min learner for the hypothesis class of Axes-aligned hyperplanes in Section 4.

Teaching in the LfS paradigm for binary classification. Algorithmic machine teaching, first introduced by [18, 37], studies the interaction between a teacher and a learner where the teacher's goal is to find an optimal sequence of training samples to teach a target hypothesis. [18] introduced a measure of teaching complexity, named teaching dimension (TD) of the hypothesis class, in the learning-from-samples (LfS) paradigm. The classical notion of TD in [18] characterized the minimum number of samples (i.e., examples) needed to teach a target hypothesis to a version-space learner who picks hypothesis within the version space *arbitrarily* (in an adversarial way). In the past two decades, several new teaching settings have been studied, driven by the motivation to lower teaching complexity and to find settings for which TD has better connections with Vapnik–Chervonenkis dimension (VCD) [38]. In particular, several new teaching models and complexity measures have been proposed for both the batch teaching settings (e.g., worst-case [18], recursive [22, 24], preference-based [25], and non-clashing models [27]) and the sequential settings (e.g., local preference-based model [26, 29]). These teaching settings, in turn, lead to different notions of TD, that we collectively refer to as LfS-TD. In recent work, [29] has characterized these different notions of TD through a unified framework of modeling learners with preference/ranking functions. In Section 5, we will build on this framework to model the learner's query functions in the LwEQ paradigm through ranking functions, allowing us to connect LwEQ-TD with LfS-TD.

Teaching in other learning settings. Within binary classification setting, teaching complexity results have been extended beyond version space learners, including models for gradient learners [39, 40], models inspired by control theory [41, 42], and models for human-centered applications [43, 44, 45, 46]. Furthermore, a recent line of research has studied robust notions of teaching in settings where the teacher has limited information about the learner's dynamics [47, 48, 49]. Given the importance of teacher-learner interactions in many real-world applications, teaching has also been studied in richer domains. In particular, teaching complexity has been investigated for imitation learning settings where the teacher provides demonstrations [50, 51, 52, 53, 54], and for reinforcement learning settings where the teacher provides reward feedback [55, 56]. We see these works as complementary to ours, and we refer the reader to see [28] for an overview.

2 Problem Setup

Teaching framework. Let $\mathcal{X}$ be a ground set of unlabeled instances and $\mathcal{Y}$ the set of labels. Let $\mathcal{H}$ be a finite class of hypotheses; each element $h \in \mathcal{H}$ is a function $h : \mathcal{X} \to \mathcal{Y}$. Here, we only consider boolean functions, and hence $\mathcal{Y} = \{0, 1\}$. Let $\mathcal{Z} \subseteq \mathcal{X} \times \mathcal{Y}$ be the ground set of labeled examples. Each element $z = (x_z, y_z) \in \mathcal{Z}$ represents a labeled example where the label is given by the target hypothesis h^*, i.e., $y_z = h^*(x_z)$. Furthermore, for any $Z \subseteq \mathcal{Z}$, we define *version space* induced by the examples Z as the subset of hypotheses $\mathcal{H}(Z) \subseteq \mathcal{H}$ that are consistent with the labels of all the examples, i.e.,

$$\mathcal{H}(Z) := \{h \in \mathcal{H} \mid \forall z = (x_z, y_z) \in Z,\ h(x_z) = y_z\}. \tag{1}$$

Equivalence queries. We consider the LwEQ paradigm where a learner seeks to identify a target hypothesis from the hypothesis class via equivalence queries. In an equivalence query, the learner asks if the current hypothesis, say $h' \in \mathcal{H}$, is equivalent to the target hypothesis h^* or not. The teacher provides a response $\mathbf{r}$ where $\mathbf{r}$ is either "yes" if $h' \equiv h^*$ or "no" along with a counterexample $z := (x_z, y_z) \in \mathcal{Z}$, such that $h'(x_z) \neq y_z$.

Learner model and query protocol. We consider a generic model of the learner that captures our assumptions about how the learner conjectures its hypothesis for equivalence queries based on the responses received from the teacher. A key aspect of this model is the learner's query function ℓ

over the hypotheses. Based on the information encoded in the inputs of this query function (i.e., the current hypothesis and the history of counterexamples), the learner will choose one hypothesis in $\mathcal{H}$. In the beginning, the learner starts with an initial hypothesis $h_0 \in \mathcal{H}$, the history is $Z_0 = \emptyset$, and the version space is $H_0 = \mathcal{H}$. At a time step $t \geq 1$, the learner first picks the hypothesis h_t as follows:

$$\ell(Z_{t-1}, h_{t-1}) \longrightarrow h_t \in \mathcal{H}(Z_{t-1}), \tag{2}$$

where Z_{t-1} is the history of counterexamples seen up until time t and $\mathcal{H}(Z_{t-1})$ is the corresponding version space. The learner then queries h_t for equivalence to h^* and receives a response $\mathbf{r}_t$. Then, the query protocol proceeds as follows: (i) if $\mathbf{r}_t$ is "yes", the learner has identified h^* and stops; (ii) otherwise $\mathbf{r}_t$ is "no" along with a counterexample z_t using which the learner updates $Z_t = Z_{t-1} \cup \{z_t\}$, and continues. We summarize this query protocol in Algorithm 1.

Algorithm 1: Query protocol between the learner and the teacher

```
Learner's initial hypothesis is h_0 ∈ H, history is Z_0 = ∅, and version space is H_0 = H;
for t = 1, 2, 3, ··· do
    learner picks h_t ∈ H(Z_{t-1}) based on Z_{t-1} and h_{t-1} as per Eq. (2);
    learner performs an equivalence query with h_t;
    teacher provides a response r_t that is either "yes" or 'no" along with a counterexample z_t;
    if r_t is "yes" then
        learner has identified h* and stops;
    else
        learner updates Z_t = Z_{t-1} ∪ {z_t};
```

We assume that both the learner and the teacher have full knowledge of $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{H}$; in addition, the teacher knows the target h^* as well as the learner's query function ℓ. In this work, we consider learner models which could be characterized by a specific query function ℓ as discussed above. These include well-known learners studied in the literature, such as a *constant* query learner (denoted as ℓ_{const}) who picks the next hypothesis h_t arbitrarily in $\mathcal{H}(Z_{t-1})$ without any preference [18, 1, 29], a *global* query learner (denoted as ℓ_{global}) who uses a global ranking over $\mathcal{H}$ to pick the next hypothesis h_t in $\mathcal{H}(Z_{t-1})$ as per Eq. (2) [25, 29], and the Max-Min learning algorithm (denoted as $\ell_{\text{Max-Min}}$) introduced in a recent work on LwEQ paradigm [14].

Complexity of teaching (i.e., the number of queries made). In this paper, we study the number of equivalence queries asked by the learner to the teacher to identify a target hypothesis in Algorithm 1, and we call it the *query complexity* or *teaching complexity for LwEQ* paradigm interchangeably. Clearly, this query complexity depends on the learner's query function ℓ and the choice of counterexamples by the teacher. In the following sections, we study this complexity for different teacher types depending on the informativeness of the provided counterexamples, as well as for different families of learner types.

3 The Query Complexity with Best-Case Teacher: LwEQ-TD

In this section, we consider the optimal teacher who picks best-case counterexamples with the objective of minimizing the learner's queries for identifying h^*. For this optimal teacher, we provide a formal characterization of teaching complexity, namely *learning-with-equivalence queries teaching dimension* (LwEQ-TD) paradigm, inspired by different notions of teaching dimension for the learning-from-samples (LfS-TD) paradigm [18, 24, 26, 29].

Notation. We denote by $\mathcal{L}$ a family of learner models—alternatively, we can think of them as a family of query functions. To begin, we fix a query function $\ell \in \mathcal{L}$ that the learner uses to pick next hypotheses for equivalence queries. Our characterization below will be based on understanding the minimal "cost" (i.e., the number of queries needed) in steering the learner from a hypothesis h with the history of counterexamples Z to the target hypothesis h^*—in Algorithm 1, h refers to h_{t-1} and Z refers to Z_{t-1} at the beginning of time t.

Minimal cost of steering. We begin by providing a recursive function that captures this cost of steering and will be key to formalize teaching complexity in the LwEQ paradigm. As is typically considered in the LfS paradigm [18, 24, 26, 29], we consider the "adversarial" perspective in how the learner breaks ties in picking the hypothesis w.r.t. its query function in Eq. (2). The optimal cost for steering the learner from the current history (Z, h) to h^* in the query protocol of Algorithm 1 is then given by:

$$D_\ell(Z, h, h^*) = \begin{cases} 0, & \text{if } \ell(Z, h) = \{h^*\} \\ 1 + \max\limits_{h' \in \ell(Z,h)} \min\limits_{z : h'(x_z) \neq y_z} D_\ell(Z \cup \{z\}, h', h^*), & \text{otherwise} \end{cases} \tag{3}$$

Note that D_ℓ in Eq. (3) has a $\max$ operator w.r.t. the learner's choice and $\min$ operator with the teacher's choice (see [29]). Furthermore, in this function, we are not counting the last query when the learner's queried hypothesis h_t equals h^* (hence, the number of queries is same as the number of counterexamples received by the learner).

LwEQ-TD. Given a fixed query function $\ell \in \mathcal{L}$, an initial hypothesis h_0, and a target hypothesis h^*, LwEQ-TD w.r.t. ℓ and h^* is the optimal cost for teaching the target hypothesis h^*:

$$\text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\ell, h^*) = D_\ell(\emptyset, h_0, h^*). \tag{4}$$

To characterize the teaching complexity for the hypothesis class, we consider the worst-case target hypothesis in $\mathcal{H}$, given by:

$$\text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\ell) = \max_{h^*} D_\ell(\emptyset, h_0, h^*). \tag{5}$$

Finally, to compare LwEQ-TD with existing notions of TD in the LfS paradigm (i.e., LfS-TD, see Footnote 1), we define LwEQ-TD for a given family of query functions $\mathcal{L}$. Based on [24, 27, 29], we define LwEQ-TD w.r.t the family $\mathcal{L}$ as the teaching complexity w.r.t. the best ℓ in that family:

$$\text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\mathcal{L}) = \min_{\ell \in \mathcal{L}} \max_{h^*} D_\ell(\emptyset, h_0, h^*). \tag{6}$$

In the following two sections, we will investigate how LwEQ-TD (the complexity of the optimal teacher in the LwEQ paradigm) compares with various other complexity notions. In Section 4, we investigate different teachers in the LwEQ paradigm, showcasing the power of the optimal teacher. In Section 5, we establish new connections of LwEQ-TD with LfS-TD.

4 The Query Complexity for Different Teachers in the LwEQ Paradigm

In this section, we study the query complexity in the LwEQ paradigm for different types of teachers. In Section 3, we characterized the query complexity for the optimal teacher who provides best-case counterexamples. The goal of this section is to compare the query complexity for the optimal teacher with other variants of teachers as discussed below. We will denote the teacher as TEQ, and we consider the following four variants of teachers:

- best-TEQ responds "yes" or "no" along with best-case counterexamples (see Section 3).
- random-TEQ responds "yes" or "no" along with random counterexamples (picked uniformly at random). As discussed in [14], we characterize the query complexity for random-TEQ as the expected number of counterexamples provided by random-TEQ in Algorithm 1.
- worst-TEQ responds "yes" or "no" along with worst-case (least informative) counterexamples. We characterize query complexity for worst-TEQ by replacing the $\min$ over the choice of counterexamples with a $\max$ in the function D_ℓ in Eq. (3).
- binary-TEQ responds "yes" or "no" without any counterexamples.

In the rest of the section, we compare the query complexity for these teachers when interacting with different learners and we characterize this complexity by the richness of their query functions.

4.1 Warm-up: Query Complexity Bounds When Teaching Different Types of Learners

In this section, we study the query complexity bounds for different teachers when teaching various learner models characterized by the richness of the underlying query functions. In particular, we consider the query functions ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$ (see Section 2). ℓ_{const} picks the next hypothesis h_t arbitrarily in $\mathcal{H}(Z_{t-1})$ without any preference [29]. ℓ_{global} uses a global ranking over

$\mathcal{H}$ to pick the next hypothesis h_t in $\mathcal{H}(Z_{t-1})$ as per Eq. (2) and is popularly studied in the LfS paradigm [25, 29]. $\ell_{\text{Max-Min}}$ uses a richer history-dependent query function, and was introduced for achieving a better query complexity when the counterexamples are chosen randomly [14]. To study the query complexity for different settings (corresponding to four teachers and three learners), we consider the hypothesis class of Axes-aligned hyperplanes, a simple yet rich hypothesis class, that generalizes the canonical hypothesis class of threshold boolean functions to higher dimensions. Next, we introduce this hypothesis class.[2]

Definition 4.1 (Axes-aligned hyperplanes)**.** *Fix an input space $\mathcal{X} = \{1, 2, \ldots, n\}^d$ in $\mathbb{R}^d$ with labels set $\mathcal{Y} = \{0, 1\}$, where an input $\mathbf{x} \in \mathcal{X}$ is a d-dimensional point in $\mathbb{R}^d$, i.e., $\mathbf{x} := (x_1, x_2, \ldots, x_d)$ such that all $x_i \in [1, n+1]$. We define the hypothesis class of Axes-aligned hyperplanes as:*

$$\mathcal{H}_{\text{axh}} := \{h \mid \exists\, i \in [1, d+1],\ j \in [n+1],\ \text{s.t.}\ \forall \mathbf{x} \in \mathcal{X},\ \ h(\mathbf{x}) = 1 \text{ if } x_i \leq j,\ \text{otherwise } 0\}\,. \quad (7)$$

We summarize the query complexity bounds for ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$ under different teaching scenarios in Table 2. We note that $|\mathcal{H}_{\text{axh}}| = d \cdot (n+1)$. To capture the global ranking of ℓ_{global}, we use a ranking function $g : \mathcal{H}_{\text{axh}} \to \{0, \ldots, |\mathcal{H}_{\text{axh}}| - 1\}$. For a target hypothesis $h^* \in \mathcal{H}_{\text{axh}}$ and a ranking function g, the bounds for ℓ_{global} in Table 2 are described using the following two quantities: (i) $t_g^* := \left|\{h \in \mathcal{H}_{\text{axh}} \text{ s.t. } g(h) \leq g(h^*)\}\right|$ and (ii) n_g^* is defined as the highest number of hypotheses aligned along one of the axis that are preferred over h^* as per the ranking function g. Note that $1 \leq t_g^* \leq d \cdot (n+1)$ and $1 \leq n_g^* \leq n+1$. In the Appendix of the supplementary, we state the bounds from Table 2 as theorems and provide detailed proofs.

