Journal of Machine Learning Research 18 (2017) 1-32 Submitted 9/16; Revised 2/17; Published 4/17

Preference-based Teaching

Ziyuan Gao GAO257@CS.UREGINA.CA
Department of Computer Science, University of Regina

Christoph Ries CHRISTOPH.RIES@RUB.DE
Department of Mathematics, Ruhr-University Bochum

Hans U. Simon HANS.SIMON@RUB.DE
Department of Mathematics, Ruhr-University Bochum

Sandra Zilles ZILLES@CS.UREGINA.CA
Department of Computer Science, University of Regina

Editor: Manfred Warmuth

Abstract

We introduce a new model of teaching named "preference-based teaching" and a corresponding complexity parameter—the preference-based teaching dimension (PBTD)—representing the worst-case number of examples needed to teach any concept in a given concept class. Although the PBTD coincides with the well-known recursive teaching dimension (RTD) on finite classes, it is radically different on infinite ones: the RTD becomes infinite already for trivial infinite classes (such as half-intervals) whereas the PBTD evaluates to reasonably small values for a wide collection of infinite classes including classes consisting of so-called closed sets w.r.t. a given closure operator, including various classes related to linear sets over $\mathbb{N}_0$ (whose RTD had been studied quite recently) and including the class of Euclidean half-spaces. On top of presenting these concrete results, we provide the reader with a theoretical framework (of a combinatorial flavor) which helps to derive bounds on the PBTD.

Keywords: teaching dimension, preference relation, recursive teaching dimension, learning half-spaces, linear sets

1. Introduction

The classical model of teaching (Shinohara and Miyano, 1991; Goldman and Kearns, 1995) formulates the following interaction protocol between a teacher and a student:

- Both of them agree on a "classification-rule system", formally given by a concept class $\mathcal{L}$.
- In order to teach a specific concept $L \in \mathcal{L}$, the teacher presents to the student a *teaching set*, i.e., a set T of labeled examples so that L is the only concept in $\mathcal{L}$ that is consistent with T.
- The student determines L as the unique concept in $\mathcal{L}$ that is consistent with T.

Goldman and Mathias (1996) pointed out that this model of teaching is not powerful enough, since the teacher is required to make *any* consistent learner successful. A challenge is to model powerful teacher/student interactions without enabling unfair "coding tricks". Intuitively, the term "coding trick" refers to any form of undesirable collusion between teacher and learner, which would reduce the learning process to a mere decoding of a code the teacher sent to the learner. There is no generally accepted definition of what constitutes a coding trick, in part because teaching an exact learner could always be considered coding to some extent: the teacher presents a set of examples which the learner "decodes" into a concept.

In this paper, we adopt the notion of "valid teacher/learner pair" introduced by Goldman and Mathias (1996). They consider their model to be intuitively free of coding tricks while it provably allows for a much broader class of interaction protocols than the original teaching model. In particular, teaching may thus become more efficient in terms of the number of examples in the teaching sets. Further definitions of how to avoid unfair coding tricks have been suggested (Zilles et al., 2011), but they were less stringent than the one proposed by Goldman and Mathias. The latter simply requests that, if the learner hypothesizes concept L upon seeing a sample set S of labeled examples, then the learner will still hypothesize L when presented with any sample set $S \cup S'$, where S' contains only examples labeled consistently with L. A coding trick would then be any form of exchange between the teacher and the learner that does not satisfy this definition of validity.

The model of recursive teaching (Zilles et al., 2011; Mazadi et al., 2014), which is free of coding tricks according to the Goldman-Mathias definition, has recently gained attention because its complexity parameter, the recursive teaching dimension (RTD), has shown relations to the VC-dimension and to sample compression (Chen et al., 2016; Doliwa et al., 2014; Moran et al., 2015; Simon and Zilles, 2015), when focusing on finite concept classes. Below though we will give examples of rather simple infinite concept classes with infinite RTD, suggesting that the RTD is inadequate for addressing the complexity of teaching infinite classes.

In this paper, we introduce a model called *preference-based teaching*, in which the teacher and the student do not only agree on a classification-rule system $\mathcal{L}$ but also on a preference relation (a strict partial order) imposed on $\mathcal{L}$. If the labeled examples presented by the teacher allow for several consistent explanations (= consistent concepts) in $\mathcal{L}$, the student will choose a concept $L \in \mathcal{L}$ that she prefers most. This gives more flexibility to the teacher than the classical model: the set of labeled examples need not distinguish a target concept L from any other concept in $\mathcal{L}$ but only from those concepts L' over which L is not preferred.[1] At the same time, preference-based teaching yields valid teacher/learner pairs according to Goldman and Mathias's definition. We will show that the new model, despite avoiding coding tricks, is quite powerful. Moreover, as we will see in the course of the paper, it often allows for a very natural design of teaching sets.

Assume teacher and student choose a preference relation that minimizes the worst-case number M of examples required for teaching any concept in the class $\mathcal{L}$. This number M is then called the preference-based teaching dimension (PBTD) of $\mathcal{L}$. In particular, we will show the following:

(i) Recursive teaching is a special case of preference-based teaching where the preference relation satisfies a so-called "finite-depth condition". It is precisely this additional condition that renders recursive teaching useless for many natural and apparently simple infinite concept classes. Preference-based teaching successfully addresses these shortcomings of recursive teaching, see Section 3. For finite classes, PBTD and RTD are equal.

(ii) A wide collection of geometric and algebraic concept classes with infinite RTD can be taught very efficiently, i.e., with low PBTD. To establish such results, we show in Section 4 that spanning sets can be used as preference-based teaching sets with positive examples only — a result that is very simple to obtain but quite useful.

1. Such a preference relation can be thought of as a kind of bias in learning: the student is "biased" towards concepts that are preferred over others, and the teacher, knowing the student's bias, selects teaching sets accordingly.

(iii) In the preference-based model, linear sets over $\mathbb{N}_0$ with origin 0 and at most k generators can be taught with k positive examples, while recursive teaching with a bounded number of positive examples was previously shown to be impossible and it is unknown whether recursive teaching with a bounded number of positive and negative examples is possible for $k \geq 4$. We also give some almost matching upper and lower bounds on the PBTD for other classes of linear sets, see Section 6.

(iv) The PBTD of halfspaces in $\mathbb{R}^d$ is upper-bounded by 6, independent of the dimensionality d (see Section 7), while its RTD is infinite.

(v) We give full characterizations of concept classes that can be taught with only one example (or with only one example, which is positive) in the preference-based model (see Section 8).

Based on our results and the naturalness of the teaching sets and preference relations used in their proofs, we claim that preference-based teaching is far more suitable to the study of infinite concept classes than recursive teaching.

Parts of this paper were published in a previous conference version (Gao et al., 2016).

2. Basic Definitions and Facts

$\mathbb{N}_0$ denotes the set of all non-negative integers and $\mathbb{N}$ denotes the set of all positive integers. A *concept class* $\mathcal{L}$ is a family of subsets over a universe $\mathcal{X}$, i.e., $\mathcal{L} \subseteq 2^{\mathcal{X}}$ where $2^{\mathcal{X}}$ denotes the powerset of $\mathcal{X}$. The elements of $\mathcal{L}$ are called *concepts*. A *labeled example* is an element of $\mathcal{X} \times \{-,+\}$. We slightly deviate from this notation in Section 7, where our treatment of halfspaces makes it more convenient to use $\{-1,1\}$ instead of $\{-,+\}$, and in Section 8, where we perform Boolean operations on the labels and therefore use $\{0,1\}$ instead of $\{-,+\}$. Elements of $\mathcal{X}$ are called *examples*. Suppose that T is a set of labeled examples. Let $T^+ = \{x \in \mathcal{X} : (x,+) \in T\}$ and $T^- = \{x \in \mathcal{X} : (x,-) \in T\}$. A set $L \subseteq \mathcal{X}$ is *consistent with* T if it includes all examples in T that are labeled "$+$" and excludes all examples in T that are labeled "$-$", i.e, if $T^+ \subseteq L$ and $T^- \cap L = \emptyset$. A set of labeled examples that is consistent with L but not with L' is said to *distinguish* L *from* L'. The classical model of teaching is then defined as follows.

Definition 1 (Shinohara and Miyano (1991); Goldman and Kearns (1995)) *A* teaching set *for a concept $L \in \mathcal{L}$ w.r.t. $\mathcal{L}$ is a set T of labeled examples such that L is the only concept in $\mathcal{L}$ that is consistent with T, i.e., T distinguishes L from any other concept in $\mathcal{L}$. Define* $\mathrm{TD}(L,\mathcal{L}) = \inf\{|T| :$ *T is a teaching set for L w.r.t. $\mathcal{L}$*$\}$. *i.e.,* $\mathrm{TD}(L,\mathcal{L})$ *is the smallest possible size of a teaching set for L w.r.t. $\mathcal{L}$. If L has no finite teaching set w.r.t. $\mathcal{L}$, then* $\mathrm{TD}(L,\mathcal{L}) = \infty$. *The number* $\mathrm{TD}(\mathcal{L}) = \sup_{L \in \mathcal{L}} \mathrm{TD}(L,\mathcal{L}) \in \mathbb{N}_0 \cup \{\infty\}$ *is called the* teaching dimension of $\mathcal{L}$.

For technical reasons, we will occasionally deal with the number $\mathrm{TD}_{min}(\mathcal{L}) = \inf_{L \in \mathcal{L}} \mathrm{TD}(L, \mathcal{L})$, i.e., the number of examples needed to teach the concept from $\mathcal{L}$ that is easiest to teach.

In this paper, we will examine a teaching model in which the teacher and the student do not only agree on a classification-rule system $\mathcal{L}$ but also on a preference relation, denoted as $\prec$, imposed on $\mathcal{L}$. We assume that $\prec$ is a *strict partial order* on $\mathcal{L}$, i.e., $\prec$ is asymmetric and transitive. The partial order that makes every pair $L \neq L' \in \mathcal{L}$ incomparable is denoted by $\prec_\emptyset$. For every $L \in \mathcal{L}$, let

$$\mathcal{L}_{\prec L} = \{L' \in \mathcal{L} : L' \prec L\}$$

be the set of concepts over which L is strictly preferred. Note that $\mathcal{L}_{\prec_\emptyset L} = \emptyset$ for every $L \in \mathcal{L}$.

As already noted above, a teaching set T of L w.r.t. $\mathcal{L}$ distinguishes L from any other concept in $\mathcal{L}$. If a preference relation comes into play, then T will be exempted from the obligation to distinguish L from the concepts in $\mathcal{L}_{\prec L}$ because L is strictly preferred over them anyway.

Definition 2 *A* teaching set for $L \subseteq X$ w.r.t. $(\mathcal{L}, \prec)$ *is defined as a teaching set for L w.r.t. $\mathcal{L} \setminus \mathcal{L}_{\prec L}$. Furthermore define*

$$\mathrm{PBTD}(L, \mathcal{L}, \prec) = \inf\{|T| : T \textit{ is a teaching set for } L \textit{ w.r.t. } (\mathcal{L}, \prec)\} \in \mathbb{N}_0 \cup \{\infty\} \ .$$

The number $\mathrm{PBTD}(\mathcal{L}, \prec) = \sup_{L \in \mathcal{L}} \mathrm{PBTD}(L, \mathcal{L}, \prec) \in \mathbb{N}_0 \cup \{\infty\}$ *is called the* teaching dimension of $(\mathcal{L}, \prec)$.

Definition 2 implies that

$$\mathrm{PBTD}(L, \mathcal{L}, \prec) = \mathrm{TD}(L, \mathcal{L} \setminus \mathcal{L}_{\prec L}) \ . \tag{1}$$

Let $L \mapsto T(L)$ be a mapping that assigns a teaching set for L w.r.t. $(\mathcal{L}, \prec)$ to every $L \in \mathcal{L}$. It is obvious from Definition 2 that T must be injective, i.e., $T(L) \neq T(L')$ if L and L' are distinct concepts from $\mathcal{L}$. The classical model of teaching is obtained from the model described in Definition 2 when we plug in the empty preference relation $\prec_\emptyset$ for $\prec$. In particular, $\mathrm{PBTD}(\mathcal{L}, \prec_\emptyset) = \mathrm{TD}(\mathcal{L})$.

We are interested in finding the partial order that is optimal for the purpose of teaching and we aim at determining the corresponding teaching dimension. This motivates the following notion:

Definition 3 *The* preference-based teaching dimension of $\mathcal{L}$ *is given by*

$$\mathrm{PBTD}(\mathcal{L}) = \inf\{\mathrm{PBTD}(\mathcal{L}, \prec) : \prec \textit{ is a strict partial order on } \mathcal{L}\} \ .$$

A relation R' on $\mathcal{L}$ is said to be an *extension of a relation* R if $R \subseteq R'$. The *order-extension principle* states that any partial order has a linear extension (Jech, 1973). The following result (whose second assertion follows from the first one in combination with the order-extension principle) is pretty obvious:

Lemma 4

1. *Suppose that $\prec'$ extends $\prec$. If T is a teaching set for L w.r.t. $(\mathcal{L}, \prec)$, then T is a teaching set for L w.r.t. $(\mathcal{L}, \prec')$. Moreover* $\mathrm{PBTD}(\mathcal{L}, \prec') \le \mathrm{PBTD}(\mathcal{L}, \prec)$.
2. $\mathrm{PBTD}(\mathcal{L}) = \inf\{\mathrm{PBTD}(\mathcal{L}, \prec) : \prec$ *is a strict linear order on* $\mathcal{L}\}$.

Recall that Goldman and Mathias (1996) suggested to avoid coding tricks by requesting that any superset S of a teaching set for a concept L remains a teaching set, if S is consistent with L. This property is obviously satisfied in preference-based teaching. A preference-based teaching set needs to distinguish a concept L from all concepts in $\mathcal{L}$ that are preferred over L. Adding more labeled examples from L to such a teaching set will still result in a set distinguishing L from all concepts in $\mathcal{L}$ that are preferred over L.

Preference-based teaching with positive examples only. Suppose that $\mathcal{L}$ contains two concepts L, L' such that $L \subset L'$. In the classical teaching model, any teaching set for L w.r.t. $\mathcal{L}$ has to employ a negative example in order to distinguish L from L'. Symmetrically, any teaching set for L' w.r.t. $\mathcal{L}$ has to employ a positive example. Thus classical teaching cannot be performed with one type of examples only unless $\mathcal{L}$ is an antichain w.r.t. inclusion. As for preference-based teaching, the restriction to one type of examples is much less severe, as our results below will show.