Learner \ Teacher	binary-TEQ (Yes/No)	worst-TEQ	random-TEQ	best-TEQ
ℓ_{const}	$d \cdot (n+1)$	$\Omega\left(d \cdot (n+1)\right)$	$\mathcal{O}\left(d \cdot \log(n+1)\right)$	2
ℓ_{global}	t_g^*	$\Omega\left(t_g^*\right)$	$\mathcal{O}\left(d \cdot \log n_g^*\right)$	1
$\ell_{\text{Max-Min}}$	(Not-Applicable)	$\Omega\left(d \cdot \log(n+1)\right)$	$\mathcal{O}\left(\log(d \cdot (n+1))\right)$	2

Table 2: Query complexity for the Axes-aligned hyperplanes considering different teachers and learners; see Section 4.1 for details.

4.2 Lower and Upper Bounds on Query Complexity

To establish lower bounds on the query complexity for three teachers (binary-TEQ, worst-TEQ, random-TEQ), we will consider learner models characterized with *optimal query functions*. In contrast to these lower bounds, we establish upper bounds on the query complexity for the best-case teacher (best-TEQ) by considering a weaker learner model characterized with a global query function (see Section 2). To study the bounds in these different teaching settings, we consider two more hypothesis classes beyond Axes-aligned hyperplanes: (i) Monotone monomials [18, 37] and (ii) Orthogonal rectangles [31, 18, 57, 58], that have been extensively studied in the LfS paradigm.

We summarize the results in Table 3 with detailed proofs deferred to the Appendix of the supplementary. First, we state the results on query complexity for Axes-aligned hyperplanes in the theorem below.

Theorem 1. *Consider the hypothesis class of Axes-aligned hyperplanes $\mathcal{H}_{\text{axh}}$ (see Eq. (7)). There exists a global learner ℓ_{global} such that best-*TEQ *achieves LwEQ-TD(ℓ_{global}) of exactly 1. In contrast, for any optimal learner, random-*TEQ *provides at least $\Omega(\log d)$ counterexamples "in expectation" and worst-*TEQ *provides at least $\Omega(d + \log(n+1))$ counterexamples "in the worst-case".*

The results in the above theorem are based on the following key ideas. The bound for best-TEQ is based on results from Section 4.1, where we showed that a specific choice of ℓ_{global} upper bounds the query complexity for best-TEQ by 1. On the other hand, as per Theorem 25 (see [14]), the query complexity for a random teacher is connected to the VC dimension of a hypothesis class independent of the learner model. We show that VCD($\mathcal{H}_{\text{axh}}$) is $\Omega(\log d)$, which entails a direct lower bound on the query complexity for random-TEQ for any learner model. Furthermore, we note that the

[2]Given two positive integers a and b where $a < b$, we use the following shorthand notation: $[a, b] = \{a, a+1, \ldots, b-1\}$ and $[b] = \{0, 1, \ldots, b-1\}$.

Hypothesis class \ Teacher	binary-TEQ (Yes/No)	worst-TEQ	random-TEQ	best-TEQ
Axes-aligned hyperplanes	$\Omega\left(d\cdot(n+1)\right)$	$\Omega\left(d+\log(n+1)\right)$	$\Omega\left(\log d\right)$	1
Monotone monomials	$\Omega\left(2^n\right)$	$\Omega\left(n\right)$	$\Omega\left(n\right)$	1
Orthogonal rectangles	$\Omega\left((n\cdot(n+1))^d\right)$	$\Omega\left(d\cdot\log(n+1)\right)$	$\Omega\left(d\right)$	2

Table 3: Lower and upper bounds on query complexity for the Axes-aligned hyperplanes, Monotone monomials, and Orthogonal rectangles when considering different teaching settings. The lower bounds for three teachers (binary-TEQ, worst-TEQ, random-TEQ) are established based on learner models characterized with *optimal query functions*; the upper bounds for the best-case teacher (best-TEQ) are established based on learner models characterized with a global query function. Additional details are provided in the proofs of Theorems 1, 2, and 3.

worst-TEQ could force any learner to query a hypothesis in every axis in addition to $\Omega\left(\log(n+1)\right)$ queries in the axis of the target hypothesis.

Now, we define the hypothesis class of Monotone monomials and then state the aforementioned query complexity bounds for this class in Theorem 2.

Definition 4.2 (Monotone monomials)**.** *Fix a set of literals $\{v_1, v_2, \ldots, v_n\}$, an input space $\mathcal{X} = \{0,1\}^n$, and labels set $\mathcal{Y} = \{0,1\}$. A monotone monomial is a negation-free conjunction of literals. If $mono(i_1, i_2, \ldots, i_k)$ denotes a monotone monomial over the set of k literals $\{v_{i_1}, v_{i_2}, \ldots, v_{i_k}\}$, then it canonically represents a hypothesis $h(\mathbf{x}) := x_{i_1} \wedge x_{i_2} \wedge \ldots \wedge x_{i_k}$ over the input space. With these notations, we define the hypothesis class of Monotone monomials as:*

$$\mathcal{H}_{\text{mono}} := \{h \,|\, \exists\ \{i_1, i_2, \ldots, i_k\} \subseteq \{1, 2, \ldots, n\},\ \text{s.t.}\ \forall \mathbf{x} \in \mathcal{X},\ \ h(\mathbf{x}) = x_{i_1} \wedge x_{i_2} \wedge \ldots \wedge x_{i_k}\,\}\,. \tag{8}$$

Theorem 2. *Consider the hypothesis class of Monotone monomials $\mathcal{H}_{\text{mono}}$ (see Eq. (8)). There exists a global learner ℓ_{global} such that best-*TEQ *achieves LwEQ-TD(ℓ_{global}) of exactly 1. In contrast, for any optimal learner, random-*TEQ *provides at least $\Omega\left(n\right)$ counterexamples "in expectation" and worst-*TEQ *provides at least $\Omega\left(n\right)$ counterexamples "in the worst-case".*

Now, we define the hypothesis class of Orthogonal rectangles and then state the aforementioned query complexity bounds for this class in Theorem 3.

Definition 4.3 (Orthogonal rectangles)**.** *Fix an input space $\mathcal{X} = \{1, \ldots, n\}^d$ in $\mathbb{R}^d$ with labels set $\mathcal{Y} = \{0,1\}$, where an input $\mathbf{x} \in \mathcal{X}$ is a d-dimensional point in $\mathbb{R}^d$, i.e, $\mathbf{x} := (x_1, x_2, \ldots, x_d)$ such that all $x_i \in [1, n+1]$. We define the class of Orthogonal rectangles as:*

$$\mathcal{H}_{\text{rec}} := \left\{ h \,\middle|\, \exists \left\{a_j, b_j\right\}_{j\in[1,d+1]} \subset [n+1]\,,\ \text{s.t.}\ \forall \mathbf{x} \in \mathcal{X},\ \ h(\mathbf{x}) = \begin{cases} 1 & \text{if}\ \ \forall j,\ a_j < x_j \leq b_j \\ 0 & \text{otherwise.} \end{cases} \right\}. \tag{9}$$

Theorem 3. *Consider the hypothesis class of Orthogonal rectangles $\mathcal{H}_{\text{rec}}$ (see Eq. (9)). There exists a global learner ℓ_{global} such that best-*TEQ *achieves LwEQ-TD(ℓ_{global}) of exactly 2. In contrast, for any optimal learner, random-*TEQ *provides at least $\Omega\left(d\right)$ counterexamples "in expectation" and worst-*TEQ *provides at least $\Omega\left(d\cdot\log(n+1)\right)$ counterexamples "in the worst-case".*

Similar to Axes-aligned hyperplanes (see Theorem 1), the lower bound results are significantly worse than the upper bound results for rich classes of Monotone monomials (see Theorem 2) and Orthogonal rectangles (see Theorem 3), thereby establishing the power of the best-case teacher. These results are summarized in Table 3.

5 Teaching Dimensions for the LwEQ and LfS Paradigms

In this section, we draw new comparisons of the notion of LwEQ-TD in the LwEQ paradigm to existing notions of TD in the LfS paradigm. Based on the teaching setting and the learner models, one gets different notions of teaching complexity and we collectively refer to these notions as LfS-TD (see Footnote 1) [18, 21, 22, 25, 27, 59, 29]. For comparisons with LfS-TD below, we consider the framework of [29] that provides a unified view of teaching settings by modeling the learners through preference/ranking functions. In the following, we first introduce a notation for these ranking functions σ and then discuss how a ranking function σ, in turn, induces a learner in the LwEQ paradigm.

Learner's query function ℓ using a framework of a ranking σ. Consider a hypothesis class $\mathcal{H}$, a version space $H \subseteq \mathcal{H}$, and hypotheses $h', h'', h \in \mathcal{H}$ such that $h', h'' \in H$. Building on the notation for a preference function σ (see [29]), we define a ranking function $\sigma : \mathcal{H} \times 2^{\mathcal{H}} \times \mathcal{H} \to \mathbb{R}$ where $\sigma(h'; H, h)$ signifies how h' is ranked in the version space H from the current hypothesis h. Thus, we say h' is *ranked* (or preferred) over h'' in the current version space H from the current hypothesis h if $\sigma(h'; H, h) \leq \sigma(h''; H, h)$, and vice versa. Similar to preference-based learners [29], a learner's query function (see Section 2) could use this ranking to pick the most preferred hypothesis for an equivalence query. In this section, we consider query functions for our learner model ℓ (see Section 2) which use a framework of a ranking function σ to pick the next hypothesis h_t based on the current history of counterexamples seen Z_{t-1} and the current hypothesis h_{t-1} as follows:

$$\ell(Z_{t-1}, h_{t-1}) \longrightarrow h_t \in \operatorname*{arg\,min}_{h' \in \mathcal{H}(Z_{t-1})} \sigma(h'; \mathcal{H}(Z_{t-1}), h_{t-1}). \tag{10}$$

In this section, we identify a learner model ℓ_σ[3] with the corresponding ranking σ and consequently use the notation LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma)$ (for fixed $\mathcal{X}, \mathcal{H}, h_0$) for LwEQ-TD (see Eq. (5), Section 3) for the learner model ℓ_σ. We denote a family of ranking functions σ as Σ. For a family of ranking functions Σ, the corresponding LwEQ teaching dimension is denoted as LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma)$ (see Eq. (6), Section 3).

For families of learner models induced by specific types of ranking functions, we connect LwEQ-TD to existing notions of LfS-TD. In the following, we consider ranking functions broadly categorized into two classes: (i) ranking functions independent of Z_{t-1} and h_{t-1}; (ii) ranking functions dependent on Z_{t-1} and/or h_{t-1}.

5.1 LwEQ Learners with Ranking Functions Independent of Z_{t-1} and h_{t-1}

These ranking functions induce learners whose next equivalence query at time t (i.e., choice of the hypothesis h_t) is independent of the history of counterexamples Z_{t-1} and independent of the hypothesis h_{t-1} (see Algorithm 1, Section 2). In Section 4.1, we discussed these learner families under the name of *contant* and *global* learners, and we formalize these families below using the ranking functions framework of Eq. (10). In particular, we introduce two families of ranking functions: (i) Σ_{const} is a family of *constant* ranking functions where $\sigma \in \Sigma_{\text{const}}$ ranks every hypothesis equally without any preference; (ii) Σ_{global} is a family of *global* ranking functions where $\sigma \in \Sigma_{\text{global}}$ ranks hypothesis based on a global preference. These two families are given below:

$$\Sigma_{\text{const}} = \{\sigma \mid \exists\, c \in \mathbb{R}, \text{ s.t.} \forall h', Z, h,\ \sigma(h'; \mathcal{H}(Z), h) = c\}. \tag{11}$$

$$\Sigma_{\text{global}} = \{\sigma \mid \exists\, g : \mathcal{H} \to \mathbb{R}, \text{ s.t.} \forall h', Z, h,\ \sigma(h'; \mathcal{H}(Z), h) = g(h')\}. \tag{12}$$

Given these ranking functions, the following theorem establishes the connection of LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{const}})$ and LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{global}})$ with existing notions of LfS-TD; also see Eq. (6) in Section 3.

Theorem 4. *Fix $\mathcal{X}, \mathcal{H}, h_0$. For learners whose query function is induced by a ranking function independent of Z_{t-1} and h_{t-1}, the corresponding LwEQ-TD is equivalent to the notions of LfS-TD as follows:*

$$\text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{const}}) = \text{wc-TD}(\mathcal{H}). \tag{13}$$

$$\text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{global}}) = \text{RTD}(\mathcal{H}). \tag{14}$$

When the rankings are independent of Z_{t-1} and h_{t-1}, Eq. (13) and Eq. (14) connect LwEQ-TD to the notions of wc-TD [18] and RTD [22, 24] in batch teaching settings (the optimal teaching sequence is invariant to its permutation) of the LfS paradigm. To show the equalities, we note that if $Z' \subseteq \mathcal{Z}$ is a teaching sequence for fixed h^* in the LfS paradigm, then there exists a permutation of Z' which forms a teaching sequence of counterexamples for the learner in the LwEQ

[3]We use the notation ℓ_σ for a learner's query function ℓ using a ranking σ.

paradigm. Furthermore, this theorem allows us to connect LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{global}})$ with VCD$(\mathcal{H})$ as RTD$(\mathcal{H}) = \mathcal{O}\left(\text{VCD}^2(\mathcal{H})\right)$ [24, 29]. A detailed proof of the theorem is in the Appendix of the supplementary.

5.2 LwEQ Learners with Ranking Functions dependent on Z_{t-1} and/or h_{t-1}

Here, we consider ranking functions that induce learners whose next equivalence query at time t (i.e., choice of the hypothesis h_t) is dependent on the history of counterexamples Z_{t-1} and/or on the hypothesis h_{t-1} (see Algorithm 1, Section 2). We formalize these families below using the ranking functions framework of Eq. (10). In particular, we introduce three families of ranking functions: (i) Σ_{gvs} is a family of *global version space* ranking functions where $\sigma \in \Sigma_{\text{gvs}}$ ranks hypotheses based on a global preference dependent on Z_{t-1} but independent of h_{t-1}; (ii) Σ_{local} is a family of *local* ranking functions where $\sigma \in \Sigma_{\text{local}}$ ranks hypotheses based on a local preference dependent only on h_{t-1}; (iii) Σ_{lvs} is a family of *local version space* ranking functions where $\sigma \in \Sigma_{\text{lvs}}$ ranks hypotheses based on a local preference dependent on Z_{t-1} and h_{t-1}. These three families are given below:

$$\Sigma_{\text{gvs}} = \left\{\sigma \,\middle|\, \exists\, g : \mathcal{H} \times 2^{\mathcal{H}} \to \mathbb{R}, \text{ s.t.} \,\forall h', Z, h,\ \sigma(h'; \mathcal{H}(Z), h) = g(h', \mathcal{H}(Z)) \right\}. \tag{15}$$

$$\Sigma_{\text{local}} = \{\sigma \mid \exists\, g : \mathcal{H} \times \mathcal{H} \to \mathbb{R}, \text{ s.t.} \,\forall h', Z, h,\ \sigma(h'; \mathcal{H}(Z), h) = g(h', h) \}. \tag{16}$$

$$\Sigma_{\text{lvs}} = \left\{\sigma \,\middle|\, \exists\, g : \mathcal{H} \times 2^{\mathcal{H}} \times \mathcal{H} \to \mathbb{R}, \text{ s.t.} \,\forall h', Z, h,\ \sigma(h'; \mathcal{H}(Z), h) = g(h', \mathcal{H}(Z), h) \right\}. \tag{17}$$

As discussed in the LfS paradigm, these ranking functions could lead to the learner and the teacher colluding to achieve arbitrarily low teaching complexity [29]. To avoid this, we consider a specific collusion-free behavior where the ranking is consistent with its choice of hypothesis. More formally,

Definition 5.1 (**Collusion-free ranking** [29])**.** *Consider a time t where the learner's current hypothesis is h_{t-1} and the history of inputs seen is Z_{t-1}. Further assume that the learner's preferred hypothesis for time t is uniquely given by* $\arg\min_{h' \in \mathcal{H}(Z_{t-1})} \sigma(h'; \mathcal{H}(Z_{t-1}), h_{t-1}) = \{\hat{h}\}$*. Let S be additional examples provided by an adversary from time t onwards. We call a ranking function σ collusion-free if for any S consistent with $\hat{h}$, it holds that* $\arg\min_{h' \in \mathcal{H}(Z_{t-1} \cup S)} \sigma(h'; \mathcal{H}(Z_{t-1} \cup S), \hat{h}) = \{\hat{h}\}$.

For the ranking functions in Eqs. (15)-(17), we consider subsets that satisfy Definition 5.1, and consequently the corresponding collusion-free families are denoted as $\Sigma^{\text{CF}}_{\text{gvs}}$, $\Sigma^{\text{CF}}_{\text{local}}$, and $\Sigma^{\text{CF}}_{\text{lvs}}$. For these ranking functions, we study the connection of LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{gvs}})$, LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{local}})$, and LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{lvs}})$ with existing notions of LfS-TD, in particular NCTD [27], local-PBTD [59, 29], and wc-TD [18, 29]. We state these connections in the following theorem and provide detailed proofs in the Appendix of the supplementary.

Theorem 5. *Fix $\mathcal{X}, \mathcal{H}, h_0$. For learners whose query function is induced by a ranking function dependent on Z_{t-1} and/or h_{t-1}, the corresponding LwEQ-TD is connected to the notions of LfS-TD as:*

$$\text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}\left(\Sigma^{\text{CF}}_{\text{gvs}}\right) = \text{NCTD}(\mathcal{H}). \tag{18}$$

$$\text{local-PBTD}_{\mathcal{X},\mathcal{H},h_0} \le \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{local}}) \le \text{wc-TD}(\mathcal{H}). \tag{19}$$

$$\text{lvs-PBTD}_{\mathcal{X},\mathcal{H},h_0} \le \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{lvs}}) \le \text{wc-TD}(\mathcal{H}). \tag{20}$$

To achieve equality in Eq. (18), we observe that the family of rankings $\Sigma^{\text{CF}}_{\text{gvs}}$ leads to a batch teaching setting in the LfS paradigm. To show this equality, we establish the permutation invariance of a teaching set $Z \subseteq \mathcal{Z}$ in the LfS paradigm to form a teaching sequence of counterexamples in the LwEQ paradigm using the collusion-freeness property of the underlying ranking function.

To show the lower bounds in Eq. (19) and Eq. (20), we note that a teaching sequence of counterexamples in the LwEQ paradigm forms a teaching sequence in the LfS paradigm. For the upper bound of wc-TD$(\mathcal{H})$, we note that a *constant* query learner (see Eq. (11)) could pick any consistent hypothesis in a version space to maximize the number of counterexamples. This observation, along with the results in Theorem 4, leads to the desired upper bound.

6 Concluding Discussions

We investigated the query complexity for the learning-with-equivalence-queries (LwEQ) paradigm when the counterexamples are provided by the optimal teacher. We introduced LwEQ-TD, a notion of teaching dimension, to characterize the complexity of teaching (i.e., the number of queries made) for the optimal teacher. We showed the power of best-case counterexamples picked by the optimal teacher, in contrast to worst-case or random counterexamples, for different hypothesis classes, including Axes-aligned hyperplanes, Monotone monomials, and Orthogonal rectangles. We further established new connections of LwEQ-TD with existing notions of TD in the learning-with-samples paradigm, including wc-TD, RTD, NCTD, and local-PBTD.

In the learning-with-queries paradigm, several works have analyzed the query complexity for a combination of different query types [1, 8, 15], such as membership queries, equivalence queries, among others. Building on our characterization of LwEQ-TD, an important research direction of future work is to investigate similar notions of TD when a learner can ask a combination of different queries. Alternatively, one could study LwEQ in the setting of *improper* equivalence queries where the learner can pick the queried hypothesis outside the hypothesis class. [31] considered improper equivalence queries in the LwEQ paradigm, leading to a reduction in query complexity for various hypothesis classes. It would be important to characterize the teaching complexity for best-case counterexamples with improper equivalence queries.

Another important research direction is to further investigate LwEQ-TD for more complex hypothesis classes, including DFA (deterministic finite automaton), NFA (nondeterministic finite-state acceptors), CFG (context-free grammars), among others. As a concrete hypothesis class, one could consider DFA_s^2 (i.e., DFA with alphabet size 2 and state size s). For this hypothesis class, the query complexity is exponential in s when considering worst-case counterexamples [13]. [60] computed VCD for a variety of hypothesis classes including DFA_s^2; this result along with query complexity bounds of the Max-Min algorithm in [14] establishes the bound of $\Theta\left(s \log s\right)$ on the expected number of random counterexamples for DFA_s^2. As future work, it would be interesting to investigate LwEQ-TD for DFA_s^2. A concrete direction is to study the query complexity when providing more structured counterexamples (e.g., by picking minimal length counterexamples as considered in [12]) as this would allow establishing an upper bound for best-case counterexamples.

References

[1] Dana Angluin. Learning Regular Sets from Queries and Counterexamples. *Information and Computation*, 75(2):87–106, 1987.

[2] Leonard Pitt. Inductive Inference, DFAs, and Computational Complexity. In *Proceedings of the International Workshop on Analogical and Inductive Inference*, AII '89, page 18–44, 1989.

[3] J. Balcázar, J. Díaz, and R. Gavaldà. Algorithms for Learning Finite Automata from Queries: A Unified View. In *Advances in Algorithms, Languages, and Complexity*, 1997.

[4] B. Bollig, P. Habermehl, C. Kern, and M. Leucker. Angluin-Style Learning of NFA. In *IJCAI*, 2009.

[5] Doron Peled, Moshe Y. Vardi, and Mihalis Yannakakis. *Black Box Checking*, pages 225–240. 1999.

[6] Susmit Jha and Sanjit Seshia. A Theory of Formal Synthesis via Inductive Learning. *Acta Informatica*, 54, 2017.

[7] Frits Vaandrager. Model Learning. *Communications of the ACM*, 60(2):86–95, 2017.

[8] Dana Angluin. Queries Revisited. *Theoretical Computer Science*, 313(2):175 – 194, 2004.

[9] Dana Angluin. Queries and Concept Learning. *Machine Learning*, 2(4):319–342, 1988.

[10] Armando Solar-Lezama, Liviu Tancau, Rastislav Bodik, Sanjit Seshia, and Vijay Saraswat. Combinatorial Sketching for Finite Programs. *SIGARCH*, 34(5):404–415, 2006.

[11] Tommaso Dreossi, Shromona Ghosh, Xiangyu Yue, Kurt Keutzer, Alberto Sangiovanni-Vincentelli, and Sanjit A. Seshia. Counterexample-Guided Data Augmentation. In *IJCAI*, page 2071–2078, 2018.

[12] Oscar Habarra and Tao Jiang. Learning Regular Languages from Counterexamples. *Journal of Computer and System Sciences*, 43(2):299–316, 1991.

[13] Dana Angluin. Negative Results for Equivalence Queries. *Machine Learning*, 5(2):121–150, 1990.

[14] Dana Angluin and Tyler Dohrn. The Power of Random Counterexamples. In *ALT*, pages 452–465, 2017.

[15] Hunter Chase and James Freitag. Bounds in Query Learning. In *COLT*, pages 1142–1160, 2020.

[16] L. G. Valiant. A Theory of the Learnable. *Communications of ACM*, 27(11):1134–1142, 1984.

[17] V. N. Vapnik. An Overview of Statistical Learning Theory. *IEEE Transactions on Neural Networks*, 10(5):988–999, 1999.

[18] Sally A. Goldman and Michael J. Kearns. On the Complexity of Teaching. *Journal of Computer and System Sciences*, 50(1):20–31, 1995.

[19] Christian Kuhlmann. On Teaching and Learning Intersection-Closed Concept Classes. In *European Conference on Computational Learning Theory*, pages 168–182, 1999.

[20] Frank J Balbach. Measuring Teachability using Variants of the Teaching Dimension. *Theoretical Computer Science*, 397(1-3):94–113, 2008.

[21] Sandra Zilles, Steffen Lange, Robert Holte, and Martin Zinkevich. Teaching Dimensions Based on Cooperative Learning. In *COLT*, pages 135–146, 2008.

[22] Sandra Zilles, Steffen Lange, Robert Holte, and Martin Zinkevich. Models of Cooperative Teaching and Learning. *JMLR*, 12:349–384, 2011.

[23] Xiaojin Zhu. Machine Teaching for Bayesian Learners in the Exponential Family. In *NeurIPS*, pages 1905–1913, 2013.

[24] Thorsten Doliwa, Gaojian Fan, Hans Ulrich Simon, and Sandra Zilles. Recursive Teaching Dimension, VC-Dimension and Sample Compression. *JMLR*, 15(1):3107–3131, 2014.

[25] Ziyuan Gao, Christoph Ries, Hans U Simon, and Sandra Zilles. Preference-Based Teaching. *JMLR*, 18(31):1–32, 2017.

[26] Yuxin Chen, Adish Singla, Oisin Mac Aodha, Pietro Perona, and Yisong Yue. Understanding the Role of Adaptivity in Machine Teaching: The Case of Version Space Learners. In *NeurIPS*, 2018.

[27] David Kirkpatrick, Hans U. Simon, and Sandra Zilles. Optimal Collusion-Free Teaching. In *ALT*, volume 98, pages 506–528, 2019.

[28] Xiaojin Zhu, Adish Singla, Sandra Zilles, and Anna N. Rafferty. An Overview of Machine Teaching. *CoRR*, abs/1801.05927, 2018.

[29] Farnam Mansouri, Yuxin Chen, Ara Vartanian, Xiaojin Zhu, and Adish Singla. Preference-based Batch and Sequential Teaching: Towards a Unified View of Models. In *NeurIPS*, pages 9195–9205, 2019.

[30] Dana Angluin, Lisa Hellerstein, and Marek Karpinski. Learning Read-Once Formulas with Queries. Technical report, USA, 1989.

[31] G. Turan and W. Maass. On the Complexity of Learning from Counterexamples. In *FOCS*, pages 262–267, 1989.

[32] Jagdeep Singh Bhatia. Simple and Fast Algorithms for Interactive Machine Learning with Random Counter-examples. *JMLR*, 22(15):1–30, 2021.

[33] Y. Freund, H. Seung, E. Shamir, and Naftali Tishby. Selective Sampling Using the Query by Committee Algorithm. *Machine Learning*, 28:133–168, 2004.

[34] D. Cohn, L. Atlas, and R. Ladner. Improving Generalization with Active Learning. *Machine Learning*, 15:201–221, 1994.

[35] Sanjoy Dasgupta. Coarse Sample Complexity Bounds for Active Learning. In *NeurIPS*, page 235–242, 2005.

[36] Vitaly Feldman. On the Power of Membership Queries in Agnostic Learning. *JMLR*, 10:163–182, 2009.

[37] Ayumi Shinohara and Satoru Miyano. Teachability in Computational Learning. *New Gen. Comput.*, 8(4):337–347, 1991.

[38] VN Vapnik and A Ya Chervonenkis. On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities. *Theory of Probability and its Applications*, 16(2):264, 1971.

[39] Weiyang Liu, Bo Dai, Ahmad Humayun, Charlene Tay, Chen Yu, Linda B. Smith, James M. Rehg, and Le Song. Iterative Machine Teaching. In *ICML*, pages 2149–2158, 2017.

[40] Weiyang Liu, Bo Dai, Xingguo Li, Zhen Liu, James M. Rehg, and Le Song. Towards Black-box Iterative Machine Teaching. In *ICML*, pages 3147–3155, 2018.

[41] Xiaojin Zhu. An Optimal Control View of Adversarial Machine Learning. *CoRR*, abs/1811.04422, 2018.

[42] Laurent Lessard, Xuezhou Zhang, and Xiaojin Zhu. An Optimal Control Approach to Sequential Machine Teaching. In *AISTATS*, pages 2495–2503, 2019.

[43] Adish Singla, Ilija Bogunovic, G Bartók, A Karbasi, and A Krause. On Actively Teaching the Crowd to Classify. In *NIPS Workshop on Data Driven Education*, 2013.

[44] Adish Singla, Ilija Bogunovic, Gábor Bartók, Amin Karbasi, and Andreas Krause. Near-Optimally Teaching the Crowd to Classify. In *ICML*, volume 32, pages 154–162, 2014.

[45] Anette Hunziker, Yuxin Chen, Oisin Mac Aodha, Manuel Gomez Rodriguez, Andreas Krause, Pietro Perona, Yisong Yue, and Adish Singla. Teaching Multiple Concepts to a Forgetful Learner. In *NeurIPS*, pages 4050–4060, 2019.

[46] Oisin Mac Aodha, Shihan Su, Yuxin Chen, Pietro Perona, and Yisong Yue. Teaching Categories to Human Learners with Visual Explanations. In *CVPR*, pages 3820–3828, 2018.

[47] Sanjoy Dasgupta, Daniel Hsu, Stefanos Poulis, and Xiaojin Zhu. Teaching a Black-Box Learner. In *ICML*, pages 1547–1555, 2019.

[48] Rati Devidze, Farnam Mansouri, Luis Haug, Yuxin Chen, and Adish Singla. Understanding the Power and Limitations of Teaching with Imperfect Knowledge. In *IJCAI*, pages 2647–2654.