A teaching set T for $L \in \mathcal{L}$ w.r.t. $(\mathcal{L}, \prec)$ is said to be *positive* if it does not make use of negatively labeled examples, i.e., if $T^- = \emptyset$. In the sequel, we will occasionally identify a positive teaching set T with T^+. A positive teaching set for L w.r.t. $(\mathcal{L}, \prec)$ can clearly not distinguish L from a proper superset of L in $\mathcal{L}$. Thus, the following holds:

Lemma 5 *Suppose that $L \mapsto T^+(L)$ maps each $L \in \mathcal{L}$ to a positive teaching set for L w.r.t. $(\mathcal{L}, \prec)$. Then $\prec$ must be an extension of $\supset$ (so that proper subsets of a set L are strictly preferred over L) and, for every $L \in \mathcal{L}$, the set $T^+(L)$ must distinguish L from every proper subset of L in $\mathcal{L}$.*

Define

$$\mathrm{PBTD}^+(L, \mathcal{L}, \prec) = \inf\{|T| : T \text{ is a positive teaching set for } L \text{ w.r.t. } (\mathcal{L}, \prec)\} \ . \tag{2}$$

The number $\mathrm{PBTD}^+(\mathcal{L}, \prec) = \sup_{L \in \mathcal{L}} \mathrm{PBTD}^+(L, \mathcal{L}, \prec)$ (possibly ∞) is called the *positive teaching dimension of* $(\mathcal{L}, \prec)$. The *positive preference-based teaching dimension of* $\mathcal{L}$ is then given by

$$\mathrm{PBTD}^+(\mathcal{L}) = \inf\{\mathrm{PBTD}^+(\mathcal{L}, \prec) : \prec \text{ is a strict partial order on } \mathcal{L}\} \ . \tag{3}$$

Monotonicity. A complexity measure K that assigns a number $K(\mathcal{L}) \in \mathbb{N}_0$ to a concept class $\mathcal{L}$ is said to be *monotonic* if $\mathcal{L}' \subseteq \mathcal{L}$ implies that $K(\mathcal{L}') \leq K(\mathcal{L})$. It is well known (and trivial to see) that TD is monotonic. It is fairly obvious that PBTD is monotonic, too:

Lemma 6 PBTD *and* PBTD^+ *are monotonic.*

As an application of monotonicity, we show the following result:

Lemma 7 *For every finite subclass $\mathcal{L}'$ of $\mathcal{L}$, we have* $\mathrm{PBTD}(\mathcal{L}) \geq \mathrm{PBTD}(\mathcal{L}') \geq \mathrm{TD}_{min}(\mathcal{L}')$.

Proof The first inequality holds because PBTD is monotonic. The second inequality follows from the fact that a finite partially ordered set must contain a minimal element. Thus, for any fixed choice of $\prec$, $\mathcal{L}'$ must contain a concept L' such that $\mathcal{L}'_{\prec L'} = \emptyset$. Hence,

$$\mathrm{PBTD}(\mathcal{L}', \prec) \geq \mathrm{PBTD}(L', \mathcal{L}', \prec) \stackrel{(1)}{=} \mathrm{TD}(L', \mathcal{L}' \setminus \mathcal{L}'_{\prec L'}) = \mathrm{TD}(L', \mathcal{L}') \geq \mathrm{TD}_{min}(\mathcal{L}') \ .$$

Since this holds for any choice of $\prec$, we get $\mathrm{PBTD}(\mathcal{L}') \geq \mathrm{TD}_{min}(\mathcal{L}')$, as desired. ∎

3. Preference-based versus Recursive Teaching

The preference-based teaching dimension is a relative of the recursive teaching dimension. In fact, both notions coincide on finite classes, as we will see shortly. We first recall the definitions of the recursive teaching dimension and of some related notions (Zilles et al., 2011; Mazadi et al., 2014).

A *teaching sequence for* $\mathcal{L}$ is a sequence of the form $\mathcal{S} = (\mathcal{L}_i, d_i)_{i \geq 1}$ where $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3, \ldots$ form a partition of $\mathcal{L}$ into non-empty sub-classes and, for every $i \geq 1$, we have that

$$d_i = \sup_{L \in \mathcal{L}_i} \mathrm{TD}\left(L, \mathcal{L} \setminus \cup_{j=1}^{i-1} \mathcal{L}_j\right) \ . \tag{4}$$

If, for every $i \geq 1$, d_i is the supremum over all $L \in \mathcal{L}_i$ of the smallest size of a *positive teaching set* for L w.r.t. $\cup_{j \geq i} \mathcal{L}_j$ (and $d_i = \infty$ if some $L \in \mathcal{L}_i$ does not have a positive teaching set w.r.t. $\cup_{j \geq i} \mathcal{L}_j$), then $\mathcal{S}$ is said to be a *positive teaching sequence for* $\mathcal{L}$. The *order* of a teaching sequence or a positive teaching sequence $\mathcal{S}$ (possibly ∞) is defined as $\mathrm{ord}(\mathcal{S}) = \sup_{i \geq 1} d_i$. The *recursive teaching dimension of* $\mathcal{L}$ (possibly ∞) is defined as the order of the teaching sequence of lowest order for $\mathcal{L}$. More formally, $\mathrm{RTD}(\mathcal{L}) = \inf_{\mathcal{S}} \mathrm{ord}(\mathcal{S})$ where $\mathcal{S}$ ranges over all teaching sequences for $\mathcal{L}$. Similarly, $\mathrm{RTD}^+(\mathcal{L}) = \inf_{\mathcal{S}} \mathrm{ord}(\mathcal{S})$, where $\mathcal{S}$ ranges over all positive teaching sequences for $\mathcal{L}$. Note that the following holds for every $\mathcal{L}' \subseteq \mathcal{L}$ and for every teaching sequence $\mathcal{S} = (\mathcal{L}_i, d_i)_{i \geq 1}$ for $\mathcal{L}'$ such that $\mathrm{ord}(\mathcal{S}) = \mathrm{RTD}(\mathcal{L}')$:

$$\mathrm{RTD}(\mathcal{L}) \geq \mathrm{RTD}(\mathcal{L}') = \mathrm{ord}(\mathcal{S}) \geq d_1 = \sup_{L \in \mathcal{L}_1} \mathrm{TD}(L, \mathcal{L}') \geq \mathrm{TD}_{min}(\mathcal{L}') \ . \tag{5}$$

Note an important difference between PBTD and RTD: while $\mathrm{RTD}(\mathcal{L}) \geq \mathrm{TD}_{min}(\mathcal{L}')$ for *all* $\mathcal{L}' \subseteq \mathcal{L}$, in general the same holds for PBTD only when restricted to finite $\mathcal{L}'$, cf. Lemma 7. This difference will become evident in the proof of Lemma 10.

The *depth* of $L \in \mathcal{L}$ w.r.t. a strict partial order imposed on $\mathcal{L}$ is defined as the length of the longest chain in $(\mathcal{L}, \prec)$ that ends with the $\prec$-maximal element L (resp. as ∞ if there is no bound on the length of these chains). The recursive teaching dimension is related to the preference-based teaching dimension as follows:

Lemma 8 $\mathrm{RTD}(\mathcal{L}) = \inf_{\prec} \mathrm{PBTD}(\mathcal{L}, \prec)$ *and* $\mathrm{RTD}^+(\mathcal{L}) = \inf_{\prec} \mathrm{PBTD}^+(\mathcal{L}, \prec)$ *where* $\prec$ *ranges over all strict partial orders on* $\mathcal{L}$ *that satisfy the following "finite-depth condition": every* $L \in \mathcal{L}$ *has a finite depth w.r.t.* $\prec$.

The following is an immediate consequence of Lemma 8 and the trivial observation that the finite-depth condition is always satisfied if $\mathcal{L}$ is finite:

Corollary 9 $\mathrm{PBTD}(\mathcal{L}) \leq \mathrm{RTD}(\mathcal{L})$*, with equality if* $\mathcal{L}$ *is finite.*

While $\mathrm{PBTD}(\mathcal{L})$ and $\mathrm{RTD}(\mathcal{L})$ refer to the same finite number when $\mathcal{L}$ is finite, there are classes for which the RTD is infinity and yet the PBTD is finite, as Lemma 10 will show.

Lemma 10 *There exists an infinite class* $\mathcal{L}_\infty$ *of VC-dimension* 1 *such that* $\mathrm{PBTD}^+(\mathcal{L}_\infty) = 1$ *and* $\mathrm{RTD}(\mathcal{L}_\infty) = \infty$.

Proof Let $\mathcal{L}_\infty$ be the family of closed half-intervals over $[0,1)$, i.e., $\mathcal{L}_\infty = \{[0,a] : 0 \le a < 1\}$. We first prove that $\mathrm{PBTD}^+(\mathcal{L}_\infty) = 1$. Consider the preference relation given by $[0,b] \prec [0,a]$ iff $a < b$. Then, for each $0 \le a < 1$, we have

$$\mathrm{PBTD}([0,a], \mathcal{L}_\infty, \prec) \stackrel{(1)}{=} \mathrm{TD}([0,a], \{[0,b] :\ 0 \le b \le a\}) = 1$$

because the single example $(a, +)$ suffices for distinguishing $[0,a]$ from any interval $[0,b]$ with $b < a$.

It was observed by Moran et al. (2015) already that $\mathrm{RTD}(\mathcal{L}_\infty) = \infty$ because every teaching set for some $[0,a]$ must contain an infinite sequence of distinct reals that converges from above to a. Thus, using Equation (5) with $\mathcal{L}' = \mathcal{L}$, we have $\mathrm{RTD}(\mathcal{L}_\infty) \ge \mathrm{TD}_{min}(\mathcal{L}_\infty) = \infty$. ■

Remark 11 *For every $k \ge 1$, there exists an infinite class $\mathcal{L}_k$ such that $\mathrm{PBTD}^+(\mathcal{L}_k) = 1$ and $\mathrm{RTD}(\mathcal{L}_k) = k$; see (Gao et al., 2017, Lemma 6).*

4. Preference-based Teaching with Positive Examples Only

The main purpose of this section is to relate positive preference-based teaching to "spanning sets" and "closure operators", which are well-studied concepts in the computational learning theory literature. Let $\mathcal{L}$ be a concept class over the universe $\mathcal{X}$. We say that $S \subseteq \mathcal{X}$ is a *spanning set* of $L \in \mathcal{L}$ w.r.t. $\mathcal{L}$ if $S \subseteq L$ and any set in $\mathcal{L}$ that contains S must contain L as well.[2] In other words, L is the unique smallest concept in $\mathcal{L}$ that contains S. We say that $S \subseteq \mathcal{X}$ is a *weak spanning set* of $L \in \mathcal{L}$ w.r.t. $\mathcal{L}$ if $S \subseteq L$ and S is not contained in any proper subset of L in $\mathcal{L}$.[3] We denote by $I(\mathcal{L})$ (resp. $I'(\mathcal{L})$) the smallest number k such that every concept $L \in \mathcal{L}$ has a spanning set (resp. a weak spanning set) w.r.t. $\mathcal{L}$ of size at most k. Note that S is a spanning set of L w.r.t. $\mathcal{L}$ iff S distinguishes L from all concepts in $\mathcal{L}$ except for supersets of L, i.e., iff S is a positive teaching set for L w.r.t. $(\mathcal{L}, \supset)$. Similarly, S is a weak spanning set of L w.r.t. $\mathcal{L}$ iff S distinguishes L from all its proper subsets in $\mathcal{L}$ (which is necessarily the case when S is a positive teaching set). These observations can be summarized as follows:

$$I'(\mathcal{L}) \le \mathrm{PBTD}^+(\mathcal{L}) \le \mathrm{PBTD}^+(\mathcal{L}, \supset) \le I(\mathcal{L}) \ . \tag{6}$$

The last two inequalities are straightforward. The inequality $I'(\mathcal{L}) \le \mathrm{PBTD}^+(\mathcal{L})$ follows from Lemma 5, which implies that no concept L can have a preference-based teaching set T smaller than its smallest weak spanning set. Such a set T would be consistent with some proper subset of L, which is impossible by Lemma 5.

Suppose $\mathcal{L}$ is intersection-closed. Then $\cap_{L \in \mathcal{L}: S \subseteq L} L$ is the unique smallest concept in $\mathcal{L}$ containing S. If $S \subseteq L_0$ is a weak spanning set of $L_0 \in \mathcal{L}$, then $\cap_{L \in \mathcal{L}: S \subseteq L} L = L_0$ because, on the one hand, $\cap_{L \in \mathcal{L}: S \subseteq L} L \subseteq L_0$ and, on the other hand, no proper subset of L_0 in $\mathcal{L}$ contains S. Thus the distinction between spanning sets and weak spanning sets is blurred for intersection-closed classes:

2. This generalizes the classical definition of a spanning set (Helmbold et al., 1990), which is given w.r.t. intersection-closed classes only.
3. Weak spanning sets have been used in the field of recursion-theoretic inductive inference under the name "tell-tale sets" (Angluin, 1980).

Lemma 12 *Suppose that $\mathcal{L}$ is intersection-closed. Then* $I'(\mathcal{L}) = \mathrm{PBTD}^+(\mathcal{L}) = I(\mathcal{L})$.

Example 1 *Let $\mathcal{R}_d$ denote the class of d-dimensional axis-parallel hyper-rectangles (= d-dimensional boxes). This class is intersection-closed and clearly $I(\mathcal{R}_d) = 2$. Thus* $\mathrm{PBTD}^+(\mathcal{R}_d) = 2$.

A mapping $\mathrm{cl} : 2^{\mathcal{X}} \to 2^{\mathcal{X}}$ is said to be a *closure operator* on the universe $\mathcal{X}$ if the following conditions hold for all sets $A, B \subseteq \mathcal{X}$:

$$A \subseteq B \Rightarrow \mathrm{cl}(A) \subseteq \mathrm{cl}(B) \text{ and } A \subseteq \mathrm{cl}(A) = \mathrm{cl}(\mathrm{cl}(A)) \ .$$

The following notions refer to an arbitrary but fixed closure operator. The set $\mathrm{cl}(A)$ is called the *closure* of A. A set C is said to be *closed* if $\mathrm{cl}(C) = C$. It follows that precisely the sets $\mathrm{cl}(A)$ with $A \subseteq \mathcal{X}$ are closed. With this notation, we observe the following lemma.