[49] Ferdinando Cicalese, Sergio Filho, Eduardo Laber, and Marco Molinaro. Teaching with Limited Information on the Learner's Behaviour. In *ICML*, 2020.

[50] Maya Cakmak and Manuel Lopes. Algorithmic and Human Teaching of Sequential Decision Tasks. In *AAAI*, volume 26, 2012.

[51] Thomas J. Walsh and Sergiu Goschin. Dynamic Teaching in Sequential Decision Making Environments. In *UAI*, page 863–872, 2012.

[52] Parameswaran Kamalaruban, Rati Devidze, Volkan Cevher, and Adish Singla. Interactive Teaching Algorithms for Inverse Reinforcement Learning. In *IJCAI*, pages 2692–2700, 2019.

[53] Daniel S Brown and Scott Niekum. Machine Teaching for Inverse Reinforcement Learning: Algorithms and Applications. In *AAAI*, 2019.

[54] Sebastian Tschiatschek, Ahana Ghosh, Luis Haug, Rati Devidze, and Adish Singla. Learner-aware Teaching: Inverse Reinforcement Learning with Preferences and Constraints. In *NeurIPS*, 2019.

[55] Amin Rakhsha, Goran Radanovic, Rati Devidze, Xiaojin Zhu, and Adish Singla. Policy Teaching via Environment Poisoning: Training-time Adversarial Attacks against Reinforcement Learning. In *ICML*, 2020.

[56] Xuezhou Zhang, Shubham Kumar Bharti, Yuzhe Ma, Adish Singla, and Xiaojin Zhu. The Sample Complexity of Teaching by Reinforcement on Q-Learning. In *AAAI*, 2021.

[57] Kathleen Romanik. Approximate Testing and Learnability. page 327–332, 1992.

[58] Kathleen Romantik and Carl Smith. Testing Geometric Objects. *Computational Geometry*, 4(3):157–176, 1994.

[59] Yuxin Chen, Oisin Mac Aodha, Shihan Su, Pietro Perona, and Yisong Yue. Near-Optimal Machine Teaching via Explanatory Teaching Sets. In *AISTATS*, 2018.

[60] Y. Ishigami and Sei'ichi Tani. The VC-Dimensions of Finite Automata with n States. In *ALT*, 1993.

[61] Z. Gao, D. Kirkpatrick, Christoph Ries, H. Simon, and Sandra Zilles. Preference-based Teaching of Unions of Geometric Objects. In *ALT*, 2017.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Section 4 and Section 5.
   (b) Did you describe the limitations of your work? [Yes] We discuss how our results can be extended further as future work directions in Section 6.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] All proofs can be found in appendices.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A List of Appendices

Now, we list the appendices which provide the proofs of our theoretical results in full detail. The appendices are summarized as follows:

- Appendix B discusses different types of learners and contains proofs of the formal results in Section 4.1.
  - Appendix B.1 discusses learner models characterized by query functions ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$.
  - Appendix B.2 defines the hypothesis class of Threshold functions and provides query complexity bounds for different teachers when teaching learner models characterized by query functions ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$.
  - Appendix B.3 contains proofs of the formal results shown in Table 2 for Axes-aligned hyperplanes.
- Appendix C contains proofs of the formal results as shown in Table 3 in Section 4.2.
  - The proof of Theorem 1 is in Appendix C.1.
  - The proof of Theorem 2 is in Appendix C.2.
  - The proof of Theorem 3 is in Appendix C.3.
- Appendix D contains proofs of the formal results in Section 5.
  - The proof of Theorem 4 is in Appendix D.1.
  - The proof of Theorem 5 is in Appendix D.2.

B Query Complexity Bounds When Teaching Different Types of Learners (Section **4.1**)

In this appendix, we provide the proofs for the bounds shown in Table 2 on the query complexity for Axes-aligned hyperplanes (Definition 4.1). We divide the appendix as follows: Appendix B.1 discusses learner models characterized by query functions ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$; Appendix B.2 discusses Threshold functions and their query complexity bounds for different teachers; Appendix B.3 provides the proofs to entries in Table 2.

B.1 Learner Types

In the following, we discuss learners characterized by the query functions ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$ used for the query complexity bounds in Table 2.

Constant query learner (ℓ_{const}). A constant query learner (ℓ_{const}) picks the next hypothesis h_t arbitrarily in $\mathcal{H}(Z_{t-1})$ without any preference [18, 1, 29]. To analyze the query complexity for different hypothesis classes we assume that the learner picks the worst-case hypothesis in $\mathcal{H}(Z_{t-1})$ to query for equivalence. Since all the constant query learners have the same query function (see Section 2) the query complexity bounds are identified for the entire family of constant learners. We characterize a family of constant query learners in Section 5.

Global query learner (ℓ_{global}). A global query learner (ℓ_{global}) uses a global ranking over $\mathcal{H}$ to pick the next hypothesis h_t in $\mathcal{H}(Z_{t-1})$ as per Eq. (2) [25, 29]. We denote the global ranking of ℓ_{global} using a function $g : \mathcal{H} \to [|\mathcal{H}|]$. Similar to a constant learner ℓ_{const}, we assume that ℓ_{global} picks the worst ranked hypothesis according to the function g for querying (ties are broken arbitrarily). We characterize a family of global query learners in Section 5.

Max-Min query learner ($\ell_{\text{Max-Min}}$). In the LwEQ paradigm, [14] considered the setting where a learner queries a hypothesis for equivalence and receives a random counterexample picked from a known probability distribution and introduced the Max-Min learning algorithm. A Max-Min query learner picks the next hypothesis h_t based on the current history of counterexamples Z_{t-1} but *agnostic* of the current hypothesis h_{t-1} as follows:

$$h_t \in \underset{h \in \mathcal{H}(Z_{t-1})}{\arg\max} \; \min_{h' \in \mathcal{H}(Z_{t-1}) \setminus \{h\}} E(h, h') \tag{21}$$

where $E(h, h')$ is given by the expected fraction[4] of hypothesis eliminated from $\mathcal{H}(Z_{t-1})$ if an equivalence query is posed with hypothesis h and target hypothesis h', assuming counterexamples are picked from a known probability distribution. In this work, we assume that the underlying probability distribution is uniform. Furthermore, to evaluate a Max-Min query learner in the worst-case teaching scenario (similar for best-case teaching), i.e when worst-TEQ provides counterexamples (see *Column 2* in Table 2) we assume that the learner computes $E(h, h')$ modeling a uniform distribution over the set of valid counterexamples. We discuss the query complexity bounds for $\ell_{\text{Max-Min}}$ under different teaching scenarios in Appendix B.3.

B.2 Threshold Functions

Now, we discuss Threshold functions and study their query complexity bounds under different teaching scenarios in the LwEQ paradigm. Results and insights gained in this section would be used in many other results in our work. First, we define the hypothesis class of Threshold functions as follows:

Definition B.1 (Threshold functions). *Fix an input space $\mathcal{X} = \{1, \cdots, n\}$ in $\mathbb{R}$ with labels set $\mathcal{Y} = \{0, 1\}$. We define the hypothesis class of Threshold functions as:*

$$\mathcal{H}_{\text{threshold}} := \{h \,|\, \exists\, j \in [n+1], \text{ s.t. } \forall x \in \mathcal{X},\;\; h(x) = 1 \text{ if } x \leq j, \text{ otherwise } 0\}\,. \tag{22}$$

[4] Query function remains the same if we compute the expected number, i.e. expected fraction $\times$ total number of hypotheses in the version space.

We note that if a teacher provides two counterexamples corresponding to the threshold j of a target hypothesis $h^* \in \mathcal{H}_{\text{threshold}}$, then it eliminates any other hypothesis in the class. Thus, it is straightforward that LwEQ-TD$_{\mathcal{X},\mathcal{H}_{\text{threshold}},h_0}(\ell_{\text{const}})$ = 2. On the other hand, providing just the rightmost example classified as 1 is sufficient for a global learner ℓ_{global}, i.e. LwEQ-TD$_{\mathcal{X},\mathcal{H}_{\text{threshold}},h_0}(\ell_{\text{global}})$ = 1.

In the following theorem, we establish query complexity bounds for worst-TEQ and random-TEQ when teaching a constant query learner ℓ_{const}.

Theorem 6. *Consider the hypothesis space of $\mathcal{H}_{\text{threshold}}$ over the input space $\mathcal{X}$. Fix a constant query learner ℓ_{const}. "In the worst-case", ℓ_{const} queries at least*[5] $\Omega\left(n+1\right)$ *times to worst-*TEQ. *Furthermore, ℓ_{const} queries at most $\mathcal{O}\left(\log(n+1)\right)$ times "in expectation" to random-*TEQ.

Proof. In the worst-case teaching scenario, it is easy to note that worst-TEQ acts as an adversary and provides counterexamples chosen from the leftmost examples. Thus, "in the worst-case" ℓ_{const} queries at least $\Omega\left(n+1\right)$ times to identify the target.

Now, we show the bound for random-TEQ when teaching a constant learner ℓ_{const}. For random-TEQ which provides random counterexamples (sampled uniformly) to the queried hypothesis, we first show that the expected number of counterexamples has the form $\sum_{i=1}^{n} \frac{1}{i}$, which in turn establishes the bound. Consider the random variable X_n to be the number of counterexamples given by random-TEQ for the successful identification of a target hypothesis. Define $W(n)$ to be the expected number of counterexamples provided "in the worst-case" to ℓ_{const} for Threshold functions $\mathcal{H}_{\text{threshold}}$. Thus, we note $W(n) = \mathbb{E}\left[X_n\right]$. Using induction, we show that $W(n) = \sum_{i=1}^{n} \frac{1}{i}$. Note, $W(1) = 1$ and $W(2) = 1 + \frac{1}{2}$. Assume the induction statement for $n = k$. Now, we would prove the statement for $n = k+1$. Consider the following:

$$
\begin{aligned}
& W(k+1) \\
&= \sum_{i=1}^{k+1} i \cdot \mathcal{P}_{\mathcal{U}_{k+1}}(X_{k+1} = i) && (23)\\
&= \sum_{i=1}^{k+1} i \cdot \Big(\mathcal{P}_{\mathcal{U}_{k+1}}(x_1, X_k = i-1) + \mathcal{P}_{\mathcal{U}_{k+1}}(X_k = i)\Big) && (24)\\
&= \frac{1}{k+1} \cdot \sum_{i=1}^{k+1} i \cdot \mathcal{P}_{\mathcal{U}_k}(X_k = i-1) + \sum_{i=1}^{k} i \cdot \mathcal{P}_{\mathcal{U}_{k+1}}(X_k = i) && (25)\\
&= \frac{1}{k+1} \cdot \left(\sum_{i=1}^{k+1} (i-1) \cdot \mathcal{P}_{\mathcal{U}_k}(X_k = i-1) + \sum_{i=1}^{k+1} \mathcal{P}_{\mathcal{U}_k}(X_k = i-1)\right) + \sum_{i=1}^{k} i \cdot \mathcal{P}_{\mathcal{U}_{k+1}}(X_k = i) && (26)\\
&= \frac{1}{k+1} \cdot \sum_{i=1}^{k} i \cdot \mathcal{P}_{\mathcal{U}_k}(X_k = i) + \frac{1}{k+1} \cdot \sum_{i=1}^{k} \mathcal{P}_{\mathcal{U}_k}(X_k = i) + \frac{k}{k+1} \cdot \sum_{i=1}^{k} i \cdot \mathcal{P}_{\mathcal{U}_k}(X_k = i) && (27)\\
&= \frac{1}{k+1} \cdot W(k) + \frac{1}{k+1} \cdot 1 + \frac{k}{k+1} \cdot W(k) && (28)\\
&= W(k) + \frac{1}{k+1} && (29)
\end{aligned}
$$

$\mathcal{U}_k$ denotes the discrete uniform distribution over k samples. Eq. (23) follows using the definition of expectation. In Eq. (24), we decompose based on the first example x_1 is either chosen or not. Eq. (25), Eq. (26), and Eq. (27) follow by switching the probability space from the uniform distribution $\mathcal{U}_{k+1}$ to $\mathcal{U}_k$. In Eq. (28) we apply the induction statement. Thus, we have proven that $W(n) = \sum_{i=1}^{n} \frac{1}{i}$. Since $\sum_{i=1}^{n} \frac{1}{i} \leq \mathcal{O}\left(\log(n+1)\right)$, thus random-TEQ provides at the most $\mathcal{O}\left(\log(n+1)\right)$ counterexamples "in the worst-case" to steer ℓ_{const} to the target hypothesis in $\mathcal{H}_{\text{threshold}}$. □

Now, we discuss the query complexity bounds for a Max-Min query learner. The bounds achieved would be useful in analyzing the case of Axes-aligned hyperplanes in Appendix B.3.

[5] We use the size of the hypothesis class in the bounds.

Min-Max query learner for Threshold functions. To study the query complexity of Threshold functions we construct an elimination graph $G_{\mathsf{elim}}(\mathcal{H}_{\text{threshold}}, \mathcal{U}_n)$ (see **Definition 7** [14]) which is an $(n+1) \times (n+1)$ matrix with rows and columns marked with ordered hypotheses in $\mathcal{H}_{\text{threshold}}$ and every entry corresponding to $h, h' \in \mathcal{H}_{\text{threshold}}$ is $E(h, h')$ (see Appendix B.1). First, we represent the hypothesis class $\mathcal{H}_{\text{threshold}}$ as a boolean matrix $C_{\text{n+1}}$ as

$$
\begin{array}{c}
\\ h_0 \\ h_1 \\ \vdots \\ \vdots \\ \vdots \\ h_{\text{n}}
\end{array}
\begin{array}{c}
\begin{array}{cccccc} x_1 & x_2 & \cdot & \cdot & \cdot & x_n \end{array} \\
\begin{pmatrix}
0 & 0 & \cdots & \cdots & \cdots & 0 \\
1 & 0 & \cdots & \cdots & \cdots & 0 \\
1 & 1 & \cdots & \cdots & \cdots & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots & 0 \\
1 & 1 & 1 & 1 & 1 & 0 \\
1 & 1 & \cdots & \cdots & \cdots & 1
\end{pmatrix}
\end{array}
$$

which has $n+1$ rows for the hypotheses and n columns for the examples. Now, we would construct the **elimination graph** corresponding to $C_{\text{n+1}}$ for uniform distribution over the examples. It is not very difficult to note that the elimination graph $G_{\mathsf{elim}}(C_{\text{n+1}}, \mathcal{U}_n)$ has the form:

$$
\begin{array}{c}
\\ h_0 \\ h_1 \\ h_2 \\ \vdots \\ \vdots \\ h_{\text{n}}
\end{array}
\begin{array}{c}
\begin{array}{cccccc} h_0 & h_1 & h_2 & h_3 & \cdots & h_{\text{n}} \end{array} \\
\begin{pmatrix}
0 & \frac{1}{(n+1)} & \frac{3}{2(n+1)} & \frac{2}{(n+1)} & \cdots & \frac{1}{2} \\
\frac{n}{(n+1)} & 0 & \frac{2}{(n+1)} & \frac{5}{2(n+1)} & \cdots & \frac{n}{2(n-1)} - \frac{1}{(n-1)(n+1)} \\
\frac{(2n-1)}{2(n+1)} & \frac{(n-1)}{(n+1)} & 0 & \frac{3}{(n+1)} & \cdots & \frac{n}{2(n-2)} - \frac{1}{(n-2)(n+1)} \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
\frac{1}{2} & \cdot & \cdots & \cdots & \cdots & 0
\end{pmatrix}
\end{array}
$$

We note that $\ell_{\text{Max-Min}}$ picks the middle threshold function for querying based on Eq. (21). Thus, $\ell_{\text{Max-Min}}$ performs a *binary search* over the set of consistent hypotheses. Thus, in the worst-case teaching scenario $\ell_{\text{Max-Min}}$ queries at the least $\Omega(\log(n+1))$ times. Since binary search algorithm over Threshold functions is optimal in terms of query complexity (use induction on j where $n = 2^j$), thus $\ell_{\text{Max-Min}}$ is an optimal learner for Threshold functions. But we also note that even in the presence of best-TEQ the learner $\ell_{\text{Max-Min}}$ has to query at least twice for the worst-case threshold function, and hence LwEQ-TD$_{\mathcal{X}, \mathcal{H}_{\text{threshold}}, h_0}(\ell_{\text{Max-Min}})$ = 2 as shown for a constant query learner.