Lemma 13 *Let $\mathcal{C}$ be the set of all closed subsets of $\mathcal{X}$ under some closure operator* cl, *and let $L \in \mathcal{C}$. If $L = \mathrm{cl}(S)$, then S is a spanning set of L w.r.t. $\mathcal{C}$.*

Proof Suppose $L' \in \mathcal{C}$ and $S \subseteq L'$. Then $L = \mathrm{cl}(S) \subseteq \mathrm{cl}(L') = L'$. ■

For every closed set $L \in \mathcal{L}$, let $s_{cl}(L)$ denote the size (possibly ∞) of the smallest set $S \subseteq \mathcal{X}$ such that $\mathrm{cl}(S) = L$. With this notation, we get the following (trivial but useful) result:

Theorem 14 *Given a closure operator, let $\mathcal{C}[m]$ be the class of all closed subsets $C \subseteq \mathcal{X}$ with $s_{cl}(C) \le m$. Then* $\mathrm{PBTD}^+(\mathcal{C}[m]) \le \mathrm{PBTD}^+(\mathcal{C}[m], \supset) \le m$. *Moreover, this holds with equality provided that $\mathcal{C}[m] \setminus \mathcal{C}[m-1] \neq \emptyset$.*

Proof The inequality $\mathrm{PBTD}^+(\mathcal{C}[m], \supset) \le m$ follows directly from Equation (6) and Lemma 13. Pick a concept $C_0 \in \mathcal{C}[m]$ such that $s_{cl}(C_0) = m$. Then any subset S of C_0 of size less than m spans only a proper subset of C_0, i.e., $\mathrm{cl}(S) \subset C_0$. Thus S does not distinguish C_0 from $\mathrm{cl}(S)$. However, by Lemma 5, any preference-based learner must strictly prefer $\mathrm{cl}(S)$ over C_0. It follows that there is no positive teaching set of size less than m for C_0 w.r.t. $\mathcal{C}[m]$. ■

Many natural classes can be cast as classes of the form $\mathcal{C}[m]$ by choosing the universe and the closure operator appropriately; the following examples illustrate the usefulness of Theorem 14 in that regard.

Example 2 *Let*

$$\mathrm{LINSET}_k = \{\langle G \rangle : (G \subset \mathbb{N}) \wedge (1 \le |G| \le k)\}$$

where $\langle G \rangle = \left\{ \sum_{g \in G} a(g) g : a(g) \in \mathbb{N}_0 \right\}$. *In other words,* LINSET_k *is the set of all non-empty linear subsets of $\mathbb{N}_0$ that are generated by at most k generators. Note that the mapping $G \mapsto \langle G \rangle$ is a closure operator over the universe $\mathbb{N}_0$. Since obviously* $\mathrm{LINSET}_k \setminus \mathrm{LINSET}_{k-1} \neq \emptyset$, *we obtain* $\mathrm{PBTD}^+(\mathrm{LINSET}_k) = k$.

Example 3 *Let $\mathcal{X} = \mathbb{R}^2$ and let $\mathcal{C}_k$ be the class of convex polygons with at most k vertices. Defining* $\mathrm{cl}(S)$ *to be the convex closure of S, we obtain $\mathcal{C}[k] = \mathcal{C}_k$ and thus* $\mathrm{PBTD}^+(\mathcal{C}_k) = k$.

Example 4 *Let $\mathcal{X} = \mathbb{R}^n$ and let $\mathcal{C}_k$ be the class of polyhedral cones that can be generated by k (or less) vectors in $\mathbb{R}^n$. If we take* $\mathrm{cl}(S)$ *to be the conic closure of $S \subseteq \mathbb{R}^n$, then $\mathcal{C}[k] = \mathcal{C}_k$ and thus* $\mathrm{PBTD}^+(\mathcal{C}_k) = k$.

5. A Convenient Technique for Proving Upper Bounds

In this section, we give an alternative definition of the preference-based teaching dimension using the notion of an "admissible mapping". Given a concept class $\mathcal{L}$ over a universe $\mathcal{X}$, let T be a mapping $L \mapsto T(L) \subseteq \mathcal{X} \times \{-, +\}$ that assigns a set $T(L)$ of labeled examples to every set $L \in \mathcal{L}$ such that the labels in $T(L)$ are consistent with L. The *order* of T, denoted as $\mathrm{ord}(T)$, is defined as $\sup_{L\in\mathcal{L}} |T(L)| \in \mathbb{N} \cup \{\infty\}$. Define the mappings T^+ and T^- by setting $T^+(L) = \{x : (x, +) \in T(L)\}$ and $T^-(L) = \{x : (x, -) \in T(L)\}$ for every $L \in \mathcal{L}$. We say that T is *positive* if $T^-(L) = \emptyset$ for every $L \in \mathcal{L}$. In the sequel, we will occasionally identify a positive mapping $L \mapsto T(L)$ with the mapping $L \mapsto T^+(L)$. The symbol "$+$" as an upper index of T will always indicate that the underlying mapping T is positive.
The following relation will help to clarify under which conditions the sets $(T(L))_{L\in\mathcal{L}}$ are teaching sets w.r.t. a suitably chosen preference relation:

$$R_T = \{(L, L') \in \mathcal{L} \times \mathcal{L} : \ (L \neq L') \wedge (L \text{ is consistent with } T(L'))\} \ .$$

The transitive closure of R_T is denoted as $\mathrm{trcl}(R_T)$ in the sequel. The following notion will play an important role in this paper:

Definition 15 *A mapping $L \mapsto T(L)$ with L ranging over all concepts in $\mathcal{L}$ is said to be* admissible for $\mathcal{L}$ *if the following holds:*

1. *For every $L \in \mathcal{L}$, L is consistent with $T(L)$.*
2. *The relation $\mathrm{trcl}(R_T)$ is asymmetric (which clearly implies that R_T is asymmetric too).*

If T is admissible, then $\mathrm{trcl}(R_T)$ is transitive and asymmetric, i.e., $\mathrm{trcl}(R_T)$ is a strict partial order on $\mathcal{L}$. We will therefore use the notation $\prec_T$ instead of $\mathrm{trcl}(R_T)$ whenever T is known to be admissible.

Lemma 16 *Suppose that T^+ is a positive admissible mapping for $\mathcal{L}$. Then the relation $\prec_{T^+}$ on $\mathcal{L}$ extends the relation $\supset$ on $\mathcal{L}$. More precisely, the following holds for all $L, L' \in \mathcal{L}$:*

$$L' \subset L \Rightarrow (L, L') \in R_{T^+} \Rightarrow L \prec_{T^+} L' \ .$$

Proof If T^+ is admissible, then L' is consistent with $T^+(L')$. Thus $T^+(L') \subseteq L' \subset L$ so that L is consistent with $T^+(L')$ too. Therefore $(L, L') \in R_{T^+}$, i.e., $L \prec_{T^+} L'$. ■

The following result clarifies how admissible mappings are related to preference-based teaching:

Lemma 17 *For each concept class $\mathcal{L}$, the following holds:*

$$\mathrm{PBTD}(\mathcal{L}) = \inf_T \mathrm{ord}(T) \ \textit{and} \ \mathrm{PBTD}^+(\mathcal{L}) = \inf_{T^+} \mathrm{ord}(T^+)$$

where T ranges over all mappings that are admissible for $\mathcal{L}$ and T^+ ranges over all positive mappings that are admissible for $\mathcal{L}$.

Proof We restrict ourselves to the proof for $\mathrm{PBTD}(\mathcal{L}) = \inf_T \mathrm{ord}(T)$ because the equation $\mathrm{PBTD}^+(\mathcal{L}) = \inf_{T^+} \mathrm{ord}(T^+)$ can be obtained in a similar fashion. We first prove that $\mathrm{PBTD}(\mathcal{L}) \le \inf_T \mathrm{ord}(T)$. Let T be an admissible mapping for $\mathcal{L}$. It suffices to show that, for every $L \in \mathcal{L}$, $T(L)$ is a teaching set for L w.r.t. $(\mathcal{L}, \prec_T)$. Suppose $L' \in \mathcal{L} \setminus \{L\}$ is consistent with $T(L)$. Then $(L', L) \in R_T$ and thus $L' \prec_T L$. It follows that $\prec_T$ prefers L over all concepts $L' \in \mathcal{L} \setminus \{L\}$ that are consistent with $T(L)$. Thus T is a teaching set for L w.r.t. $(\mathcal{L}, \prec_T)$, as desired.

We now prove that $\inf_T \mathrm{ord}(T) \le \mathrm{PBTD}(\mathcal{L})$. Let $\prec$ be a strict partial order on $\mathcal{L}$ and let T be a mapping such that, for every $L \in \mathcal{L}$, $T(L)$ is a teaching set for L w.r.t. $(\mathcal{L}, \prec)$. It suffices to show that T is admissible for $\mathcal{L}$. Consider a pair $(L', L) \in R_T$. The definition of R_T implies that $L' \ne L$ and that L' is consistent with $T(L)$. Since $T(L)$ is a teaching set w.r.t. $(\mathcal{L}, \prec)$, it follows that $L' \prec L$. Thus, $\prec$ is an extension of R_T. Since $\prec$ is transitive, it is even an extension of $\mathrm{trcl}(R_T)$. Because $\prec$ is asymmetric, $\mathrm{trcl}(R_T)$ must be asymmetric, too. It follows that T is admissible. ■

6. Preference-based Teaching of Linear Sets

Some work in computational learning theory (Abe, 1989; Gao et al., 2015; Takada, 1992) is concerned with learning *semi-linear sets*, i.e., unions of linear subsets of $\mathbb{N}^k$ for some fixed $k \ge 1$, where each linear set consists of exactly those elements that can be written as the sum of some constant vector c and a linear combination of the elements of some fixed set of generators, see Example 2. While semi-linear sets are of common interest in mathematics in general, they play a particularly important role in the theory of formal languages, due to *Parikh's theorem*, by which the so-called Parikh vectors of strings in a context-free language always form a semi-linear set (Parikh, 1966).

A recent study (Gao et al., 2015) analyzed computational teaching of classes of linear subsets of $\mathbb{N}$ (where $k = 1$) and some variants thereof, as a substantially simpler yet still interesting special case of semi-linear sets. In this section, we extend that study to preference-based teaching.

Within the scope of this section, all concept classes are formulated over the universe $\mathcal{X} = \mathbb{N}_0$. Let $G = \{g_1, \ldots, g_k\}$ be a finite subset of $\mathbb{N}$. We denote by $\langle G \rangle$ resp. by $\langle G \rangle_+$ the following sets:

$$\langle G \rangle = \left\{ \sum_{i=1}^{k} a_i g_i :\ a_1, \ldots, a_k \in \mathbb{N}_0 \right\} \text{ and } \langle G \rangle_+ = \left\{ \sum_{i=1}^{k} a_i g_i :\ a_1, \ldots, a_k \in \mathbb{N} \right\} .$$

We will determine (at least approximately) the preference-based teaching dimension of the following concept classes over $\mathbb{N}_0$:

$$\begin{aligned}
\mathrm{LINSET}_k &= \{\langle G \rangle :\ (G \subset \mathbb{N}) \wedge (1 \le |G| \le k)\} \ . \\
\mathrm{CF\text{-}LINSET}_k &= \{\langle G \rangle :\ (G \subset \mathbb{N}) \wedge (1 \le |G| \le k) \wedge (\gcd(G) = 1)\} \ . \\
\mathrm{NE\text{-}LINSET}_k &= \{\langle G \rangle_+ :\ (G \subset \mathbb{N}) \wedge (1 \le |G| \le k)\} \ . \\
\mathrm{NE\text{-}CF\text{-}LINSET}_k &= \{\langle G \rangle_+ :\ (G \subset \mathbb{N}) \wedge (1 \le |G| \le k) \wedge (gcd(G) = 1)\} \ .
\end{aligned}$$

A subset of $\mathbb{N}_0$ whose complement in $\mathbb{N}_0$ is finite is said to be *co-finite*. The letters "CF" in CF-LINSET mean "co-finite". The concepts in LINSET_k have the algebraic structure of a monoid w.r.t. addition. The concepts in $\mathrm{CF\text{-}LINSET}_k$ are also known as "numerical semigroups" (Rosales and García-Sánchez, 2009). A zero coefficient $a_j = 0$ erases g_j in the linear combination $\sum_{i=1}^{k} a_i g_i$. Coefficients from $\mathbb{N}$ are non-erasing in this sense. The letters "NE" in "NE-LINSET" mean "non-erasing".

The *shift-extension* $\mathcal{L}'$ of a concept class $\mathcal{L}$ over the universe $\mathbb{N}_0$ is defined as follows:

$$\mathcal{L}' = \{c + L :\ (c \in \mathbb{N}_0) \wedge (L \in \mathcal{L})\}\ . \tag{7}$$

The following bounds on RTD and RTD$^+$ (for sufficiently large values of k)[4] are known from (Gao et al., 2015):

	RTD$^+$	RTD
LINSET_k	$= \infty$	?
$\mathrm{CF\text{-}LINSET}_k$	$= k$	$\in \{k-1, k\}$
$\mathrm{NE\text{-}LINSET}'_k$	$= k+1$	$\in \{k-1, k, k+1\}$

Here $\mathrm{NE\text{-}LINSET}'_k$ denotes the shift-extension of $\mathrm{NE\text{-}LINSET}_k$.

The following result shows the corresponding bounds with PBTD in place of RTD:

Theorem 18 *The bounds in the following table are valid:*

	PBTD$^+$	PBTD
LINSET_k	$= k$	$\in \{k-1, k\}$
$\mathrm{CF\text{-}LINSET}_k$	$= k$	$\in \{k-1, k\}$
$\mathrm{NE\text{-}LINSET}_k$	$\in \left[\left\lfloor \frac{k-1}{2} \right\rfloor : k\right]$	$\in \left[\left\lfloor \frac{k-1}{2} \right\rfloor : k\right]$
$\mathrm{NE\text{-}CF\text{-}LINSET}_k$	$\in \left[\left\lfloor \frac{k-1}{2} \right\rfloor : k\right]$	$\in \left[\left\lfloor \frac{k-1}{2} \right\rfloor : k\right]$

Moreover

$$\mathrm{PBTD}^+(\mathcal{L}') = k+1\ \wedge\ \mathrm{PBTD}(\mathcal{L}') \in \{k-1, k, k+1\} \tag{8}$$

holds for all $\mathcal{L} \in \{\mathrm{LINSET}_k, \mathrm{CF\text{-}LINSET}_k, \mathrm{NE\text{-}LINSET}_k, \mathrm{NE\text{-}CF\text{-}LINSET}_k\}$.

Note that the equation $\mathrm{PBTD}^+(\mathrm{LINSET}_k) = k$ was already proven in Example 2, using the fact that $G \mapsto \langle G \rangle$ is a closure operator. Since $G \mapsto \langle G \rangle_+$ is not a closure operator, we give a separate argument to prove an upper bound of k on $\mathrm{PBTD}^+(\mathrm{NE\text{-}LINSET}_k)$ (see Lemma 37 in Appendix A). All other upper bounds in Theorem 18 are then easy to derive. The lower bounds in Theorem 18 are much harder to obtain. A complete proof of Theorem 18 will be given in Appendix A.

Remark 19 *The lower bound on* $\mathrm{PBTD}^+(\mathrm{NE\text{-}LINSET}_k)$ *and* $\mathrm{PBTD}^+(\mathrm{NE\text{-}CF\text{-}LINSET}_k)$ *may be improved to* $k-1$*; see (Gao et al., 2017, Theorem 2, Appendix A.3).*

4. For instance, $\mathrm{RTD}^+(\mathrm{LINSET}_k) = \infty$ holds for all $k \geq 2$ and $\mathrm{RTD}(\mathrm{LINSET}_k) = ?$ (where "?" means "unknown") holds for all $k \geq 4$.

7. Preference-based Teaching of Halfspaces

In this section, we study preference-based teaching of halfspaces. We will denote the all-zeros vector as $\vec{0}$. The vector with 1 in coordinate i and with 0 in the remaining coordinates is denoted as $\vec{e}_i$. The dimension of the Euclidean space in which these vectors reside will always be clear from the context. The sign of a real number x (with value 1 if $x > 0$, value -1 if $x < 0$, and value 0 if $x = 0$) is denoted by $\mathrm{sign}(x)$.

Suppose that $w \in \mathbb{R}^d \setminus \{\vec{0}\}$ and $b \in \mathbb{R}$. The *(positive) halfspace induced by w and b* is then given by

$$H_{w,b} = \{x \in \mathbb{R}^d :\ w^\top x + b \geq 0\} \ .$$

Instead of $H_{w,0}$, we simply write H_w. Let $\mathcal{H}_d$ denote the class of d-dimensional Euclidean halfspaces:

$$\mathcal{H}_d = \{H_{w,b} :\ w \in \mathbb{R}^d \setminus \{\vec{0}\} \wedge b \in \mathbb{R}\} \ .$$

Similarly, $\mathcal{H}_d^0$ denotes the class of d-dimensional homogeneous Euclidean halfspaces:

$$\mathcal{H}_d^0 = \{H_w :\ w \in \mathbb{R}^d \setminus \{\vec{0}\}\} \ .$$

Let S_{d-1} denote the $(d-1)$-dimensional unit sphere in $\mathbb{R}^d$. Moreover $S_{d-1}^+ = \{x \in S_{d-1} : x_d > 0\}$ denotes the "northern hemisphere". If not stated explicitly otherwise, we will represent homogeneous halfspaces with normalized vectors residing on the unit sphere. We remind the reader of the following well-known fact:

Remark 20 *The orthogonal group in dimension d (i.e., the multiplicative group of orthogonal $(d \times d)$-matrices) acts transitively on S_{d-1} and it conserves the inner product.*

We now prove a helpful lemma, stating that each vector w^* in the northern hemisphere may serve as a representative for some homogeneous halfspace H_u in the sense that all other elements of H_u in the northern hemisphere have a strictly smaller d-th component than w^*. This will later help to teach homogeneous halfspaces with a preference that orders vectors by the size of their last coordinate.

Lemma 21 *Let $d \geq 2$, let $0 < h \leq 1$ and let $R_{d,h} = \{w \in S_{d-1} : w_d = h\}$. With this notation the following holds. For every $w^* \in R_{d,h}$, there exists $u \in \mathbb{R}^d \setminus \{\vec{0}\}$ such that*

$$(w^* \in H_u) \wedge (\forall w \in (S_{d-1}^+ \cap H_u) \setminus \{w^*\} : w_d < h) \ . \tag{9}$$

Proof For $h = 1$, the statement is trivial, since $R_{d,1} = \{\vec{e}_d\}$. So let $h < 1$.

Because of Remark 20, we may assume without loss of generality that the vector $w^* \in R_{d,h}$ equals $(0, \ldots, 0, \sqrt{1-h^2}, h)$. It suffices therefore to show that, with this choice of w^*, the vector $u = (0, \ldots, 0, w_d^*, -w_{d-1}^*)$ satisfies (9). Note that $w \in H_u$ iff $\langle u, w\rangle = w_d^* w_{d-1} - w_{d-1}^* w_d \geq 0$. Since $\langle u, w^*\rangle = 0$, we have $w^* \in H_u$. Moreover, it follows that

$$S_{d-1}^+ \cap H_u = \left\{ w \in S_{d-1}^+ : \frac{w_{d-1}}{w_d} \geq \frac{w_{d-1}^*}{w_d^*} > 0 \right\} \ .$$

It is obvious that no vector $w \in S_{d-1}^+ \cap H_u$ can have a d-th component w_d exceeding $w_d^* = h$ and that setting $w_d = h = w_d^*$ forces the settings $w_{d-1} = w_{d-1}^* = \sqrt{1-h^2}$ and $w_1 = \ldots = w_{d-2} = 0$.

Consequently, (9) is satisfied, which concludes the proof. ■

With this lemma in hand, we can now prove an upper bound of 2 for the preference-based teaching dimension of the class of homogeneous halfspaces, independent of the underlying dimension d.

Theorem 22 $\mathrm{PBTD}(\mathcal{H}_1^0) = \mathrm{TD}(\mathcal{H}_1^0) = 1$ *and, for every* $d \geq 2$*, we have* $\mathrm{PBTD}(\mathcal{H}_d^0) \leq 2$.

Proof Clearly, $\mathrm{PBTD}(\mathcal{H}_1^0) = \mathrm{TD}(\mathcal{H}_1^0) = 1$ since $\mathcal{H}_1^0$ consists of the two sets $\{x \in \mathbb{R} : x \geq 0\}$ and $\{x \in \mathbb{R} : x \leq 0\}$.
Suppose now that $d \geq 2$. Let w^* be the target weight vector (i.e., the weight vector that has to be taught). Under the following conditions, we may assume without loss of generality that $w_d^* \neq 0$:

- For any $0 < s_1 < s_2$, the student prefers any weight vector that ends with s_2 zero coordinates over any weight vector that ends with only s_1 zero coordinates.
- If the target vector ends with (exactly) s zero coordinates, then the teacher presents only examples ending with (at least) s zero coordinates.

In the sequel, we specify a student and a teacher such that these conditions hold, so that we will consider only target weight vectors w^* with $w_d^* \neq 0$.
The student has the following preference relation:

- Among the weight vectors w with $w_d \neq 0$, the student prefers vectors with larger values of $|w_d|$ over those with smaller values of $|w_d|$.

The teacher will use two examples. The first one is chosen as

$$\begin{cases} (-\vec{e}_d, -) & \text{if } w_d^* > 0 \\ (\vec{e}_d, -) & \text{if } w_d^* < 0 \end{cases} .$$

This example reveals whether the unknown weight vector $w^* \in S_{d-1}$ has a strictly positive or a strictly negative d-th component. For reasons of symmetry, we may assume that $w_d^* > 0$. We are now precisely in the situation that is described in Lemma 21. Given w^* and $h = w_d^*$, the teacher picks as a second example $(u, +)$ where $u \in \mathbb{R}^d \setminus \{\vec{0}\}$ has the properties described in the lemma. It follows immediately that the student's preferences will make her choose the weight vector w^*. ■

The upper bound of 2 given in Theorem 22 is tight, as is stated in the following lemma.

Lemma 23 *For every* $d \geq 2$*, we have* $\mathrm{PBTD}(\mathcal{H}_d^0) \geq 2$.

Proof We verify this lemma via Lemma 7, by providing a finite subclass $\mathcal{F}$ of $\mathcal{H}_2^0$ such that $\mathrm{TD}_{min}(\mathcal{F}) = 2$. Let $\mathcal{F} = \{H_w : \vec{0} \neq w \in \{-1, 0, 1\}^2\}$. It is easy to verify that each of the 8 halfspaces in $\mathcal{F}$ has a teaching dimension of 2 with respect to $\mathcal{F}$. This example can be extended to higher dimensions in the obvious way. ■

We thus conclude that the class of homogeneous halfspaces has a preference-based teaching dimension of 2, independent of the dimensionality $d \geq 2$.

Corollary 24 *For every $d \geq 2$, we have* $\mathrm{PBTD}(\mathcal{H}_d^0) = 2$.

By contrast, we will show next that the recursive teaching dimension of the class of homogeneous halfspaces grows with the dimensionality.

Theorem 25 *For any $d \geq 2$,* $\mathrm{TD}(\mathcal{H}_d^0) = \mathrm{RTD}(\mathcal{H}_d^0) = d + 1$.

Proof Assume by normalization that the target weight vector has norm 1, i.e., it is taken from S_{d-1}. Remark 20 implies that all weight vectors in S_{d-1} are equally hard to teach. It suffices therefore to show that $\mathrm{TD}(H_{\vec{e}_1}, \mathcal{H}_d^0) = d + 1$.

We first show that $\mathrm{TD}(H_{\vec{e}_1}, \mathcal{H}_d^0) \leq d + 1$. Define $u = -\sum_{i=2}^{d} \vec{e}_i$. We claim that $T = \{(\vec{e}_i, +) : 2 \leq i \leq d\} \cup \{(u, +), (\vec{e}_1, +)\}$ is a teaching set for $H_{\vec{e}_1}$ w.r.t. $\mathcal{H}_d^0$. Consider any $w \in S_{d-1}$ such that H_w is consistent with T. Note that $w_i = \langle \vec{e}_i, w \rangle \geq 0$ for all $i \in \{2, \ldots, d\}$ and $\langle u, w \rangle = -\sum_{i=2}^{d} w_i \geq 0$ together imply that $w_i = 0$ for all $i \in \{2, \ldots, d\}$ and therefore $w = \pm\vec{e}_1$. Furthermore, $w_1 = \langle w, \vec{e}_1 \rangle \geq 0$, and so $w = \vec{e}_1$, as required.

Now we show that $\mathrm{TD}(H_{\vec{e}_1}, \mathcal{H}_d^0) \geq d + 1$ holds for all $d \geq 2$. It is easy to see that two examples do not suffice for distinguishing $\vec{e}_1 \in \mathbb{R}^2$ from all weight vectors in S_1. In other words, $\mathrm{TD}(H_{\vec{e}_1}, \mathcal{H}_2^0) \geq 3$. Suppose now that $d \geq 3$. It is furthermore easy to see that a teaching set T which distinguishes $\vec{e}_1$ from all weight vectors in S_{d-1} must contain at least one positive example u that is orthogonal to $\vec{e}_1$. The inequality $\mathrm{TD}(H_{\vec{e}_1}, \mathcal{H}_d^0) \geq d + 1$ is now obtained inductively because the example $(u, +) \in T$ leaves open a problem that is not easier than teaching $\vec{e}_1$ w.r.t. the $(d-2)$-dimensional sphere $\{x \in S_{d-1} : x \perp u\}$. ∎

We have thus established that the class of homogeneous halfspaces has a recursive teaching dimension growing linearly with d, while its preference-based teaching dimension is constant. In the case of general (i.e., not necessarily homogeneous) d-dimensional halfspaces, the difference between RTD and PBTD is even more extreme. On the one hand, by generalizing the proof of Lemma 10, it is easy to see that $\mathrm{RTD}(\mathcal{H}_d) = \infty$ for all $d \geq 1$. On the other hand, we will show in the remainder of this section that $\mathrm{PBTD}(\mathcal{H}_d) \leq 6$, independent of the value of d.

We will assume in the sequel (by way of normalization) that an inhomogeneous halfspace has a bias $b \in \{\pm 1\}$. We start with the following result:

Lemma 26 *Let $w^* \in \mathbb{R}^d$ be a vector with a non-trivial d-th component $w_d^* \neq 0$ and let $b^* \in \{\pm 1\}$ be a bias. Then there exist three examples labeled according to H_{w^*,b^*} such that the following holds. Every weight-bias pair (w, b) consistent with these examples satisfies $b = b^*$, $\mathrm{sign}(w_d) = \mathrm{sign}(w_d^*)$ and*

$$\begin{cases} |w_d| \geq |w_d^*| & \text{if } b^* = -1 \\ |w_d| \leq |w_d^*| & \text{if } b^* = +1 \end{cases} . \tag{10}$$

Proof Within the proof, we use the label "1" instead of "+" and the label "−1" instead of "−". The pair (w, b) denotes the student's hypothesis for the target weight-bias pair (w^*, b^*). The examples shown to the student will involve the unknown quantities w^* and b^*. Each example will lead to a new constraint on w and b. We will see that the collection of these constraints reveals the required information. We proceed in three stages:

1. The first example is chosen as $(\vec{0}, b^*)$. The pair (w, b) can be consistent with this example only if $b = -1$ in the case that $b^* = -1$ and $b \in \{0, 1\}$ in the case that $b^* = 1$.

2. The next example is chosen as $\vec{a}_2 = -\frac{2b^*}{w^*_d} \cdot \vec{e}_d$ and labeled "$-b^*$". Note that $\langle w^*, \vec{a}_2\rangle + b^* = -b^*$. We obtain the following new constraint:

$$\langle w, \vec{a}_2\rangle + b = \begin{cases} -2\frac{w_d}{w^*_d} + \overbrace{b}^{\in\{0,1\}} < 0 & \text{if } b^* = 1 \\ +2\frac{w_d}{w^*_d} + \underbrace{b}_{=-1} \geq 0 & \text{if } b^* = -1 \end{cases} .$$

The pair (w, b) with $b = b^*$ if $b^* = -1$ and $b \in \{0, 1\}$ if $b^* = 1$ can satisfy the above constraint only if the sign of w_d equals the sign of w^*_d.