Remarks. For the hypothesis class of Threshold functions, there could be ties when finding $\arg\max_{h \in \mathcal{H}(Z_{t-1})} \min_{h' \in \mathcal{H}(Z_{t-1}) \setminus \{h\}} E(h, h')$ (rhs of Eq. (21)). A simple case is when the size n of the input space is 1. At time step $t = 0$, we obtain the (2×2) elimination graph G_{elim} such that $E(h_0, h_1) = E(h_1, h_0) = \frac{1}{2}$. This observation could be generalized for any odd natural number n. In particular, one can show that at time step $t = 0$, both the hypotheses $h_{\frac{n-1}{2}}$ and $h_{\frac{n+1}{2}}$ maximize Eq. (21) such that $E(h_{\frac{n-1}{2}}, h_{\frac{n+1}{2}}) = E(h_{\frac{n+1}{2}}, h_{\frac{n-1}{2}}) = \frac{1}{2}$.

In the case of Threshold functions, the intuition that $\ell_{\text{Max-Min}}$ turns out to be performing binary search is based on the computation of the elimination graph G_{elim}, as shown above. At time step $t = 0$, it is clear that $\ell_{\text{Max-Min}}$ (using Eq. (21)) picks the middle threshold function. For time step $t > 0$, we observe that the version space is a continuous interval of threshold functions (i.e., if h, h' are in the version space then every threshold in between h and h' is also in the version space). Hence, we can again use the computation of the elimination graph as shown for time step $t = 0$. This is the main intuition behind the query.

B.3 Analysis for Axes-aligned Hyperplanes

In this section, we provide the proofs to the query complexity bounds for ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$ under different teaching scenarios as shown in Table 2.

In the rest of the section, we refer a hypothesis $h \in \mathcal{H}_{\text{axh}}$ based on the axis of alignment i and the index j, where $i \in [1, d+1]$, $j \in [n+1]$ such that $\forall \mathbf{x} \in \mathcal{X}, \;\; h(\mathbf{x}) = 1$ if $x_i \leq j$, otherwise 0.

For the sake of continuity, we redefine some of the quantities discussed in Section 4.1. We state the bounds in terms of $|\mathcal{H}_{\text{axh}}| = d \cdot (n+1)$. To capture the global ranking of ℓ_{global} we use the function $g : \mathcal{H}_{\text{axh}} \to [|\mathcal{H}_{\text{axh}}|]$. For a target hypothesis $h^* \in \mathcal{H}_{\text{axh}}$ and ranking function g, the bounds for ℓ_{global} are described using the quantities—$t_g^* := |\{h \in \mathcal{H}_{\text{axh}} \mid g(h) \leq g(h^*)\}|$ (note $0 \leq t_g^* \leq d \cdot (n+1)$). We denote the set $\{h \in \mathcal{H}_{\text{axh}} \mid g(h) \leq g(h^*)\}$ as $\mathcal{H}^*$. Now, we redefine the quantity n_g^* as $n_g^* := \max_{m \in [1,d+1]} |\{h \in \mathcal{H}^* : h \text{ is aligned along axis } m\}|$. Note, $0 \leq n_g^* \leq n+1$.

To show the worst-case teaching bounds we don't impose any restrictions on the learners ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$. On the other hand, our bound for best-case teaching in the case of global learners is for a specific learner ℓ_{global} depending on the target hypothesis h^*. To establish the results, we state the following theorems and provide their proofs[6]: Theorem 7 states the bounds for worst-TEQ, Theorem 8 states the bounds for random-TEQ, and Theorem 9 states the bounds for best-TEQ.

Theorem 7 (Worst-case teaching)**.** *Consider the hypothesis class of Axes-aligned hyperplanes $\mathcal{H}_{\text{axh}}$ (see Eq. (7)). In the LwEQ paradigm, the following bounds on the query complexity hold for worst-*TEQ*:*

1. *For a constant query learner ℓ_{const} the query complexity is lower bounded by $\Omega(d \cdot (n+1))$.*
2. *For a global query learner ℓ_{global} the query complexity is lower bounded by $\Omega(t_g^*)$.*
3. *For a Max-Min learner $\ell_{\text{Max-Min}}$ the query complexity is lower bounded by $\Omega(d \cdot \log(n+1))$.*

Proof. Constant and Global query learners: Consider counterexamples of the form:

$$\mathbf{x} = \Big(0, \cdots, 0, \underbrace{j}_{\text{k-th component}}, 0, \cdots, 0\Big), \;\; \mathbf{x}' = \Big(n, \cdots, n, \underbrace{j}_{\text{k-th component}}, n, \cdots, n\Big). \tag{30}$$

$(\mathbf{x}, 1)$ or $(\mathbf{x}', 0)$ eliminates at most 1 hyperplane aligned along an axis $p \neq k$. Using these counterexamples worst-TEQ provides at least $\Omega(n+1)$ counterexamples "in the worst-case" along an axis (see Threshold functions in Appendix B.2) thus for $\mathcal{H}_{\text{axh}}$ the query complexity is lower bounded by $\Omega(|\mathcal{H}_{\text{axh}}|) = \Omega(d \cdot (n+1))$. Using a similar argument one achieves the lower bound of $\Omega(|\mathcal{H}^*|) = \Omega(t_g^*)$ for global query learner ℓ_{global}.

Max-Min query learner: Fix an arbitrary target hypothesis h^* aligned along an axis i^*. Consider hypotheses $h, h' \in \mathcal{H}_{\text{axh}}$. Assume that h is aligned along axis k and indexed at $i \in [1, n+1]$, whereas h' is aligned along axis k' and indexed at $j \in [1, n+1]$.

To prove the lower bound we show the following: i) worst-TEQ "could" pick counterexamples so that less than *half* of the consistent hypotheses along an axis are eliminated in the version space, and ii) "in the worst-case" $\ell_{\text{Max-Min}}$ has to query so that all the hypotheses aligned along any axis are eliminated in the final version space to locate the target hypothesis.

We note that if $\ell_{\text{Max-Min}}$ queries h for equivalence such that target hypothesis is aligned along a different axis, counterexamples based on Eq. (30) ensure at most one hypothesis aligned along $p \neq k$ is eliminated as well as worst-TEQ could pick $(\mathbf{x}, 1)$ or $(\mathbf{x}', 0)$ such that at most *half* consistent hypotheses aligned along k are eliminated. This completes the proof for i).

Assume that $\ell_{\text{Max-Min}}$ locates the target hypothesis h^* in the query protocol of Algorithm 1 so that $\mathcal{H}_{\text{axh}}(Z)$ contains hypotheses aligned along an axis other than i^* where Z is the history of counterexamples received. We pick such a hypothesis $\hat{h} \in \mathcal{H}_{\text{axh}}(Z)$. Note $\hat{h}$ is consistent with the counterexamples in Z. Thus if $\hat{h}$ is chosen as the "target" hypothesis, the learner $\ell_{\text{Max-Min}}$ would be fooled in querying h^*, and thus the query complexity increases. So, the assumption is invalidated. So, Z must not have any consistent hypothesis aligned along an axis other than i^*. This completes the proof of ii).

Using i) and ii) we show that $\ell_{\text{Max-Min}}$ queries at the least $\Omega(\log(n+1))$[7] times in an axis other than i^*. Since we know that for Threshold functions worst-TEQ provides at the least $\Omega(\log(n+1))$

[6] We skip a theorem for binary-TEQ as the results are self-explanatory.

[7] We use the number of hypotheses aligned along an axis, i.e. $(n+1)$.

counterexamples to a Max-Min learner $\ell_{\text{Max-Min}}$, thus even along the axis i^* (one dimensional threshold functions), the Max-Min learner $\ell_{\text{Max-Min}}$ queries at the least $\Omega\left(\log(n+1)\right)$ times. Hence, we have shown that the query complexity for $\ell_{\text{Max-Min}}$ is lower bounded by $\Omega\left(d \cdot \log(n+1)\right)$. $\square$

In the following, we state the theorem for random-TEQ when teaching ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$.

Theorem 8 (Random teaching)**.** *Consider the hypothesis class of Axes-aligned hyperplanes $\mathcal{H}_{\text{axh}}$ (see Eq. (7)). In the LwEQ paradigm, the following bounds on the query complexity, i.e. the number of counterexamples provided "in expectation" hold for random-*TEQ*:*

1. *For a constant query learner ℓ_{const} the query complexity is upper bounded by $\mathcal{O}\left(d \cdot \log(n+1)\right)$.*
2. *For a global query learner ℓ_{global} the query complexity is upper bounded by $\mathcal{O}\left(d \cdot \log n^*_g\right)$.*
3. *For a Max-Min learner $\ell_{\text{Max-Min}}$ the query complexity is upper bounded by $\mathcal{O}\left(\log(d \cdot (n+1))\right)$.*

Proof. First, we note that the upper bound on the query complexity for a Max-Min learner follows directly using the $\mathcal{O}\left(\log|\mathcal{H}|\right)$ (for finite hypothesis class $\mathcal{H}$) bound on the number of counterexamples in the random case as shown in [14]. So, we provide the proofs for ℓ_{const} and ℓ_{global}.

Constant query learner: To show the upper bound we argue that, if $h \in \mathcal{H}_{\text{axh}}$ aligned along the axis k (dimension), is the next hypothesis picked by the learner ℓ_{const} for equivalence query then the expected number of hypothesis eliminated in the current version space, say H, when random-TEQ provides counterexamples to h is at the least *half* of the consistent hypotheses along the axis k in the version space H. This is sufficient to yield the upper bound.

We argue for $d = 2$. Similar analysis works for $d > 2$ as the probability mass is integrated to the inputs corresponding to the case of $d = 2$. Consider a hypothesis $h' \not\equiv h \in \mathcal{H}_{\text{axh}}$. We note that when h and h' are aligned along the same axis, then using the analysis of Threshold functions (see Appendix B.2), it could be easily shown that "in expectation" a random counterexample to h for h' eliminates at least half of the hyperplanes in the version space. We denote by $\mathbf{R}^k_{h,h'}$ the expected number of consistent hyperplanes eliminated along the axis k when random-TEQ provides counterexamples to h for h' where h' is aligned along an axis other than k. Notice that the set of counterexamples are composed of inputs in two blocks ('+' for label 1 and '-' for label 0) as shown in 2D representation of the hypothesis class below.

$$
i \left(\begin{array}{cc|ccc}
+ & + & & & \\
+ & + & & & \\
 & & & & \\
 & & - & - & - \\
 & & - & - & -
\end{array}\right)
\quad j
$$

We compute $\mathbf{R}^k_{h,h'}$ as follows:

$$
\mathbf{R}^k_{h,h'} = \frac{(n-j)\cdot i \cdot (2n - i + 1) + j \cdot (n-i)\cdot(n+i+1)}{2\left[i\cdot(n-j) + j\cdot(n-i)\right]} \tag{31}
$$

where h and h' are the i-th and j-th indexed hyperplanes in their respective axes. $\mathbf{R}^k_{h,h'}$ Here we show the computation when none of the hyperplanes aligned along an axis are eliminated. Change of parameters achieves the same result in the generic case. We need to show that $\mathbf{R}^k_{h,h'} \geq \frac{n+1}{2}$ irrespective of the choice of i and j. We note that:

$$
\begin{aligned}
& (n-j)\cdot i\cdot(2n-i+1) + j\cdot(n-i)\cdot(n+i+1) \\
= \; & (n-j)\cdot i\cdot(n+1) + (n-j)\cdot i\cdot(n-i) + j\cdot(n-i)\cdot(n+1) + j\cdot(n-i).i \\
\geq \; & (n+1)\cdot i\cdot(n-j) + (n+1)\cdot j\cdot(n-i)
\end{aligned}
$$

Thus, $\mathbf{R}^k_{h,h'} \geq \frac{n+1}{2}$. This implies that if the constant query learner ℓ_{const} receives a counterexample to h then at the least *half* of the consistent hypotheses along k is eliminated from the current version space. Thus, the learner ℓ_{const} queries at most $\mathcal{O}\left(\log(n+1)\right)$ times to random-TEQ to either locate the target or eliminate all the hypotheses along an axis. Since there are d dimensions, thus random-TEQ provides at most $\mathcal{O}(d \cdot \log(n+1))$ counterexamples "in expectation".

Global query learner: Now, we show the query complexity bound for a global query learner ℓ_{global} based on the quantity n_g^*. Using Eq. (31) we note that in expectation at least half of the consistent hypotheses are eliminated along an axis if a random counterexample is provided. Since for any axis k there are at most n_g^* consistent hypotheses in $\mathcal{H}^*$ thus for a global query learner ℓ_{global} random-TEQ provides at the most $\mathcal{O}\left(d \cdot \log n_g^*\right)$ counterexamples "in expectation". □

In the following, we state the theorem for best-TEQ when teaching ℓ_{const}, ℓ_{global}, and $\ell_{\text{Max-Min}}$.