3. The third example is chosen as the example $\vec{a}_3 = -\frac{b^*}{w^*_d} \cdot \vec{e}_d$ with label "1". Note that $\langle w^*, \vec{a}_3\rangle^* + b^* = 0$. We obtain the following new constraint:

$$\langle w, \vec{a}_3\rangle = -\frac{b^* w_d}{w^*_d} + b \geq 0 \ .$$

Given that w is already constrained to weight vectors satisfying $\mathrm{sign}(w_d) = \mathrm{sign}(w^*_d)$, we can safely replace w_d/w^*_d by $|w_d|/|w^*_d|$. This yields $|w_d|/|w^*_d| \leq b$ if $b^* = 1$ and $|w_d|/|w^*_d| \geq -b$ if $b^* = -1$. Since b is already constrained as described in stage 1 above, we obtain $|w_d|/|w^*_d| \leq b \in \{0, 1\}$ if $b^* = 1$ and $|w_d|/|w^*_d| \geq -b = 1$ if $b^* = -1$. The weight-bias pair (w, b) satisfies these constraints only if $b = b^*$ and if (10) is valid.

The assertion of the lemma is immediate from this discussion. ■

Theorem 27 $\mathrm{PBTD}(\mathcal{H}_d) \leq 6$.

Proof As in the proof of Lemma 26, we use the label "1" instead of "+" and the label "-1" instead of "$-$". As in the proof of Theorem 22, we may assume without loss of generality that the target weight vector $w^* \in \mathbb{R}^d$ satisfies $w^*_d \neq 0$. The proof will proceed in stages. On the way, we specify six rules which determine the preference relation of the student.

Stage 1 is concerned with teaching homogeneous halfspaces given by w^* (and $b^* = 0$). The student respects the following rules:

Rule 1: She prefers any pair $(w, 0)$ over any pair (w', b) with $b \neq 0$. In other words, any homogeneous halfspace is preferred over any non-homogeneous halfspace.

Rule 2: Among homogeneous halfspaces, her preferences are the same as the ones that were used within the proof of Theorem 22 for teaching homogeneous halfspaces.

Thus, if $b^* = 0$, then we can simply apply the teaching protocol for homogeneous halfspaces. In this case, w^* can be taught at the expense of only two examples.

Stage 1 reduces the problem to teaching inhomogeneous halfspaces given by (w^*, b^*) with $b^* \neq 0$. We assume, by way of normalization, that $b^* \in \{\pm 1\}$, but note that w^* can now not be assumed to be of unit (or any other fixed) length.

In **stage 2**, the teacher presents three examples in accordance with Lemma 26. It follows that the student will take into consideration only weight-bias pairs (w, b) such that the constraints $b = b^*$, $\mathrm{sign}(w_d) = \mathrm{sign}(w^*_d)$ and (10) are satisfied. The following rule will then induce the constraint $w_d = w^*_d$:

Rule 3: Among the pairs (w, b) such that $w_d \neq 0$ and $b \in \{\pm 1\}$, the student's preferences are as follows. If $b = -1$ (resp. $b = 1$), then she prefers vectors w with a smaller (resp. larger) value of $|w_d|$ over those with a larger (resp. smaller) value of $|w_d|$.

Thanks to Lemma 26 and thanks to Rule 3, we may from now on assume that $b = b^*$ and $w_d = w_d^*$. In the sequel, let w^* be decomposed according to $w^* = (\vec{w}_{d-1}^*, w_d^*) \in \mathbb{R}^{d-1} \times \mathbb{R}$. We think of $\vec{w}_{d-1}$ as the student's hypothesis for $\vec{w}_{d-1}^*$.

Stage 3 is concerned with the special case where $\vec{w}_{d-1}^* = \vec{0}$. The student will automatically set $\vec{w}_{d-1} = \vec{0}$ if we add the following to the student's rule system:

Rule 4: Given that the values for w_d and b have been fixed already (and are distinct from 0), the student prefers weight-bias pairs with $\vec{w}_{d-1} = \vec{0}$ over any weight-bias pair with $\vec{w}_{d-1} \neq \vec{0}$.

Stage 3 reduces the problem to teaching (w^*, b^*) with fixed non-zero values for w_d and b^* (known to the student) and with $\vec{w}_{d-1}^* \neq \vec{0}$. Thus, essentially, only $\vec{w}_{d-1}^*$ has still to be taught. In the next stage, we will argue that the problem of teaching $\vec{w}_{d-1}^*$ is equivalent to teaching a homogeneous halfspace.

In **stage 4**, the teacher will present only examples a such that $a_d = -\frac{b^*}{w_d^*}$ so that the contribution of the d-th component to the inner product of w^* and a cancels with the bias b^*. Given this commitment for a_d, the first $d-1$ components of the examples can be chosen so as to teach the homogeneous halfspace $H_{\vec{w}_{d-1}^*}$. According to Theorem 22, this can be achieved at the expense of two more examples. Of course the student's preferences must match with the preferences that were used in the proof of this theorem:

Rule 5: Suppose that the values of w_d and b have been fixed already (and are distinct from 0) and suppose that $\vec{w}_{d-1} \neq \vec{0}$. Then the preferences for the choice of $\vec{w}_{d-1}$ match with the preferences that were used in the protocol for teaching homogeneous halfspaces.

After stage 4, the student takes into consideration only weight-bias pairs (w, b) such that $w_d = w_d^*$, $b = b^*$ and $H_{\vec{w}_{d-1}} = H_{\vec{w}_{d-1}^*}$. However, since we had normalized the bias and not the weight vector, this does not necessarily mean that $\vec{w}_{d-1} = \vec{w}_{d-1}^*$. On the other hand, the two weight vectors already coincide modulo a positive scaling factor, say

$$\vec{w}_{d-1} = s \cdot \vec{w}_{d-1}^* \text{ for some } s > 0 \ . \tag{11}$$

In order to complete the proof, it suffices to teach the L_1-norm of $\vec{w}_{d-1}^*$ to the student (because (11) and $\|\vec{w}_{d-1}\|_1 = \|\vec{w}_{d-1}^*\|_1$ imply that $\vec{w}_{d-1} = \vec{w}_{d-1}^*$). The next (and final) stage serves precisely this purpose.

As for **stage 5**, we first fix some notation. For $i = 1, \ldots, k-1$, let $\beta_i = \mathrm{sign}(w_i^*)$. Note that (11) implies that $\beta_i = \mathrm{sign}(w_i)$. Let $L = \|\vec{w}_{d-1}^*\|_1$ denote the L_1-norm of $\vec{w}_{d-1}^*$. The final example is chosen as $\vec{a}_6 = (\beta_1, \ldots, \beta_{d-1}, -(L + b^*)/w_d^*)$ and labeled "1". Note that

$$\left\langle w^*, \vec{a}_6 \right\rangle + b^* = |w_1^*| + \ldots + |w_{d-1}^*| - L = 0 \ .$$

Given that $\beta_i = \mathrm{sign}(w_i)$, $w_d = w_d^*$ and $b = b^*$, the student can derive from $\vec{a}_6$ and its label the following constraint on $\vec{w}_{d-1}$:

$$\langle w, \vec{a}_6 \rangle + b = |w_1| + \ldots + |w_{d-1}| - L \geq 0 \ .$$

In combination with the following rule, we can now force the constraint $\|\vec{w}_{d-1}\|_1 = L$:

Rule 6: Suppose that the values of w_d and b have been fixed already (and are distinct from 0) and suppose that $H_{\vec{w}_{d-1}}$ has already been fixed. Then, among the vectors representing $H_{\vec{w}_{d-1}}$, the ones with a smaller L_1-norm are preferred over the ones with a larger L_1-norm.

An inspection of the six stages reveals that at most six examples altogether were shown to the student (three in stage 2, two in stage 4, and one in stage 5). This completes the proof of the theorem. ∎

Note that Theorems 22 and 27 remain valid when we allow w to be the all-zero vector, which extends $\mathcal{H}_d^0$ by $\{\mathbb{R}^d\}$ and $\mathcal{H}_d$ by $\{\mathbb{R}^d, \emptyset\}$. $\mathbb{R}^d$ will be taught with a single positive example, and $\emptyset$ with a single negative example. The student will give the highest preference to $\mathbb{R}^d$, the second highest to $\emptyset$, and among the remaining halfspaces, the student's preferences stay the same.

8. Classes with PBTD or PBTD^+ Equal to One

In this section, we will give complete characterizations of (i) the concept classes with a positive preference-based teaching dimension of 1, and (ii) the concept classes with a preference-based teaching dimension of 1. Throughout this section, we use the label "1" to indicate positive examples and the label "0" to indicate negative examples.

Let I be a (possibly infinite) index set. We will consider a mapping $A : I \times I \to \{0,1\}$ as a binary matrix $A \in \{0,1\}^{I \times I}$. A is said to be *lower-triangular* if there exists a linear ordering $\prec$ on I such that $A(i,i') = 0$ for every pair (i,i') such that $i \prec i'$.

We will occasionally identify a set $L \subseteq \mathcal{X}$ with its indicator function by setting $L(x) = \mathbb{1}_{[x \in L]}$. For each $M \subseteq \mathcal{X}$, we define

$$M \oplus L = (L \setminus M) \cup (M \setminus L)$$

and

$$M \oplus \mathcal{L} = \{M \oplus L : L \in \mathcal{L}\} \ .$$

For $T \subseteq \mathcal{X} \times \{0,1\}$, we define similarly

$$M \oplus T = \{(x, \bar{y}) : (x,y) \in T \text{ and } x \in M\} \cup \{(x,y) \in T : x \notin M\} \ .$$

Moreover, given $M \subseteq \mathcal{X}$ and a linear ordering $\prec$ on $\mathcal{L}$, we define a linear ordering $\prec_M$ on $M \oplus \mathcal{L}$ as follows:

$$M \oplus L' \prec_M M \oplus L \Leftrightarrow \underbrace{M \oplus (M \oplus L')}_{=L'} \prec \underbrace{M \oplus (M \oplus L)}_{=L} \ .$$

Lemma 28 *With this notation, the following holds. If the mapping $\mathcal{L} \ni L \mapsto T(L) \subseteq \mathcal{X} \times \{0,1\}$ assigns a teaching set to L w.r.t. $(\mathcal{L}, \prec)$, then the mapping $M \oplus \mathcal{L} \ni M \oplus L \mapsto M \oplus T(L) \subseteq \mathcal{X} \times \{0,1\}$ assigns a teaching set to $M \oplus L$ w.r.t. $(M \oplus \mathcal{L}, \prec_M)$.*

Since this result is rather obvious, we skip its proof.

We say that $\mathcal{L}$ and $\mathcal{L}'$ are *equivalent* if $\mathcal{L}' = M \oplus \mathcal{L}$ for some $M \subseteq \mathcal{X}$ (and this clearly is an equivalence relation). As an immediate consequence of Lemma 28, we obtain the following result:

Lemma 29 *If $\mathcal{L}$ is equivalent to $\mathcal{L}'$, then $\mathrm{PBTD}(\mathcal{L}) = \mathrm{PBTD}(\mathcal{L}')$.*

The following lemma provides a necessary condition for a concept class to have a preference-based teaching dimension of one.

Lemma 30 *Suppose that $\mathcal{L} \subseteq 2^{\mathcal{X}}$ is a concept class of* $\mathrm{PBTD}\ 1$. *Pick a linear ordering $\prec$ on $\mathcal{L}$ and a mapping $\mathcal{L} \ni L \mapsto (x_L, y_L) \in \mathcal{X} \times \{0,1\}$ such that, for every $L \in \mathcal{L}$, $\{(x_L, y_L)\}$ is a teaching set for L w.r.t. $(\mathcal{L}, \prec)$. Then*

- *either every instance $x \in \mathcal{X}$ occurs at most once in $(x_L)_{L \in \mathcal{L}}$*
- *or there exists a concept $L^* \in \mathcal{L}$ that is preferred over all other concepts in $\mathcal{L}$ and x_{L^*} is the only instance from $\mathcal{X}$ that occurs twice in $(x_L)_{L \in \mathcal{L}}$.*

Proof Since the mapping T must be injective, no instance can occur twice in $(x_L)_{L \in \mathcal{L}}$ with the same label. Suppose that there exists an instance $x \in \mathcal{X}$ and concepts $L \prec L^*$ such that $x = x_L = x_{L^*}$ and, w.l.o.g., $y_L = 1$ and $y_{L^*} = 0$. Since $\{(x,1)\}$ is a teaching set for L w.r.t. $(\mathcal{L}, \prec)$, every concept $L' \succ L$ (including the ones that are preferred over L^*) must satisfy $L'(x) = 0$. For analogous reasons, every concept $L' \succ L^*$ (if any) must satisfy $L'(x) = 1$. A concept $L' \in \mathcal{L}$ that is preferred over L^* would have to satisfy $L'(x) = 0$ and $L'(x) = 1$, which is impossible. It follows that there can be no concept that is preferred over L^*. ■

The following result is a consequence of Lemmas 28 and 30.

Theorem 31 *If* $\mathrm{PBTD}(\mathcal{L}) = 1$, *then there exists a concept class $\mathcal{L}'$ that is equivalent to $\mathcal{L}$ and satisfies* $\mathrm{PBTD}(\mathcal{L}') = \mathrm{PBTD}^+(\mathcal{L}') = 1$.

Proof Pick a linear ordering $\prec$ on $\mathcal{L}$ and, for every $L \in \mathcal{L}$, a pair $(x_L, y_L) \in \mathcal{X} \times \{0,1\}$ such that $T(L) = \{(x_L, y_L)\}$ is a teaching set for L w.r.t. $(\mathcal{L}, \prec)$.

Case 1: Every instance $x \in \mathcal{X}$ occurs at most once in $(x_L)_{L \in \mathcal{L}}$.
Then choose $M = \{x_L : y_L = 0\}$ and apply Lemma 28.

Case 2: There exists a concept $L^* \in \mathcal{L}$ that is preferred over all other concepts in $\mathcal{L}$ and x_{L^*} is the only instance from $\mathcal{X}$ that occurs twice in $(x_L)_{L \in \mathcal{L}}$.
Then choose $M = \{x_L : y_L = 0 \wedge L \neq L^*\}$ and apply Lemma 28. With this choice, we obtain $M \oplus T(L) = \{(x_L, 1)\}$ for every $L \in \mathcal{L} \setminus \{L^*\}$. Since L^* is preferred over all other concepts in $\mathcal{L}$, we may teach L^* w.r.t. $(\mathcal{L}, \prec)$ by the empty set (instead of employing a possibly 0-labeled example).

The discussion shows that there is a class $\mathcal{L}'$ that is equivalent to $\mathcal{L}$ and can be taught in the preference-based model with positive teaching sets of size 1 (or size 0 in case of L^*). ■

We now have the tools required for characterizing the concept classes whose positive PBTD equals 1.

Theorem 32 $\mathrm{PBTD}^+(\mathcal{L}) = 1$ *if and only if there exists a mapping $\mathcal{L} \ni L \mapsto x_L \in \mathcal{X}$ such that the matrix $A \in \{0,1\}^{(\mathcal{L} \setminus \{\emptyset\}) \times (\mathcal{L} \setminus \{\emptyset\})}$ given by $A(L, L') = L'(x_L)$ is lower-triangular.*

Proof Suppose first that $\mathrm{PBTD}^{+}(\mathcal{L}) = 1$. Pick a linear ordering $\prec$ on $\mathcal{L}$ and, for every $L \in \mathcal{L}\setminus\{\emptyset\}$, pick $x_L \in \mathcal{X}$ such that $\{x_L\}$ is a positive teaching set for L w.r.t. $(\mathcal{L}, \prec)$.[5] If $L \prec L'$ (so that L' is preferred over L), we must have $L'(x_L) = 0$. It follows that the matrix A, as specified in the theorem, is lower-triangular.

Suppose conversely that there exists a mapping $\mathcal{L} \ni L \mapsto x_L \in \mathcal{X}$ such that the matrix $A \in \{0,1\}^{(\mathcal{L}\setminus\{\emptyset\})\times(\mathcal{L}\setminus\{\emptyset\})}$ given by $A(L, L') = L'(x_L)$ is lower-triangular, say w.r.t. the linear ordering $\prec$ on $\mathcal{L} \setminus \{\emptyset\}$. Then, for every $L \in \mathcal{L} \setminus \{\emptyset\}$, the singleton $\{x_L\}$ is a positive teaching set for L w.r.t. $(\mathcal{L}, \prec)$ because it distinguishes L from $\emptyset$ (of course) and also from every concept $L' \in \mathcal{L} \setminus \{\emptyset\}$ such that $L' \succ L$. If $\emptyset \in \mathcal{L}$, then extend the linear ordering $\prec$ by preferring $\emptyset$ over every other concept from $\mathcal{L}$ (so that $\emptyset$ is a positive teaching set for $\emptyset$ w.r.t. $(\mathcal{L}, \prec)$). ■

In view of Theorem 31, Theorem 32 characterizes every class $\mathcal{L}$ with $\mathrm{PBTD}(\mathcal{L}) = 1$ up to equivalence.

Let $\mathrm{Sg}(\mathcal{X}) = \{\{x\} : x \in \mathcal{X}\}$ denote the class of singletons over $\mathcal{X}$ and suppose that $\mathrm{Sg}(\mathcal{X})$ is a sub-class of $\mathcal{L}$ and $\mathrm{PBTD}(\mathcal{L}) = 1$. We will show that only fairly trivial extensions of $\mathrm{Sg}(\mathcal{X})$ with a preference-based dimension of 1 are possible.

Lemma 33 *Let $\mathcal{L} \subseteq 2^{\mathcal{X}}$ be a concept class of* $\mathrm{PBTD}\ 1$ *that contains* $\mathrm{Sg}(\mathcal{X})$. *Let T be an admissible mapping for $\mathcal{L}$ that assigns a labeled example $(x_L, y_L) \in \mathcal{X} \times \{0,1\}$ to each $L \in \mathcal{L}$. For $b = 0, 1$, let $\mathcal{L}^b = \{L \in \mathcal{L} : y_L = b\}$. Similarly, let $\mathcal{X}^b = \{x \in \mathcal{X} : y_{\{x\}} \in \mathcal{L}^b\}$. With this notation, the following holds:*

1. *If $L \in \mathcal{L}^1$ and $L \subset L' \in \mathcal{L}$, then $L' \in \mathcal{L}^1$.*
2. *If $L' \in \mathcal{L}^0$ and $L' \supset L \in \mathcal{L}$, then $L \in \mathcal{L}^0$.*
3. *$|\mathcal{X}^0| \le 2$. Moreover if $|\mathcal{X}^0| = 2$, then there exist $q \neq q' \in \mathcal{X}$ such that $\mathcal{X}^0 = \{q, q'\}$ and $x_{\{q\}} = q'$.*

Proof Recall that $R_T = \{(L, L') \in \mathcal{L} \times \mathcal{L} :\ (L \neq L') \wedge (L \text{ is consistent with } T(L'))\}$ and that R_T (and even the transitive closure of R_T) is asymmetric if T is admissible.

1. If $L \in \mathcal{L}^1$ and $L \subset L'$, then $y_L = 1$ so that L' is consistent with the example (x_L, y_L). It follows that $(L', L) \in R_T$. $L' \in \mathcal{L}^0$ would similarly imply that $(L, L') \in R_T$ so that R_T would not be asymmetric. This is in contradiction with the admissibility of T.
2. The second assertion in the lemma is a logically equivalent reformulation of the first assertion.
3. Suppose for the sake of contradiction that $\mathcal{X}^0$ contains three distinct points, say q_1, q_2, q_3. Since, for $i = 1, 2, 3$, T assigns a 0-labeled example to $\{q_i\}$, at least one of the remaining two points is consistent with $T(\{q_i\})$. Let G be the digraph with the nodes q_1, q_2, q_3 and with an edge from q_j to q_i iff $\{q_j\}$ is consistent with $T(\{q_i\})$. Then each of the three nodes has an indegree of at least 1. Digraphs of this form must contain a cycle so that $\mathrm{trcl}(R_T)$ is not asymmetric. This is in contradiction with the admissibility of R_T.

5. Such an x_L always exists, even if $\emptyset$ is a teaching set for L, because every superset of a teaching set for L that is still consistent with L is still a teaching set for L, cf. the discussion immediately after Lemma 4.

A similar argument holds if $\mathcal{X}^0$ contains only two distinct elements, say q and q'. If neither $x_{\{q\}} = q'$ nor $x_{\{q'\}} = q$, then $(\{q'\}, \{q\}) \in R_T$ and $(\{q\}, \{q'\}) \in R_T$ so that R_T is not asymmetric — again a contradiction to the admissibility of R_T.

■

We are now in the position to characterize those classes of PBTD one that contain all singletons.

Theorem 34 *Suppose that $\mathcal{L} \subseteq 2^{\mathcal{X}}$ is a concept class that contains $\mathrm{Sg}(\mathcal{X})$. Then $\mathrm{PBTD}(\mathcal{L}) = 1$ if and only if the following holds. Either $\mathcal{L}$ coincides with $\mathrm{Sg}(\mathcal{X})$ or $\mathcal{L}$ contains precisely one additional concept, which is either the empty set or a set of size 2.*

Proof We start with proving "$\Leftarrow$". It is well known that $\mathrm{PBTD}^+(\mathcal{L}) = 1$ for $\mathcal{L} = \mathrm{Sg}(\mathcal{X}) \cup \{\emptyset\}$: prefer $\emptyset$ over any singleton set, set $T(\emptyset) = \emptyset$ and, for every $x \in \mathcal{X}$, set $T(\{x\}) = \{(x, 1)\}$. In a similar fashion, we can show that $\mathrm{PBTD}(\mathcal{L}) = 1$ for $\mathcal{L} = \mathrm{Sg}(\mathcal{X}) \cup \{\{q, q'\}\}$ for any choice of $q \neq q' \in \mathcal{X}$. Prefer $\{q, q'\}$ over $\{q\}$ and $\{q'\}$, respectively. Furthermore, prefer $\{q\}$ and $\{q'\}$ over all other singletons. Finally, set $T(\{q, q'\}) = \emptyset$, $T(\{q\}) = \{(q', 0)\}$, $T(\{q'\}) = \{(q, 0)\}$ and, for every $x \in \mathcal{X} \setminus \{q, q'\}$, set $T(\{x\}) = \{(x, 1)\}$.

As for the proof of "$\Rightarrow$", we make use of the notions $T, x_L, y_L, \mathcal{L}^0, \mathcal{L}^1, \mathcal{X}^0, \mathcal{X}^1$ that had been introduced in Lemma 33 and we proceed by case analysis.

Case 1: $\mathcal{X}^0 = \emptyset$.

Since $\mathcal{X}^0 = \emptyset$, we have $\mathcal{X} = \mathcal{X}^1$. In combination with the first assertion in Lemma 33, it follows that $\mathcal{L} \setminus \{\emptyset\} = \mathcal{L}^1$. We claim that no concept in $\mathcal{L}$ contains two distinct elements. Assume for the sake of contradiction that there is a concept $L \in \mathcal{L}$ such that $|L| \geq 2$. It follows that, for every $q \in L$, $x_{\{q\}} = q$ and $y_{\{q\}} = 1$ so that $(L, \{q\}) \in R_T$. Moreover, there exists $q_0 \in L$ such that $x_L = q_0$ and $y_L = 1$. It follows that $(\{q_0\}, L) \in R_T$, which contradicts the fact that R_T is asymmetric.

Case 2: $\mathcal{X}^0 = \{q\}$ for some $q \in \mathcal{X}$.

Set $q' = x_{\{q\}}$ and note that $y_{\{q\}} = 0$. Moreover, since $\mathcal{X}^1 = \mathcal{X} \setminus \{q\}$, we have $x_{\{p\}} = p$ and $y_{\{p\}} = 1$ for every $p \in \mathcal{X} \setminus \{q\}$. We claim that $\mathcal{L}$ cannot contain a concept L of size at least 2 that contains an element of $\mathcal{X} \setminus \{q, q'\}$. Assume for the sake of contradiction, that there is a set L such that $|L| \geq 2$ and $p \in L$ for some $p \in \mathcal{X} \setminus \{q, q'\}$. The first assertion in Lemma 33 implies that $y_L = 1$ (because $y_{\{p\}} = 1$ and $\{p\} \subseteq L$). Since all pairs $(x, 1)$ with $x \neq q$ are already in use for teaching the corresponding singletons, we may conclude that $q \in L$ and $T(L) = \{(q, 1)\}$. This contradicts the fact that $\mathrm{trcl}(R_T)$ is asymmetric, because our discussion implies that $(L, \{p\}), (\{p\}, \{q\}), (\{q\}, L) \in R_T$. We may therefore safely assume that there is no concept of size at least 2 in $\mathcal{L}$ that has a non-empty intersection with $\mathcal{X} \setminus \{q, q'\}$. Thus, except for the singletons, the only remaining sets that possibly belong to $\mathcal{L}$ are $\emptyset$ and $\{q, q'\}$. We still have to show that not both of them can belong to $\mathcal{L}$. Assume for the sake of contradiction that $\emptyset, \{q, q'\} \in \mathcal{L}$. Since $\emptyset$ is consistent with $T(\{q\}) = \{(q', 0)\}$, we have $(\emptyset, \{q\}) \in R_T$. Clearly, $y_\emptyset = 0$. Since $\{q\}$ is consistent with every pair $(x, 0)$ except for $(q, 0)$, we must have $x_\emptyset = q$. (Otherwise, we have $(\{q\}, \emptyset) \in R_T$ and arrive at a contradiction.) Let us now inspect the possible teaching sets for $L = \{q, q'\}$. Since $\{q, q'\}$ is consistent with $T(\{q'\}) = \{(q', 1)\}$, setting $y_L = 0$ would lead to a contradiction. The example $(q', 1)$

is already in use for teaching $\{q'\}$. It is therefore necessary to set $T(L) = \{(q, 1)\}$. An inspection of the various teaching sets shows that $(\emptyset, \{q\}), (\{q\}, L), (L, \{q'\}), (\{q'\}, \emptyset) \in R_T$, which contradicts the fact that $\mathrm{trcl}(R_T)$ is asymmetric.

Case 3: $\mathcal{X}^0 = \{q, q'\}$ for some $q \neq q' \in \mathcal{X}$.
Note first that $y_{\{q\}} = y_{\{q'\}} = 0$ and $y_{\{p\}} = 1$ for every $p \in \mathcal{X} \setminus \{q, q'\}$. We claim that $\emptyset \notin \mathcal{L}$. Assume for the sake of contradiction that $\emptyset \in \mathcal{L}$. Then $(\emptyset, \{q\}), (\emptyset, \{q'\}) \in R_T$ since $\emptyset$ is consistent with the teaching sets for instances from $\mathcal{X}^0$. But then, no matter how x in $T(\emptyset) = \{(x, 0)\}$ is chosen, at least one of the sets $\{q\}$ and $\{q'\}$ will be consistent with $T(\emptyset)$ so that at least one of the pairs $(\{q\}, \emptyset)$ and $(\{q'\}, \emptyset)$ belongs to R_T. This contradicts the fact that R_T must be asymmetric. Thus $\emptyset \notin \mathcal{L}$, indeed. Now it suffices to show that $\mathcal{L}$ cannot contain a concept of size at least 2 that contains an element of $\mathcal{X} \setminus \{q, q'\}$. Assume for the sake of contradiction that there is a set $L \in \mathcal{L}$ such that $|L| \geq 2$ and $p \in L$ for some $p \in \mathcal{X} \setminus \{q, q'\}$. Observe that $(L, \{p\}) \in R_T$. Another application of the first assertion in Lemma 33 shows that $y_L = 1$ (because $y_{\{p\}} = 1$ and $p \in L$) and $x_L \in \{q, q'\}$ (because the other 1-labeled instances are already in use for teaching the corresponding singletons). It follows that one of the pairs $(\{q\}, L)$ and $(\{q'\}, L)$ belongs to R_T. The third assertion of Lemma 33 implies that $T(q) = \{(q', 0)\}$ or $T(q') = \{(q, 0)\}$. For reasons of symmetry, we may assume that $T(q) = \{(q', 0)\}$. This implies that $(\{p\}, \{q\}) \in R_T$. Let q'' be given by $T(q') = \{(q'', 0)\}$. Note that either $q'' = q$ or $q'' \in \mathcal{X} \setminus \{q, q'\}$. In the former case, we have that $(\{p\}, \{q'\}) \in R_T$ and in the latter case we have that $(\{q\}, \{q'\}) \in R_T$. Since $(\{p\}, \{q\}) \in R_T$ (which was observed above already), we conclude that in both cases, $(\{p\}, \{q\}), (\{p\}, \{q'\}) \in \mathrm{trcl}(R_T)$. Combining this with our observations above that $(L, \{p\}) \in R_T$ and that one of the pairs $(\{q\}, L)$ and $(\{q'\}, L)$ belongs to R_T, yields a contradiction to the fact that $\mathrm{trcl}(R_T)$ is asymmetric.

■

Corollary 35 *Let* $\mathcal{L} \subseteq 2^{\mathcal{X}}$ *be a concept class that contains* $\mathrm{Sg}(\mathcal{X})$. *If* $\mathrm{PBTD}(\mathcal{L}) = 1$, *then* $\mathrm{RTD}(\mathcal{L}) = 1$.