Theorem 9 (Best-case teaching)**.** *Consider the hypothesis class of Axes-aligned hyperplanes $\mathcal{H}_{\text{axh}}$ (see Eq. (7)). In the LwEQ paradigm, the following bounds on LwEQ-TD hold for best-*TEQ*:*

1. *For a constant query learner ℓ_{const}, LwEQ-TD(ℓ_{const}) = 2.*
2. *$\exists$ a global query learner ℓ_{global}, LwEQ-TD(ℓ_{global}) = 1.*
3. *For a Max-Min learner $\ell_{\text{Max-Min}}$, LwEQ-TD($\ell_{\text{Max-Min}}$) = 2.*

Proof. Constant query learner: We note that two counterexamples corresponding to the opposite sides of the target hyperplane are sufficient to fix the target hyperplane in the version space. Assume h^* be the target hypothesis aligned along an arbitrary axis k and indexed at $i \in [1, n+1]$, i.e., $\forall \mathbf{x} \in \mathcal{X}$

$$h^*(\mathbf{x}) = \begin{cases} 1 & \text{if } x_k \leq i \\ 0 & \text{otherwise} \end{cases}$$

Consider $\mathbf{x}', \mathbf{x}'' \in \mathcal{X}$ such that $x'_k = i$, $x'_k = i + 1$, and all the other coordinates are same. Note, $h^*(\mathbf{x}') = 1$ and $h^*(\mathbf{x}'') = 0$. But then any other hypothesis $h' \neq h^* \in \mathcal{H}_{\text{axh}}$ classifies both x', x'' either 0 or 1. Thus, $\mathcal{H}_{\text{axh}}(\{(\mathbf{x}', 1), (\mathbf{x}'', 0)\}) = \{h^*\}$. Since the learner ℓ_{const} arbitraily picks a hypothesis in the version space to query LwEQ-TD(ℓ_{const}) $= 2$.

Global query learner: Consider the global query learner ℓ_{global} with the global ranking $g : \mathcal{H}_{\text{axh}} \to [|\mathcal{H}_{\text{axh}}|]$ as follows: hypothesis which classifies the least number of inputs as negative, i.e. 0 is ranked highest (thus picked if consistent in the version space). Alternatively, for all $h, h' \in \mathcal{H}_{\text{axh}}$, $g(h) \leq g(h')$ if

$$|\{\mathbf{x} : \mathbf{x} \in \mathcal{X}, h(\mathbf{x}) = 1\}| \leq |\{\mathbf{x} : \mathbf{x} \in \mathcal{X}, h'(\mathbf{x}) = 1\}|$$

Now, if the target hypothesis is h^* s.t. it is aligned along the axis $i^* \in [1, d+1]$ and indexed at $j \in [n+1]$, then best-TEQ provides a counterexample $\mathbf{x} := \Big(0, \cdots, 0, \underbrace{j}_{i^*\text{-th component}}, 0, \cdots, 0\Big)$ with label 1. Note, ℓ_{global} picks h^* in the version space $\mathcal{H}_{\text{axh}}(\{(\mathbf{x}, 1)\})$. Hence, the result follows.

Max-Min query learner: We note that similar to a constant query learner, LwEQ-TD($\ell_{\text{Max-Min}}$) $= 2$ as a Max-Min query learner picks the next hypothesis prefixed by the pairwise computation of the expected number of elimination of hypotheses when an equivalence query is posed with hypothesis $h \in \mathcal{H}_{\text{axh}}$ and target hypothesis $h' \in \mathcal{H}_{\text{axh}}$ (see Appendix B.1). As noted in Appendix B.2, $\ell_{\text{Max-Min}}$ would perform a binary search to locate a target hypothesis but then the analysis for a constant query learner ℓ_{const} would apply. Hence, the result follows. □

C Lower and Upper Bounds on Query Complexity (Section 4.2)

In this appendix, we provide the proofs for the bounds shown in Table 3 on the query complexity for different hypothesis classes as discussed in Section 4.2. We divide the appendix in the following way: Appendix C.1 provides the proof to Theorem 1, Appendix C.2 provides the proof to Theorem 2, and Appendix C.3 provides the proof to Theorem 3.

First, we highlight an important theorem from [14] that connects the VC dimension of a hypothesis class and its query complexity for random-TEQ. Our proof for the lower bound on the query complexity in the random teaching scenario for any query learner involves the computation of VC dimension of the underlying hypothesis class.

Theorem 10 (Theorem 25 [14])**.** *If $\mathcal{H}$ is a hypothesis class of VC-dimension d, then any randomized learning algorithm to learn $\mathcal{H}$ must use at least an expected $\Omega\left(d\right)$ equivalence queries with random counterexamples for some target hypothesis.*

Now, we provide the proofs to the theorems in the subsequent appendices.

C.1 Axes-aligned Hyperplanes

In the following, we prove Theorem 1 which establishes lower bounds and upper bound on different teaching scenarios for Axes-aligned hyperplanes. Since the size of $\mathcal{H}_{\text{axh}}$ is $d\cdot(n+1)$ thus any learner queries at least $\Omega\left(d\cdot(n+1)\right)$ times to binary-TEQ.

Proof of Theorem 1. Worst-case teaching: Consider a powerful learner ℓ which performs the following procedure: first finds the dimension (axis) i along which target hypothesis h^* is aligned and then perform a binary search along the axis i. We show that ℓ is optimal and has a lower bound of $\Omega\left(d+\log(n+1)\right)$ on the query complexity for worst-TEQ "in the worst-case".

For the sake of contradiction, assume ℓ' learns the worst-case hypothesis in less than $2d+\log(n+1)$ counterexamples. Consider counterexamples of the form:

$$\mathbf{x} = \Big(0,\cdots,0,\underset{\text{k-th component}}{j},0,\cdots,0\Big),\;\; \mathbf{x}' = \Big(n,\cdots,n,\underset{\text{k-th component}}{j},n,\cdots,n\Big).$$

Notice $\mathcal{H}_{\text{axh}}(\{(\mathbf{x},1)\})$ or $\mathcal{H}_{\text{axh}}(\{(\mathbf{x}',0)\})$ contains at least one hypothesis aligned along axes $p\neq k$. Thus, worst-TEQ provides at the least $\Omega\left(d\right)$ counterexamples, say, $\mathcal{X}_d\subset\mathcal{X}$ such that $\forall h'\in\mathcal{H}_{\text{axh}}(\mathcal{X}_d)$, h' is aligned along the axis i (i.e. for target hypothesis). First, we argue that worst-TEQ could choose to provide at least one counterexample of the form discussed above along each axis. Consider the case when $d=2$. Assume that the target hypothesis is aligned along axis-1 but the learner queries a hypothesis aligned along axis-2. "In the worst-case", $(\mathbf{x},1)$ and $(\mathbf{x}',0)$ (where $x_2=x'_2=j$) forms valid labeled counterexamples to equivalence queries for hypotheses aligned along axis-2 and consistent with target hypothesis which is aligned along axis-1. Using a similar analysis, we generalize this to arbitrary dimension $d>2$. This implies, any learner queries at least one hypothesis along each axis.

On the other hand, for any learner, worst-TEQ provides at the least $\Omega\left(\log(n+1)\right)$ counterexamples "in the worst-case" (see Appendix B.2). Hence, we show that ℓ is optimal and worst-TEQ provides at the least $\Omega\left(d+\log(n+1)\right)$ counterexamples "in the worst-case".

Random teaching: The key to showing the lower bound of $\Omega\left(\log d\right)$ on the query complexity for random-TEQ, i.e number of counterexamples "in expectation", for any learner involves showing a lower bound on the VC dimension of Axes-aligned hyperplanes as defined in Definition 4.1.

Denote the set of combinations of $\log d$ objects as $\mathbb{C}^d$, where $\mathbb{C}^d_i$ represents i-th combination in $\mathbb{C}^d$. Consider the set of $\log d$ inputs $H_{\text{vc}}:=\{\mathbf{x}_1,\mathbf{x}_2,\cdots,\mathbf{x}_{\log d}\}$ such that for every component $i\in[1,d+1]$ the $(\log d)$-tuple $(x_{1,i},x_{2,i},\cdots,x_{\log d,i}) = (1,2,\cdots,\log d)_{\mathbb{C}^d_i}$, i.e. the tuple $(1,2,\cdots,\log d)$ sets index not in $\mathbb{C}^d_i$ as n else other keeps the values. Now, consider a $(\log d)$-tuple $\mathbf{b}$ of boolean values. We note that there exists an axis i and index j such that the corresponding hypothesis h' (i.e. $\forall\mathbf{x}\in\mathcal{X},\;\; h(\mathbf{x})=1$ if $x_i\leq j,$ otherwise 0) classifies H_{vc} as $\mathbf{b}$ (component-wise). This holds because for all $m\in[1,(\log d)+1]$ such that $\mathbf{b}_m=1$ (or $\mathbf{b}_m=0$) (m-th boolean value in $\mathbf{b}$) $x_{m,i}\leq j$ (or $x_{m,i}>j$) for the $(\log d)$-tuple $(x_{1,i},x_{2,i},\cdots,x_{\log d,i})$. This implies that H_{vc} is *shattered* by $\mathcal{H}_{\text{axh}}$. Thus, VCD($\mathcal{H}_{\text{axh}}$) = $\Omega(\log d)$, which directly gives a lower bound on the query complexity for random-TEQ, i.e. on the number of counterexamples provided "in expectation".

Best-case teaching: In Theorem 9 (see Appendix B.3) we show there exists a global query learner ℓ_{global} such that the query complexity for best-TEQ is 1, i.e. LwEQ-TD$_{\mathcal{X},\mathcal{H}_{\text{axh}},h_0} = 1$. □

C.2 Monotone Monomials

Now, we proof Theorem 2 which establishes lower bounds and upper bound on query complexity in different teaching scenarios for Monotone monomials. For the sake of clarity, we rewrite Eq. (8) here

$$\mathcal{H}_{\text{mono}} := \{h \mid \exists\ \{i_1, i_2, \cdots, i_k\} \subseteq \{1, 2, \cdots, n\},\ \text{s.t. } \forall \mathbf{x} \in \mathcal{X},\ \ h(\mathbf{x}) = x_{i_1} \wedge \cdots \wedge x_{i_k}\}.$$

We note that $|\mathcal{H}_{\text{mono}}| = 2^n$. Thus, any learner queries at least $\Omega(2^n)$ times to binary-TEQ. Furthermore, we note that any hypothesis $h \in \mathcal{H}_{\text{mono}}$ represented as $h(\mathbf{x}) := x_{i_1} \wedge x_{i_2} \wedge \cdots \wedge x_{i_k}$ for some $\{i_1, i_2, \cdots, i_k\} \subseteq \{1, 2, \cdots, n\}$ is "identified" with an input $\mathbf{x}' \in \mathcal{X}$ such that for all $m \in \{i_1, i_2, \cdots, i_k\}$, $x'_m = 1$ otherwise 0.

Proof of Theorem 2. Worst-case teaching: Consider a global learner ℓ which ranks hypothesis based on the number of "dependent" literals, i.e. for all $h, h' \in \mathcal{H}_{\text{mono}}$, $g(h) \leq g(h')$ (where g is the global ranking for ℓ) if h has at least as many dependent literals as h'. We show that ℓ is optimal and has a lower bound of $\Omega(n)$ on the query complexity for worst-TEQ "in the worst-case".

First, we show that ℓ asks at least $\Omega(n)$ queries to worst-TEQ "in the worst-case". We use the representation of a monomial in terms of its set of literals. Fix a starting hypothesis $h_0 := \text{mono}(j_1, j_2, \cdots, j_m)$. Consider the following illustration for query protocol of Algorithm 1:

$$\underbrace{\text{mono}(j_1, j_2, \cdots, j_m)}_{\text{starting monomial } h_0} \xrightarrow{t=0} \underbrace{\text{mono}(1, 2, \cdots, n)}_{\ell \text{ queries for equivalence}} \xrightarrow{t=1} \cdots \xrightarrow{t=t'} \underbrace{\text{mono}(i_1, i_2, \cdots, i_k)}_{\text{target monomial}}$$

Notice that the counterexample $(\mathbf{x}_1, y_1)$ to $\text{mono}(1, 2, \cdots, n)$ has to be a positive example (i.e, $h^*(\mathbf{x}_1) = y_1 = 1$) as $\{i_1, i_2, \cdots, i_k\} \subset \{1, 2, \cdots, n\}$. Also, worst-TEQ picks $\mathbf{x}_1$ such that the number of 1's in the string is $n-1$ to maximize the query complexity. Thus, it is straightforward that at each time step worst-TEQ provides a counterexample such that the subset of literals of the corresponding hypothesis (or monomial) is a superset of $\{i_1, i_2, \cdots, i_k\}$ but has size 1 less than the subset size for the current hypothesis (monomial). Essentially, the set $\{1, 2, \cdots, n\} \setminus \{i_1, i_2, \cdots, i_k\}$ is eliminated before the learner ℓ queries h^* for equivalence. This implies that ℓ performs at least $\Omega(n)$ queries to identify a target "in the worst-case" if worst-TEQ provides counterexamples.

Now, we argue there doesn't exist a learner ℓ' with a better query complexity. For the sake of contradiction, assume that some learner ℓ' achieves a query complexity less than n for worst-TEQ. Note that ℓ' can't query the target hypothesis, at some time step, such that there exists a monomial $\text{mono}(i_1, i_2, \cdots, i_k, i_{k+1})$ where $\{i_1, i_2, \cdots, i_k\} \subsetneq \{i_1, i_2, \cdots, i_k, i_{k+1}\}$ in the current version space. This holds because, if Z' is the current history of counterexamples provided by ℓ', then Z' is also a valid set of counterexamples for $\text{mono}\{i_1, i_2, \cdots, i_k, i_{k+1}\}$ as a target hypothesis. Thus, worst-TEQ fools the learner ℓ' in this case. So, ℓ' queries monomials in a way to eliminate at least all the supersets of $\{i_1, i_2, \cdots, i_k\}$. But if a monomial corresponding to a superset, say $\text{mono}\{j_1, j_2, \cdots, j_{k'}\}$, is queried for equivalence, worst-TEQ could follow the strategy described for the learner ℓ. So, monomials with the set of literals, say S' such that $\{i_1, i_2, \cdots, i_k\} \subseteq S' \subseteq \{j_1, j_2, \cdots, j_{k'}\}$ remain consistent. Since "in the worst-case" ℓ' has to query successively to reduce the size of the set of literals S' of the queried monomial, ℓ' performs at least $\Omega(n)$ queries. So, our assumption on the query complexity for worst-TEQ teaching ℓ' is wrong, thus ℓ is optimal.