Proof According to Theorem 34, either L coincides with $\mathrm{Sg}(\mathcal{X})$ or $\mathcal{L}$ contains precisely one additional concept that is $\emptyset$ or a set of size 2. The partial ordering $\prec$ on $\mathcal{L}$ that is used in the first part of the proof of Theorem 34 (proof direction "$\Leftarrow$") is easily compiled into a recursive teaching plan of order 1 for $\mathcal{L}$.[6] ■

The characterizations proven above can be applied to certain geometric concept classes.

Consider a class $\mathcal{L}$, consisting of bounded and topologically closed objects in the d-dimensional Euclidean space, that satisfies the following condition: for every pair $(A, B) \in \mathbb{R}^d$, there is exactly one object in $\mathcal{L}$, denoted as $L_{A,B}$ in the sequel, such that $A, B \in L$ and such that $\|A - B\|$ coincides with the diameter of L. This assumption implies that $|\mathcal{L} \setminus \mathrm{Sg}(\mathbb{R}^d)| = \infty$. By setting $A = B$, it furthermore implies $\mathrm{Sg}(\mathbb{R}^d) \subseteq \mathcal{L}$. Let us prefer objects with a small diameter over objects with a larger diameter. Then, obviously, $\{A, B\}$ is a positive teaching set for $L_{A,B}$. Because

6. This also follows from Lemma 8 and the fact that there are no chains of a length exceeding 2 in $(\mathcal{L}, \prec)$.

of $|\mathcal{L} \setminus \mathrm{Sg}(\mathbb{R}^d)| = \infty$, $\mathcal{L}$ does clearly not satisfy the condition in Theorem 34, which is necessary for $\mathcal{L}$ to have a PBTD of 1. We may therefore conclude that $\mathrm{PBTD}(\mathcal{L}) = \mathrm{PBTD}^+(\mathcal{L}) = 2$.

The family of classes with the required properties is rich and includes, for instance, the class of d-dimensional balls as well as the class of d-dimensional axis-parallel rectangles.

9. Conclusions

Preference-based teaching uses the natural notion of preference relation to extend the classical teaching model. The resulting model is (i) more powerful than the classical one, (ii) resolves difficulties with the recursive teaching model in the case of infinite concept classes, and (iii) is at the same time free of coding tricks even according to the definition by Goldman and Mathias (1996). Our examples of algebraic and geometric concept classes demonstrate that preference-based teaching can be achieved very efficiently with naturally defined teaching sets and based on intuitive preference relations such as inclusion. We believe that further studies of the PBTD will provide insights into structural properties of concept classes that render them easy or hard to learn in a variety of formal learning models.

We have shown that spanning sets lead to a general-purpose construction for preference-based teaching sets of only positive examples. While this result is fairly obvious, it provides further justification of the model of preference-based teaching, since the teaching sets it yields are often intuitively exactly those a teacher would choose in the classroom (for instance, one would represent convex polygons by their vertices, as in Example 3). It should be noted, too, that it can sometimes be difficult to establish whether the upper bound on PBTD obtained this way is tight, or whether the use of negative examples or preference relations other than inclusion yield smaller teaching sets. Generally, the choice of preference relation provides a degree of freedom that increases the power of the teacher but also increases the difficulty of establishing lower bounds on the number of examples required for teaching.

Acknowledgements. Sandra Zilles was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), in the Discovery Grant and Canada Research Chairs programs. We thank the anonymous referees for their numerous thoughtful comments, which greatly helped to improve the presentation of the paper.

References

N. Abe. Polynomial learnability of semilinear sets. In *Proceedings of the 2nd Annual Conference on Learning Theory (COLT)*, pages 25–40, 1989.

D. Angluin. Inductive inference of formal languages from positive data. *Information and Control*, 45(2):117–135, 1980.

X. Chen, Y. Cheng, and B. Tang. On the recursive teaching dimension of vc classes. In *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, pages 2164–2171, 2016.

T. Doliwa, G. Fan, H. U. Simon, and S. Zilles. Recursive teaching dimension, VC-dimension, and sample compression. *Journal of Machine Learning Research*, 15:3107–3131, 2014.

Z. Gao, H. U. Simon, and S. Zilles. On the teaching complexity of linear sets. In *Proceedings of the 26th International Conference on Algorithmic Learning Theory (ALT)*, pages 102–116, 2015.

Z. Gao, C. Ries, H. U. Simon, and S. Zilles. Preference-based teaching. In *Proceedings of the 29th Conference on Learning Theory (COLT)*, pages 971–997, 2016.

Z. Gao, C. Ries, H. U. Simon, and S. Zilles. Preference-based teaching. *ArXiv e-prints*, February 2017. URL `https://arxiv.org/pdf/1702.02047.pdf`.

S. A. Goldman and M. J. Kearns. On the complexity of teaching. *Journal of Computer and System Sciences*, 50:20–31, 1995.

S. A. Goldman and H. D. Mathias. Teaching a smarter learner. *Journal of Computer and System Sciences*, 52:255–267, 1996.

D. Helmbold, R. Sloan, and M. K. Warmuth. Learning nested differences of intersection-closed concept classes. *Machine Learning*, pages 165–196, 1990.

T. J. Jech. *The Axiom of Choice*. North-Holland Pub. Co., Amsterdam, 1973.

Z. Mazadi, Z. Gao, and S. Zilles. Distinguishing pattern languages with membership examples. In *Proceedings of the 8th International Conference on Language and Automata Theory and Applications (LATA)*, pages 528–540, 2014.

S. Moran, A. Shpilka, A. Wigderson, and A. Yehudayoff. Compressing and teaching for low VC-dimension. In *Proceedings of the 56th Annual Symposium on the Foundations of Computer Science (FOCS)*, pages 40–51, 2015.

R. J. Parikh. On context-free languages. *Journal of the ACM*, 13(4):570–581, 1966.

J. G. Rosales and P. A. García-Sánchez. *Numerical Semigroups*. Springer, 2009.

A. Shinohara and S. Miyano. Teachability in computational learning. *New Generation Computing*, 8(4):337–347, 1991.

H. U. Simon and S. Zilles. Open problem: Recursive teaching dimension versus VC dimension. In *Proceedings of the 28th Annual Conference on Learning Theory (COLT)*, pages 1770–1772, 2015.

Y. Takada. Learning semilinear sets from examples and via queries. *Theoretical Computer Science*, 104(2):207–233, 1992.

S. Zilles, S. Lange, R. Holte, and M. Zinkevich. Models of cooperative teaching and learning. *Journal of Machine Learning Research*, 12:349–384, 2011.

Appendix A. Proof of Theorem 18

In Section A.1, we present a general result which helps to verify the upper bounds in Theorem 18. These upper bounds are then derived in Section A.2. Section A.3 is devoted to the derivation of the lower bounds.

A.1 The Shift Lemma

In this section, we assume that $\mathcal{L}$ is a concept class over a universe $\mathcal{X} \in \{\mathbb{N}_0, \mathbb{Q}_0^+, \mathbb{R}_0^+\}$. We furthermore assume that 0 is contained in every concept $L \in \mathcal{L}$. We can extend $\mathcal{L}$ to a larger class, namely the shift-extension $\mathcal{L}'$ of $\mathcal{L}$, by allowing each of its concepts to be shifted by some constant which is taken from $\mathcal{X}$:

$$\mathcal{L}' = \{c + L : (c \in \mathcal{X}) \wedge (L \in \mathcal{L})\} \ .$$

The next result states that this extension has little effect only on the complexity measures PBTD and PBTD^+:

Lemma 36 (Shift Lemma) *With the above notation and assumptions, the following holds:*

$$\mathrm{PBTD}(\mathcal{L}) \le \mathrm{PBTD}(\mathcal{L}') \le 1{+}\mathrm{PBTD}(\mathcal{L}) \ \textit{and} \ \mathrm{PBTD}^+(\mathcal{L}) \le \mathrm{PBTD}^+(\mathcal{L}') \le 1{+}\mathrm{PBTD}^+(\mathcal{L}) \ .$$

Proof It suffices to verify the inequalities $\mathrm{PBTD}(\mathcal{L}') \le 1 + \mathrm{PBTD}(\mathcal{L})$ and $\mathrm{PBTD}^+(\mathcal{L}') \le 1 + \mathrm{PBTD}^+(\mathcal{L})$ because the other inequalities hold by virtue of monotonicity. Let T be an admissible mapping for $\mathcal{L}$. It suffices to show that T can be transformed into an admissible mapping T' for $\mathcal{L}'$ such that $\mathrm{ord}(T') \le 1 + \mathrm{ord}(T)$ and such that T' is positive provided that T is positive. To this end, we define T' as follows:

$$T'(c + L) = \{(c, +)\} \cup \{(c + x, b) : (x, b) \in T(L)\} \ .$$

Obviously $\mathrm{ord}(T') \le 1 + \mathrm{ord}(T)$. Note that $c \in c+L$ because of our assumption that 0 is contained in every concept in $\mathcal{L}$. Moreover, since the admissibility of T implies that L is consistent with $T(L)$, the above definition of $T'(c+L)$ makes sure that $c+L$ is consistent with $T'(c+L)$. It suffices therefore to show that the relation $\mathrm{trcl}(R_{T'})$ is asymmetric. Consider a pair $(c'+L', c+L) \in R_{T'}$. By the definition of $R_{T'}$, it follows that $c'+L'$ is consistent with $T'(c+L)$. Because of $(c,+) \in T'(c+L)$, we must have $c' \le c$. Suppose that $c' = c$. In this case, L' must be consistent with $T(L)$. Thus $L' \prec_T L$. This reasoning implies that $(c' + L', c + L) \in R_{T'}$ can happen only if either $c' < c$ or $(c' = c) \wedge (L' \prec_T L)$. Since $\prec_T$ is asymmetric, we may now conclude that $\mathrm{trcl}(R_{T'})$ is asymmetric, as desired. Finally note that, according to our definition above, the mapping T' is positive provided that T is positive. This concludes the proof. ■

A.2 The Upper Bounds in Theorem 18

We remind the reader that the equality $\mathrm{PBTD}^+(\mathrm{LINSET}_k) = k$ was stated in Example 2. We will show in Lemma 37 that $\mathrm{PBTD}^+(\text{NE-LINSET}_k) \le k$. In combination with the Shift Lemma, this implies that $\mathrm{PBTD}^+(\mathrm{LINSET}'_k) \le k + 1$ and $\mathrm{PBTD}^+(\text{NE-LINSET}'_k) \le k + 1$. All remaining upper bounds in Theorem 18 follow now by virtue of monotonicity.

Lemma 37 $\mathrm{PBTD}^+(\text{NE-LINSET}_k) \le k$.

Proof We want to show that there is a preference relation for which k positive examples suffice to teach any concept in NE-LINSET_k. To this end, let $G = \{g_1, \ldots, g_\ell\}$ be a generator set with $\ell \le k$ where $g_1 < \ldots < g_\ell$. We use $\mathrm{sum}(G) = g_1 + \ldots + g_\ell$ to denote the sum of all generators in G. We say that g_i is a *redundant generator* in G if $g_i \in \langle\{g_1, \ldots, g_{i-1}\}\rangle$. Let $G^* = \{g_1^*, \ldots, g_{\ell^*}^*\} \subseteq G$

with $g_1^* < \ldots < g_{\ell^*}^*$ be the set of non-redundant generators in G and let $\mathrm{tuple}(G) = (g_1^*, \ldots, g_{\ell^*}^*)$ be the corresponding ordered sequence. Then G^* is an independent subset of G generating the same linear set as G when allowing zero coefficients, i.e., we have $\langle G^* \rangle = \langle G \rangle$ (although $\langle G^* \rangle_+ \neq \langle G \rangle_+$ whenever G^* is a proper subset of G).

To define a suitable preference relation, let $G, \widehat{G}$ be generator sets of size k or less with $\mathrm{tuple}(G) = (g_1^*, \ldots, g_{\ell^*}^*)$ and $\mathrm{tuple}(\widehat{G}) = (\widehat{g}_1^*, \ldots, \widehat{g}_{\widehat{\ell}^*}^*)$. Let the student prefer G over $\widehat{G}$ if any of the following conditions is satisfied:

Condition 1: $\mathrm{sum}(G) > \mathrm{sum}(\widehat{G})$.

Condition 2: $\mathrm{sum}(G) = \mathrm{sum}(\widehat{G})$ and $\mathrm{tuple}(G)$ is lexicographically greater than $\mathrm{tuple}(\widehat{G})$ without having $\mathrm{tuple}(\widehat{G})$ as prefix.

Condition 3: $\mathrm{sum}(G) = \mathrm{sum}(\widehat{G})$ and $\mathrm{tuple}(G)$ is a proper prefix of $\mathrm{tuple}(\widehat{G})$.

To teach a concept $\langle G \rangle \in$ NE-LINSET$_k$ with $\mathrm{sum}(G) = g$ and $\mathrm{tuple}(G) = (g_1^*, \ldots, g_{\ell^*}^*)$, one uses the teaching set

$$S = \{(g, +), (g + g_1^*, +), \ldots, (g + g_{h^*}^*, +)\}$$

where

$$h = \begin{cases} \ell^* - 1 & \text{if } G^* = G \\ \ell^* & \text{if } G^* \subset G \end{cases} . \tag{12}$$

Note that S contains at most $|G| \leq k$ examples. Let $\widehat{G}$ with $\left\langle \widehat{G} \right\rangle_+ \in$ NE-LINSET$_k$ denote the generator set that is returned by the student. Clearly $\left\langle \widehat{G} \right\rangle$ satisfies $\mathrm{sum}(\widehat{G}) = g$ since

- concepts with larger generator sums are inconsistent with $(g, +)$, and
- concepts with smaller generator sums have a lower preference (compare with Condition 1 above).

It follows that $g + g_i^* \in \left\langle \widehat{G} \right\rangle_+$ is equivalent to $g_i^* \in \left\langle \widehat{G} \right\rangle = \left\langle \widehat{G}^* \right\rangle$. We conclude that the smallest generator in $\mathrm{tuple}(\widehat{G})$ equals g_1^* since

- a smallest generator in $\mathrm{tuple}(\widehat{G})$ that is greater than g_1^* would cause an inconsistency with $(g + g_1^*, +)$, and
- a smallest generator in $\mathrm{tuple}(\widehat{G})$ that is smaller than g_1^* would have a lower preference (compare with Condition 2 above).