Random-case teaching: The key to showing the lower bound of $\Omega(n)$ on the query complexity for random-TEQ, i.e number of counterexamples "in expectation", for any learner involves showing a lower bound on the VC dimension of the hypothesis class of Monotone monomials as defined in Definition 4.2. Consider the set M_{vc} defined as:

$$M_{\text{vc}} := \{\mathbf{x} \in \mathcal{X} : \exists\, !i \in [1, n+1] \text{ s.t. } x_i = 0\}$$

We show that M_{vc} is *shattered* by $\mathcal{H}_{\text{mono}}$. We note that for all $h \not\equiv h' \in \mathcal{H}_{\text{mono}}$, $h(M_{\text{vc}}) \neq$ [8] $h'(M_{\text{vc}})$. To show this consider a literal $v_{i'}$ present in h but not in h'. Now, for $\mathbf{x} \in M_{\text{vc}}$ such that $x_{i'} = 0$,

[8] We define h on a subset of $\mathcal{X}$ as a tuple of boolean values.

$h(\mathbf{x}) = 0$ but $h'(\mathbf{x}) = 1$. Since $|\mathcal{H}_{\text{mono}}| = 2^n$ and $M_{\text{vc}} = n$, we have shown that M_{vc} is shattered by $\mathcal{H}_{\text{mono}}$, leading to VCD($\mathcal{H}_{\text{mono}}$) = n. Thus, for any optimal learner the lower bound on query complexity for random-TEQ, i.e. on the number of counterexamples provided "in expectation" is $\Omega(n)$.

Best-case teaching: Consider the global query learner ℓ_{global} with the global ranking $g : \mathcal{H}_{\text{mono}} \to [|\mathcal{H}_{\text{mono}}|]$ as follows: ranks hypothesis based on the number of "dependent" literals, i.e. if $h, h' \in \mathcal{H}_{\text{mono}}$ then $g(h) \leq g(h')$ if h has at least as many dependent literals as h'. At the beginning, ℓ_{global} queries $h(\mathbf{x}) \equiv x_1 \wedge x_1 \wedge \cdots \wedge x_n$. Then the optimal teacher best-TEQ provides $(\mathbf{x}', 1)$ as a counterexample where $\mathbf{x}'$ identifies h^* as discussed above. Now, the highest ranked monomial in $\mathcal{H}_{\text{mono}}(\{(\mathbf{x}', 1)\})$ is h^* which the learner ℓ_{global} queries for equivalence. Since best-TEQ provides exactly 1 counterexample the worst-case hypothesis in $\mathcal{H}_{\text{mono}}$, thus LwEQ-TD$_{\mathcal{X}, \mathcal{H}_{\text{mono}}, h_0}(\ell_{\text{global}}) = 1$ (for a fixed $h_0 \in \mathcal{H}_{\text{mono}}$). Hence, the result follows. □

C.3 Orthogonal Rectangles

Now, we prove Theorem 3 which establishes lower bounds and upper bound on query complexity in different teaching scenarios for Orthogonal rectangles. We note that the size of $\mathcal{H}_{\text{rec}}$ is of the order $\Omega\left((n \cdot (n+1))^d\right)$, thus any learner queries at least $\Omega\left((n \cdot (n+1))^d\right)$ times to binary-TEQ.

Proof of Theorem 3. Worst-case teaching: Similar to the case of worst-case teaching for Axes-aligned hyperplanes (see Definition 4.1), we show that any optimal learner has a lower bound of $\Omega(d \cdot \log n)$ on the query complexity for worst-TEQ "in the worst-case".

Consider a target hypothesis $h^* \in \mathcal{H}_{\text{rec}}$ such that for all $j \in [1, d+1], a_j = 0$. The choice of b_j is shown later. Now, consider counterexamples of the form:

$$\mathbf{x} = \Big(0, \cdots, 0, \underbrace{m}_{\text{k-th component}}, 0, \cdots, 0\Big).$$

Notice that labeled counterexamples $(\mathbf{x}, 0)$ or $(\mathbf{x}, 1)$ to $h' \in \mathcal{H}_{\text{rec}}$ doesn't affect the choice of b_j (doesn't get fixed for consistent hypotheses in the version space) for all $j \ (\neq k) \in [1, d+1]$. This implies that worst-TEQ could provide counterexamples such that any learner has to query hypotheses for equivalence along each axis $i \in [1, d+1]$. Using the analysis of Threshold functions (see Appendix B.2) we know that the worst-TEQ provides at least $\Omega(\log(n+1))$ counterexamples to any learner along an axis. but, we could pick worst-case b_j's for a target hypothesis h^* such that worst-TEQ provides at least $\Omega(\log(n+1))$ counterexamples in every axis, implying that for any learner worst-TEQ provides at the least $\Omega(d \cdot \log(n+1))$ counterexamples "in the worst-case".

Random teaching: The key to showing the lower bound of $\Omega(d)$ on the query complexity for random-TEQ i.e number of counterexamples "in expectation", for any learner involves showing a lower bound on the VC dimension of the hypothesis class of Orthogonal rectangles as defined in Definition 4.3. Consider the set O_{vc} defined as:

$$O_{\text{vc}} := \{\mathbf{x} \in \mathcal{X} : \exists i \in [1, d+1], \text{ s.t. } x_i = \lfloor n/2 \rfloor \text{ and } \forall j \neq i, x_j = 0\}$$

We note that $|O_{\text{vc}}| = d$. Now, we show that O_{vc} is shattered by $\mathcal{H}_{\text{rec}}$. Notice that for any d-tuple $\mathbf{b}$ of boolean values, one can fix the choices of $\{a_j, b_j\}_{j \in [1, d+1]} \subset [n+1]$ to obtain a hypothesis $h \in \mathcal{H}_{\text{rec}}$ where $a_j = 0$, and $b_j = \lfloor n/2 \rfloor$ if $\mathbf{b}_i = 1$ else $b_j = \lfloor n/2 \rfloor - 1$. This gives $h(O_{\text{vc}}) = \mathbf{b}$. So, we have shown that O_{vc} is *shattered* by $\mathcal{H}_{\text{rec}}$, which implies VCD($\mathcal{H}_{\text{rec}}$) = $\Omega(d)$. Thus, for any optimal learner the lower bound on the query complexity for random-TEQ, i.e. on the number of counterexamples provided "in expectation" is $\Omega(d)$.

Best-case teaching: We show that there is a global query learner ℓ_{global} which achieves the required query complexity bound of 2. In order to establish this, we note for fixed $h_0, \mathcal{X}, \mathcal{H}$, LwEQ-TD$_{\mathcal{X}, \mathcal{H}, h_0}(\Sigma_{\sigma_{\text{global}}})$ = RTD($\mathcal{H}$) (see Theorem 4, Section 5) for the family of ranking functions $\Sigma_{\sigma_{\text{global}}}$ (see Section 5) which induce global query learners. We need to show that RTD($\mathcal{H}_{\text{rec}}$) = 2. [25] showed that the TD (teaching dimension) notion of preference-based teaching dimension (PBTD) of a finite hypothesis class $\mathcal{H}$ is the same as the RTD of the class $\mathcal{H}$ (see **Corollary 9** [25]). On the other hand, [61] showed that PBTD($\mathcal{H}_{\text{rec}}$) = 2 (see **Example 1** [61]). Thus, there exists a global query learner ℓ_{global} such that LwEQ-TD$_{\mathcal{X}, \mathcal{H}_{\text{rec}}, h_0}(\ell_{\text{global}}) = 2$. □

D Teaching Dimensions for the LwEQ and LfS Paradigms (Section 5)

In this appendix, we prove our main results for Section 5, i.e., Theorem 4 and Theorem 5, which establish connections of LwEQ-TD in the learning-with-equivalence-queries (LwEQ) paradigm to existing notions of TD in the learning-from-samples (LfS) paradigm. We divide the appendix based on the connections established for learner types induced by a specific family of ranking functions.

First, we note that the framework of ranking functions (see Eq. (10), Section 5) leads to the definition of a preference function as discussed in [29]. Thus, a ranking function σ could be, alternatively, used as a preference function to define LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma)$ (for fixed $\mathcal{X},\mathcal{H},h_0$) for a learner induced by a preference function σ as shown in [29] in the LfS paradigm. We illustrate this as follows:

σ (ranking function) $\rightsquigarrow$ σ (preference function)

$\uparrow$ $\downarrow$

Eq. (10) (LwEQ paradigm) **Protocol 1** [29] (LfS paradigm)

With this understanding, we would interchangeably talk about teaching settings in the LfS paradigm for learner models (in the LfS paradigm) induced by a framework of ranking functions. Our proof technique uses the unification of teaching settings in the LfS paradigm [29].

D.1 LwEQ Learners with Ranking Functions Independent of Z_{t-1} and h_{t-1}

For fixed $\mathcal{X},\mathcal{H},h_0$, we prove Theorem 4 establishing connections for LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{const}})$ and LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{global}})$.

Proof of Theorem 4. First, we show the connection for the family of *constant* learners Σ_{const}. We note that any sequential teaching set $Z \subseteq \mathcal{Z}$ in the LfS paradigm is permutation invariant. For arbitrary h_0 and h^*, if $Z := \{z_1, z_2, \cdots, z_l\}$ is the minimal teaching set for teaching in the LfS paradigm then any permutation of $z_1, z_2, \cdots, z_l$ is also a teaching set (Σ_{const} leads to a batched teaching setting in the LfS paradigm). Since all the hypotheses are preferred equally, they are eliminated in the teaching protocol at some time step. This implies that in the query protocol of Algorithm 1 in the LwEQ paradigm, at any step t, the teacher could pick $z_t \in Z$ such that z_t forms a counterexample to the queried hypothesis $h_t \in \ell_{\sigma_{\text{const}}}(Z_{t-1}, h_{t-1})$. Thus, the teacher could always permute Z to respond as counterexamples in Algorithm 1 so as to steer the learner from h_0 to h^*. Since h_0 and h^* are arbitrary, thus LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}}) \leq$ LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}})$.

Because any history of counterexamples $Z \subseteq \mathcal{Z}$ used in the query protocol of Algorithm 1 to steer the learner from h_0 to h^*, forms a solution for teaching in the LfS paradigm, we have LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}}) =$ LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}})$. Now, we note the following:

$$\min_{\sigma_{\text{const}} \in \Sigma_{\text{const}}} \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}}) = \min_{\sigma_{\text{const}} \in \Sigma_{\text{const}}} \text{LfS-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}}) \tag{32}$$

$$\implies \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{const}}) = \text{LfS-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{const}}) \tag{33}$$

Eq. (32) is straightforward since for any $\sigma_{\text{const}} \in \Sigma_{\text{const}}$, LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}}) =$ LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}})$. Eq. (33) follows using the definition of LwEQ-TD for a family of query functions Σ_{const} and the definition of LfS-TD for a family of preference functions Σ_{const} (see [29]). Now, we note that using [29] (see Theorem 1) shows that LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{const}}) =$ wc-TD$(\mathcal{H})$, which implies LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{const}}) =$ wc-TD$(\mathcal{H})$.

Now, for a *global* ranking function $\sigma_{\text{global}} \in \Sigma_{\text{global}}$ we note that any hypothesis which is not ranked strictly over the target hypothesis h^* is never queried for equivalence. Thus, $H^* = \{h \in \mathcal{H} \mid g(h) \leq g(h^*)\}$ where g is the global function for σ_{global}, needs to be eliminated both in the query protocol in the LwEQ paradigm and the teaching protocol in the LfS paradigm. Using similar ideas as discussed above for *constant* query learners, we note that LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{global}}) =$ LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{global}})$, and hence LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{global}}) =$ LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{global}})$. Since LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{global}}) =$ RTD$(\mathcal{H})$ (Theorem 1, [29]), we show LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{global}}) =$ RTD$(\mathcal{H})$. $\square$

Remark 1. When $\mathcal{H}$ is singleton, LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\ell) = 0$ as per the query complexity defined in Eq. (5) for a query learner ℓ. In the proof above, we implicitly assumed that $\mathcal{H}$ is not singleton.

Remark 2. Note that if there exists a target hypothesis $h^* \in \mathcal{H}$ for which $\ell(\emptyset, h_0) = \{h^*\}$, then LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\ell, h^*) = 0$ as per the query complexity defined in Eq. (4); however, the teaching complexity in the LfS paradigm (see **Protocol 1** [29]) leads to LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\ell, h^*) = 1$. Since we are interested in analyzing $\max_{h^*\in\mathcal{H}}$ LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\ell, h^*)$, the proofs focus on $h^* \in \mathcal{H}$ for which LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\ell, h^*) > 0$.

D.2 LwEQ Learners with Ranking Functions Dependent on Z_{t-1} and/or h_{t-1}

For fixed $\mathcal{X}, \mathcal{H}, h_0$, we prove Theorem 5 establishing connections for LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{gvs}})$, LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{local}})$, and LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{lvs}})$.

Proof of Theorem 5. *Connection for $\Sigma^{\text{CF}}_{\text{gvs}}$*. Fix an arbitrary global version space learner characterized with a ranking function $\sigma_{\text{gvs}} \in \Sigma^{\text{CF}}_{\text{gvs}}$. Furthermore, assume that g identifies σ_{gvs} in $\Sigma^{\text{CF}}_{\text{gvs}}$ (see Eq. (15), Section 5.2).

First, we note that any history of counterexamples $Z \subseteq \mathcal{Z}$ used in the query protocol of Algorithm 1 to steer the learner from h_0 to h^*, forms a solution for teaching in the LfS paradigm for the preference function σ_{gvs} (see **Protocol 1** [29]). Denote $|Z|$ by k and $Z_k := Z$. Now, we know that $\ell_{\sigma_{\text{gvs}}}(Z_k, h_k) = \arg\min_{h'\in\mathcal{H}(\{Z_k\})} \sigma_{\text{gvs}}(h'; \mathcal{H}(\{Z_k\}), h_k) = \{h^*\}$. But then for all h, $\arg\min_{h'\in\mathcal{H}(\{Z_k\})} \sigma(h'; \mathcal{H}(\{Z_k\}), h) = \arg\min_{h'\in\mathcal{H}(\{Z_k\})} g(h'; \mathcal{H}(\{Z_k\}))$. This implies that Z forms a teaching set in the teaching protocol as discussed in [29] in the LfS paradigm. Since h^* is arbitrarily picked, we have LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{gvs}}) \leq$ LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{gvs}})$.

Assume $Z' := \{z_1, z_2, \cdots, z_l\}$ is the minimal teaching set for teaching in the LfS paradigm. The optimal teacher (LwEQ paradigm) could pick a permutation of Z' to provide as counterexamples in the LwEQ paradigm. At time t, if Z_{t-1} is the current history of counterexamples (assuming $Z_{t-1} \subseteq Z'$) and h_{t-1} is the current hypothesis, then in the query protocol of Algorithm 1 the learner picks the next hypothesis as follows

$$h_t \in \underset{h'\in\mathcal{H}(Z_{t-1})}{\arg\min}\ \sigma_{\text{gvs}}(h'; \mathcal{H}(Z_{t-1}), h_{t-1}) = \underset{h'\in\mathcal{H}(Z_{t-1})}{\arg\min}\ g(h', \mathcal{H}(Z_{t-1}))$$

Notice that either $i)$ $h_t = h^*$ or $ii)$ $\exists z \in Z'$ s.t. z is a counterexample to h_t. Assume that neither $i)$ nor $ii)$ hold. Then $\nexists z \in Z'$ such that z is a counterexample to h_t. This contradicts the collusion-freeness of σ_{gvs} (see Definition 5.1) because

$$\underset{h'\in\mathcal{H}(Z_{t-1})}{\arg\min}\ \sigma_{\text{gvs}}(h'; \mathcal{H}(Z_{t-1}), h_{t-1}) = \{h_t\},$$

but

$$\underset{h'\in\mathcal{H}(Z_{t-1}\cup(Z'\setminus Z_{t-1}))}{\arg\min}\ \sigma_{\text{gvs}}(h'; \mathcal{H}(Z'), h_t) = \{h^*\}$$

which contradicts the fact that h_t is consistent with $Z' \setminus Z_{t-1}$. Hence, LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{gvs}}) \leq$ LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{gvs}})$. Since we picked σ_{gvs} arbitrarily, we have

$$\min_{\sigma_{\text{gvs}}\in\Sigma^{\text{CF}}_{\text{gvs}}} \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{gvs}}) = \min_{\sigma_{\text{gvs}}\in\Sigma^{\text{CF}}_{\text{gvs}}} \text{LfS-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{gvs}})$$

$$\overset{\text{(Eq. (6), Section 3)}}{\Longrightarrow} \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{gvs}}) = \text{LfS-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{gvs}})$$

$$\overset{\text{(Theorem 1 [29])}}{\Longrightarrow} \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{gvs}}) = \text{NCTD}(\mathcal{H})$$

The last equation yields the desired result.

Connection for $\Sigma^{\text{CF}}_{\text{local}}$. First, we prove the lower bound on LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{local}})$ as stated in Eq. (19). To show the bound, we note that LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma^{\text{CF}}_{\text{local}})$:= local-PBTD$_{\mathcal{X},\mathcal{H},h_0}$ (see [29]).

Fix h_0 and h^* as the starting and target hypotheses. Now, fix the best ranking function $\sigma^*_{\text{local}} := \arg\min_{\sigma_{\text{local}}\in\Sigma^{\text{CF}}_{\text{local}}}$ LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{local}})$, i.e. the best ranking function which minimizes the

LwEQ-TD for the family of ranking functions $\Sigma_{\text{local}}^{\text{CF}}$. Assume the optimal teacher provides the following counterexamples to the learner $\ell_{\sigma_{\text{local}}^*}$ in the query protocol of Algorithm 1:

$$h_0 \xrightarrow{\emptyset} h_1 \xrightarrow{Z_1} h_2 \cdots \xrightarrow{Z_{k+1}} h^* \tag{34}$$

where $Z_i := Z_{i-1} \cup \{z_{i-1}\}$ for the counterexample z_{i-1} to h_{i-1}. We show that there exists a choice of a local ranking function $\sigma'_{\text{local}} \in \Sigma_{\text{local}}^{\text{CF}}$ such that the history of counterexamples Z_{k+1} provided by the optimal teacher in Algorithm 1 as shown above forms a sequence of teaching examples for a learner characterized with the local ranking σ'_{local} in the LfS paradigm (see **Protocol 1** [29]).

Define σ'_{local} as follows: for all $h' \in \mathcal{H}(\{z_1\})$, $\sigma'_{\text{local}}(h'; \mathcal{H}(\{z_1\}), h_0) = \sigma^*_{\text{local}}(h'; \mathcal{H}(\{z_1\}), h_1)$ (see Eq. (34)) where z_1 is the counterexample provided by the optimal teacher to h_1 (hypothesis queried for equivalence with the current history of examples $\emptyset$, and the current hypothesis h_0) for the learner characterized by the ranking function σ^*_{local} in the query protocol of Algorithm 1 as described above; otherwise $\sigma'_{\text{local}} \equiv \sigma^*_{\text{local}}$. Thus, $\arg\min_{h' \in \mathcal{H}(\{z_1\})} \sigma'_{\text{local}}(h'; \mathcal{H}(\{z_1\}), h_0) = \{h_2\}$ (see Eq. (34)). Since h^* is arbitrary, thus σ'_{local} could be defined over all the triplets (h^*, z_1, h_2), i.e. for any choice of target h^*, as shown in Eq. (34). Notice that σ'_{local} is a valid local ranking function as σ^*_{local} is a local ranking function. Similarly, the collusion-freeness of σ'_{local} follows from its definition and the collusion-freeness of σ^*_{local}.

For the learner characterized by σ'_{local} in the LfS paradigm, the optimal teacher could provide the counterexamples Z_{k+1} sequentially for teaching in **Protocol 1** [29]. This is illustrated as follows:

$$h_0 \xrightarrow{Z_1} h_2 \xrightarrow{Z_2} h_3 \cdots \xrightarrow{Z_{k+1}} h^* \tag{35}$$

Upon receiving the examples z_i, the learner picks the next hypothesis h_{i+1} as shown above. Thus, Z_{k+1} forms a teaching set in the LfS paradigm. This implies that LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma'_{\text{local}}, h^*) \le |Z_{k+1}|$. Since h^* is arbitrarily chosen, thus we have LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma'_{\text{local}}) \le$ LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma^*_{\text{local}})$. But this implies

$$\min_{\sigma_{\text{local}} \in \Sigma_{\text{local}}^{\text{CF}}} \text{LfS-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{local}}) \le \text{LfS-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma'_{\text{local}}) \le \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma^*_{\text{local}})$$

$$\overset{(\sigma^*_{\text{local}} \text{ minimizes LwEQ-TD}(\cdot))}{\Longrightarrow} \text{LfS-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{local}}^{\text{CF}}) \le \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{local}}^{\text{CF}})$$

Thus, from the last equation we have local-PBTD$_{\mathcal{X},\mathcal{H},h_0} \le$ LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{local}}^{\text{CF}})$.

The upper bound on LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{local}})$ follows by noting that LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{local}}) =$ wc-TD$(\mathcal{H})$ for a ranking function $\sigma_{\text{const}} \in \Sigma_{\text{const}}$ (as shown in Theorem 4) and that *constant* query learner could be adversarial compared to the *local* learners.

We would show the inequality using induction. Fix a local ranking function $\sigma_{\text{local}} \in \Sigma_{\text{local}}^{\text{CF}}$, and a constant ranking function $\sigma_{\text{const}} \in \Sigma_{\text{const}}$. Now, consider arbitrary h_0 and h^* as the starting and target hypotheses respectively. The induction statement is on the worst-case optimal teaching sequence size for steering the learner from a starting to a target hypothesis when the learner is characterized with a constant ranking function. The inequality holds for $D_{\sigma_{\text{const}}}(Z \subseteq \mathcal{Z}, h_0, h^*) = 0$. By definition, for any $Z \subseteq \mathcal{Z}, h_0, h^*$, we have $D_{\sigma_{\text{local}}}(Z, h_0, h^*) \le D_{\sigma_{\text{const}}}(Z, h_0, h^*) = 0$ as $\ell_{\sigma_{\text{local}}}(Z, h_0) = \ell_{\sigma_{\text{const}}}(Z, h_0) = \{h^*\}$. Following the induction statement, assume that $D_{\sigma_{\text{local}}}(Z \subseteq \mathcal{Z}, h_0, h^*) \le D_{\sigma_{\text{const}}}(Z \subseteq \mathcal{Z}, h_0, h^*)$ whenever $D_{\sigma_{\text{const}}}(Z \subseteq \mathcal{Z}, h_0, h^*) \le k$. Now, we need to show the inequality when $D_{\sigma_{\text{const}}}(\emptyset, h_0, h^*) = k+1$ (similar argument holds for $Z \subseteq \mathcal{Z}$). We unfold the recursion for $D_{\sigma_{\text{const}}}$:

$$D_{\sigma_{\text{const}}}(\emptyset, h_0, h^*) = 1 + \max_{h' \in \ell_{\sigma_{\text{const}}}(\emptyset, h_0)} \min_{z: h'(x_z) \neq y_z} D_{\sigma_{\text{const}}}(\{z\}, h', h^*) \tag{36}$$

We note that for all $h' \in \ell_{\sigma_{\text{const}}}(\emptyset, h_0)$,

$$\min_{z: h'(x_z) \neq y_z} D_{\sigma_{\text{const}}}(\{z\}, h', h^*) \ge \min_{z: h'(x_z) \neq y_z} D_{\sigma_{\text{local}}}(\{z\}, h', h^*).$$

Since $\ell_{\sigma_{\text{local}}}(\emptyset, h_0) \subseteq \ell_{\sigma_{\text{const}}}(\emptyset, h_0)$ we have:

$$\max_{h' \in \ell_{\sigma_{\text{const}}}(\emptyset, h_0)} \min_{z: h'(x_z) \neq y_z} D_{\sigma_{\text{const}}}(\{z\}, h', h^*) \ge \max_{h' \in \ell_{\sigma_{\text{local}}}(\emptyset, h_0)} \min_{z: h'(x_z) \neq y_z} D_{\sigma_{\text{local}}}(\{z\}, h', h^*)$$

Plugging this into Eq. (36) implies

$$D_{\sigma_{\text{const}}}(\emptyset, h_0, h^*) \geq D_{\sigma_{\text{local}}}(\emptyset, h_0, h^*)$$

Since h^* is arbitrary we get LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}}) \geq$ LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{local}})$ for arbitrary σ_{const} and σ_{local}. Using the definition of LwEQ-TD for a family of ranking functions (Eq. (6) and Section 5), we get

$$\min_{\sigma_{\text{const}} \in \Sigma_{\text{const}}} \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{const}}) \geq \min_{\sigma_{\text{local}} \in \Sigma_{\text{local}}^{\text{CF}}} \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{local}})$$

Hence we get LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{local}}^{\text{CF}}) \leq$ wc-TD$(\mathcal{H})$ using Theorem 4.

Connection for $\Sigma_{\text{lvs}}^{\text{CF}}$. Using similar arguments as used to show the bounds in Eq. (19) for the family of local ranking functions $\Sigma_{\text{local}}^{\text{CF}}$, the bounds in Eq. (20) follow as well. □

Is there a gap between lvs-PBTD$_{\mathcal{X},\mathcal{H},h_0}$ and LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{lvs}}^{\text{CF}})$? In Theorem 5, we showed that lvs-PBTD$_{\mathcal{X},\mathcal{H},h_0} \leq$ LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{lvs}}^{\text{CF}})$. An interesting research question could be to understand the gap between lvs-PBTD$_{\mathcal{X},\mathcal{H},h_0}$ and LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{lvs}}^{\text{CF}})$. We answer this question partially and leave the unresolved part for future work. We present a problem instance that suggests that there could be a strict gap between lvs-PBTD$_{\mathcal{X},\mathcal{H},h_0}$ and LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\Sigma_{\text{lvs}}^{\text{CF}})$.

Theorem 11. *For learners whose query function is induced by a ranking function dependent on Z_{t-1} and/or h_{t-1}, there exists a problem instance of $\mathcal{X}, \mathcal{H}, h_0, \sigma_{\text{lvs}} \in \Sigma_{\text{lvs}}^{\text{CF}}$, such that*

$$\text{LfS-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{lvs}}) \ll \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{lvs}})$$

i.e., LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{lvs}})$ *is much lower than* LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{lvs}})$.

Proof. Consider a d-dimensional lattice of finite size n (isomorphic to a cube in $\mathbb{R}^d$ of length n with positive integer coordinates). $\mathcal{H}$ and $\mathcal{X}$ correspond to the nodes of the lattice. The hypothesis h_v corresponding to node v is identified as one which classifies $v' \neq v$ as positive and v as negative. We consider learners characterized with the ranking function $\sigma_{\text{lvs}} \in \Sigma_{\text{lvs}}^{\text{CF}}$ such that it moves to a close-by hypothesis measured via ℓ_1 (Manhattan) distance. It is not very difficult to note that LwEQ-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{lvs}})$ is at least $\Omega\left(n^d\right)$.

Now, we argue that LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{lvs}})$ is upper bounded by $\mathcal{O}\left(n \cdot d^2\right)$. Consider the hypotheses[9] $h_0 = \Big(\underbrace{0,0,0,0,\cdots 0,0,0,0}_{d}\Big)$ and $h^* = \Big(\underbrace{n,n,n,n,\cdots n,n,n,n}_{d}\Big)$. The teacher could try to steer the learner in the following sequential form:

$$h_0 \dashrightarrow h_1 \dashrightarrow h_2 \dashrightarrow \cdots \dashrightarrow h_i \dashrightarrow \cdots \dashrightarrow h_{d-1} \dashrightarrow h_d = h^*$$

where $h_i = \Big(\underbrace{n,n,n,n,\cdots,n,n,n}_{i},0,0,\cdots,0\Big)$. We show that with $\mathcal{O}\left(n \cdot d\right)$ teaching examples the learner could be taught to move from h_i to h_{i+1} for any $i \in [n]$. Notice that the learner prefers to stay on h_v for a node unless v is provided as a positive example to steer it away. Thus, the teacher first provides $\left\{(0,i_1,i_2,i_3,\cdots,i_{d-1}) \,\middle|\, \sum_{j=1}^{d-1} i_j \leq 1 \text{ s.t. } \exists\, !k \; i_k \neq 0\right\}$. Then, the teacher provides $\{(a,0,0,\cdots,0)\}_{a=1}^{n}$ in the sequential order. Therefore, in $\mathcal{O}\left(n \cdot d\right)$ examples the learner moves from h_0 to h_1. Using a similar strategy, the teacher requires at most $\mathcal{O}\left(n \cdot d^2\right)$ examples ($\mathcal{O}\left(n \cdot d\right)$ for each dimension) to steer the learner to the target h^*. It can be easily verified that the same holds for any h_0 and h^* and thus LfS-TD$_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{lvs}})$ is $\mathcal{O}\left(n \cdot d^2\right)$. Hence, for the hypothesis class $\mathcal{H}$ of a lattice in $\mathbb{R}^d$ of size n with the input space $\mathcal{X} := \{1,2,\cdots,n\}^d$ we get,

$$\mathcal{O}\left(n \cdot d^2\right) = \text{LfS-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{lvs}}) \ll \text{LwEQ-TD}_{\mathcal{X},\mathcal{H},h_0}(\sigma_{\text{lvs}}) = \Omega\left(n^d\right).$$

□

[9] We represent each hypothesis by the node that identifies it.