Assume inductively that the $i - 1$ smallest generators in $\mathrm{tuple}(\widehat{G})$ are $g_1^*, \ldots, g_{i-1}^*$. Since $g_i^* \notin \left\langle \{g_1^*, \ldots, g_{i-1}^*\} \right\rangle$, we may apply a reasoning that is similar to the above reasoning concerning g_1^* and conclude that the i'th smallest generator in $\mathrm{tuple}(\widehat{G})$ equals g_i^*. The punchline of this discussion is that the sequence $\mathrm{tuple}(\widehat{G})$ starts with $g_1^*, \ldots, g_h^*$ with h given by (12). Let $G' = G \setminus G^*$ be the set of redundant generators in G and note that

$$g - \sum_{i=1}^{h} g_i^* = \begin{cases} g_{\ell^*}^* & \text{if } G^* = G \\ \sum_{g' \in G'} g' & \text{if } G^* \subset G \end{cases} .$$

Let $\widehat{G}' = \widehat{G} \setminus \{g_1^*, \ldots, g_h^*\}$. We proceed by case analysis:

Case 1: $G^* = G$.

Since $\widehat{G}$ is consistent with $(g, +)$, we have $\sum_{g' \in \widehat{G}'} g' = g^*_{\ell^*}$. Since $g^*_{\ell^*} \notin \left\langle \{g^*_1, \ldots, g^*_{\ell^*-1}\} \right\rangle$, the set $\widehat{G}'$ must contain an element that cannot be generated by $g^*_1, \ldots, g^*_{\ell^*-1}$. Given the preferences of the student (compare with Condition 2), she will choose $\widehat{G}' = \{g^*_{\ell^*}\}$. It follows that $\widehat{G} = G$.

Case 2: $G^* \subset G$.

Here, we have $\sum_{g' \in \widehat{G}'} g' = \sum_{g' \in G'} g'$. Given the preferences of the student (compare with Condition 3), she will choose $\widehat{G}$ such that $\widehat{G}^* = G^*$ and $\widehat{G}'$ consists of elements from $\langle G^* \rangle$ that sum up to $\sum_{g' \in G'} g'$ (with $\widehat{G}' = \left\{ \sum_{g' \in G'} g' \right\}$ among the possible choices). Clearly, $\left\langle \widehat{G} \right\rangle_+ = \langle G \rangle_+$.

Thus, in both cases, the student comes up with the right hypothesis. ■

A.3 The Lower Bounds in Theorem 18

The lower bounds in Theorem 18 are an immediate consequence of the following result:

Lemma 38 *The following lower bounds are valid:*

$$\begin{aligned} \mathrm{PBTD}^+(\text{NE-CF-LINSET}'_k) &\geq k+1 \ . && (13) \\ \mathrm{PBTD}(\text{NE-CF-LINSET}'_k) &\geq k-1 \ . && (14) \\ \mathrm{PBTD}(\text{NE-CF-LINSET}_k) &\geq \frac{k-1}{2} \ . && (15) \\ \mathrm{PBTD}(\text{CF-LINSET}_k) &\geq k-1 \ . && (16) \end{aligned}$$

This lemma can be seen as an extension and a strengthening of a similar result in (Gao et al., 2015) where the following lower bounds were shown:

$$\begin{aligned} \mathrm{RTD}^+(\text{NE-LINSET}'_k) &\geq k+1 \ . \\ \mathrm{RTD}(\text{NE-LINSET}'_k) &\geq k-1 \ . \\ \mathrm{RTD}(\text{CF-LINSET}_k) &\geq k-1 \ . \end{aligned}$$

The proof of Lemma 38 builds on some ideas that are found in (Gao et al., 2015) already, but it requires some elaboration to obtain the stronger results.

We now briefly explain why the lower bounds in Theorem 18 directly follow from Lemma 38. Note that the lower bound $k-1$ in (8) is immediate from (14) and a monotonicity argument. This is because $\text{NE-LINSET}'_k \supseteq \text{NE-CF-LINSET}'_k$ as well as $\text{LINSET}'_k \supseteq \text{CF-LINSET}'_k \supseteq \text{NE-CF-LINSET}'_k$. Note furthermore that $\mathrm{PBTD}^+(\text{CF-LINSET}'_k) \geq k+1$ because of (13) and a monotonicity argument. Then the Shift Lemma implies that $\mathrm{PBTD}^+(\text{CF-LINSET}_k) \geq k$. All remaining lower bounds in Theorem 18 are obtained from these observations by virtue of monotonicity.

The proof of Theorem 18 can therefore be accomplished by proving Lemma 38. It turns out that the proof of this lemma is quite involved. We will present in Section A.3.1 some theoretical prerequisites. Sections A.3.2 and A.3.3 are devoted to the actual proof of the lemma.

A.3.1 SOME BASIC CONCEPTS IN THE THEORY OF NUMERICAL SEMIGROUPS

Recall from Section 6 that $\langle G\rangle = \left\{\sum_{g\in G} a(g)g : a(g) \in \mathbb{N}_0\right\}$. The elements of G are called *generators* of $\langle G\rangle$. A set $P \subset \mathbb{N}$ is said to be *independent* if none of the elements in P can be written as a linear combination (with coefficients from $\mathbb{N}_0$) of the remaining elements (so that $\langle P'\rangle$ is a proper subset of $\langle P\rangle$ for every proper subset P' of P). It is well known (Rosales and García-Sánchez, 2009) that independence makes generating systems unique, i.e., if P, P' are independent, then $\langle P\rangle = \langle P'\rangle$ implies that $P = P'$. Moreover, for every independent set P, the following implication is valid:

$$(S \subseteq \langle P\rangle \wedge P \not\subseteq S) \Rightarrow (\langle S\rangle \subset \langle P\rangle) \ . \tag{17}$$

Let $P = \{a_1, \ldots, a_k\}$ be independent with $a_1 = \min P$. It is well known[7] and easy to see that the residues of $a_1, a_2, \ldots, a_k$ modulo a_1 must be pairwise distinct (because, otherwise, we would obtain a dependence). If a_1 is a prime and $|P| \geq 2$, then the independence of P implies that $\gcd(P) = 1$. Thus the following holds:

Lemma 39 *If $P \subset \mathbb{N}$ is an independent set of cardinality at least 2 and $\min P$ is a prime, then* $\gcd(P) = 1$.

In the remainder of the paper, the symbols P and P' are reserved for denoting independent sets of generators.

It is well known that $\langle G\rangle$ is co-finite iff $\gcd(G) = 1$ (Rosales and García-Sánchez, 2009). Let P be a finite (independent) subset of $\mathbb{N}$ such that $\gcd(P) = 1$. The largest number in $\mathbb{N} \setminus \langle P\rangle$ is called the *Frobenius number of* P and is denoted as $F(P)$. It is well known (Rosales and García-Sánchez, 2009) that

$$F(\{p, q\}) = pq - p - q \tag{18}$$

provided that $p, q \geq 2$ satisfy $\gcd(p, q) = 1$.

A.3.2 PROOF OF (13)

The shift-extension of NE-CF-LINSET$_k$ is (by way of definition) the following class:

$$\text{NE-CF-LINSET}'_k = \{c + \langle P\rangle_+ : \ (c \in \mathbb{N}_0) \wedge (P \subset \mathbb{N}) \wedge (|P| \leq k) \wedge (\gcd(P) = 1)\} \ . \tag{19}$$

It is easy to see that this can be written alternatively in the form

$$\text{NE-CF-LINSET}'_k = \left\{ N + \langle P\rangle : \ N \in \mathbb{N}_0 \wedge P \subset \mathbb{N} \wedge |P| \leq k \wedge \gcd(P) = 1 \wedge \sum_{p\in P} p \leq N \right\} \tag{20}$$

where N in (20) corresponds to $c + \sum_{p\in P} p$ in (19).

For technical reasons, we define the following subfamilies of NE-CF-LINSET$'_k$. For each $N \geq 0$, let

$$\text{NE-CF-LINSET}'_k[N] = \{N + L : L \in \text{LINSET}_k[N]\}$$

7. E.g., see (Rosales and García-Sánchez, 2009)

where

$$\mathrm{LINSET}_k[N] = \left\{ \langle P \rangle \in \mathrm{LINSET}_k : (\gcd(P) = 1) \wedge \left(\sum_{p \in P} p \le N \right) \right\} \ .$$

In other words, NE-CF-LINSET$'_k[N]$ is the subclass consisting of all concepts in NE-CF-LINSET$'_k$ (written in the form (20)) whose constant is N.

A central notion for proving (13) is the following one:

Definition 40 *Let $k, N \ge 2$ be integers. We say that a set $L \in$ NE-CF-LINSET$'$ is (k,N)-special if it is of the form $L = N + \langle P \rangle$ such that the following holds:*

1. *P is an independent set of cardinality k and $\min P$ is a prime (so that $\gcd(P) = 1$ according to Lemma 39, which furthermore implies that $\langle P \rangle$ is co-finite).*
2. *Let $q(P)$ denote the smallest prime that is greater than $F(P)$ and greater than $\max P$. For $a = \min P$ and $r = 0, \ldots, a-1$, let*

$$t_r(P) = \min\{s \in \langle P \rangle : s \equiv r \pmod a\} \text{ and } t_{max}(P) = \max_{0 \le r \le a-1} t_r(P) \ .$$

Then

$$N \ge k(a + t_{max}(P)) \text{ and } N \ge q(P) + \sum_{p \in P \setminus \{a\}} p \ . \tag{21}$$

We need at least k positive examples in order to distinguish a (k,N)-special set from all its proper subsets in NE-CF-LINSET$'_k[N]$, as the following result shows:

Lemma 41 *For all $k \ge 2$, the following holds. If $L \in$ NE-CF-LINSET$'$ is (k,N)-special, then $L \in$ NE-CF-LINSET$'[N]$ and $I'(L, \text{NE-CF-LINSET}_k[N]) \ge k$.*

Proof Suppose that $L = N + \langle P \rangle$ is of the form as described in Definition 40. Let $P = \{a, a_2 \ldots, a_k\}$ with $a = \min P$. For the sake of simplicity, we will write t_r instead of $t_r(P)$ and t_{max} instead of $t_{max}(P)$. The independence of P implies that $t_{a_i \bmod a} = a_i$ for $i = 2, \ldots, k$. It follows that $t_{max} \ge \max P$. Since, by assumption, $N \ge k \cdot t_{max}$, it becomes obvious that $L \in$ NE-CF-LINSET$'[N]$.

Assume by way of contradiction that the following holds:

(A) There is a weak spanning set S of size $k-1$ for L w.r.t. NE-CF-LINSET$'_k[N]$.

Since N is contained in any concept from NE-CF-LINSET$'_k[N]$, we may assume that $N \notin S$ so that S is of the form $S = \{N + x_1, \ldots, N + x_{k-1}\}$ for integers $x_i \ge 1$. For $i = 1, \ldots, k-1$, let $r_i = x_i \bmod a \in \{0, 1, \ldots, a-1\}$. It follows that each x_i is of the form $x_i = q_i a + t_{r_i}$ for some integer $q_i \ge 0$. Let $X = \{x_1, \ldots, x_{k-1}\}$. We proceed by case analysis:

Case 1: $X \subseteq \{a_2, \ldots, a_k\}$ (so that, in view of $|X| = k-1$, we even have $X = \{a_2, \ldots, a_k\}$).
Let $L' = N + \langle X \rangle$. Then $S \subseteq L'$. Note that $X \subseteq P$ but $P \not\subseteq X$. We may conclude from (17) that $\langle X \rangle \subset \langle P \rangle$ and, therefore, $L' \subset L$. Thus L' is a proper subset of L which contains S. Note that (21) implies that $N \ge \sum_{i=2}^{k} a_i = \sum_{i=1}^{k-1} x_i$. If $\gcd(X) = 1$, then $L' \in$ NE-CF-LINSET$[N]$ and we have an immediate contradiction to the above assumption (A).

Otherwise, if $\gcd(X) \geq 2$, then we define $L'' = N + \langle X \cup \{q(P)\} \rangle$. Note that $S \subseteq L' \subseteq L''$. Since $q(P) > F(P)$, we have $X \cup \{q(P)\} \subseteq \langle P \rangle$ and, since $q(P) > \max P$, we have $P \not\subseteq X \cup \{q(P)\}$. We may conclude from (17) that $\langle X \cup \{q(P)\} \rangle \subset \langle P \rangle$ and, therefore, $L'' \subset L$. Thus, L'' is a proper subset of L which contains S. Because $X = \{a_2, \ldots, a_k\}$ and $q(P)$ is a prime that is greater than $\max P$, it follows that $\gcd(X \cup \{q(P)\}) = 1$. In combination with (21), it easily follows now that $L'' \in \text{NE-CF-LINSET}[N]$. Putting everything together, we arrive at a contradiction to the assumption (A).

Case 2: $X \not\subseteq \{a_2, \ldots, a_k\}$.

If $r_i = 0$ for $i = 1, \ldots, k-1$, then each x_i is a multiple of a. In this case, $N + \langle a, q(P) \rangle$ is a proper subset of $L = N + \langle P \rangle$ that is consistent with S, which yields a contradiction. We may therefore assume that there exists $i' \in \{1, \ldots, k-1\}$ such that $r_{i'} \neq 0$. From the case assumption, $X \not\subseteq \{a_2, \ldots, a_k\}$, it follows that there must exist an index $i'' \in \{1, \ldots, k-1\}$ such that $q_{i''} \geq 1$ or $t_{r_{i''}} \notin \{a_2, \ldots, a_k\}$. For $i = 1, \ldots, k-1$, let $q'_i = \min\{q_i, 1\}$ and $x'_i = q'_i a + t_{r_i}$. Note that $q'_{i''} = 1$ iff $q_{i''} \geq 1$. Define $L'' = N + \langle X' \rangle$ for $X' = \{a, x'_1, \ldots, x'_{k-1}\}$ and observe the following. First, the set L'' clearly contains S. Second, the choice of $x'_1, \ldots, x'_{k-1}$ implies that $X' \subseteq \langle P \rangle$. Third, it easily follows from $q'_{i''} = 1$ or $t_{r_{i''}} \notin \{a_2, \ldots, a_k\}$ that $P \not\subseteq \{a, x'_1, \ldots, x'_{k-1}\}$. We may conclude from (17) that $\langle X' \rangle \subset \langle P \rangle$ and, therefore, $L'' \subset L$. Thus, L'' is a proper subset of L which contains S. Since $r_{i'} \neq 0$ and a is a prime, it follows that $\gcd(a, x'_{i'}) = 1$ and, therefore, $\gcd(X') = 1$. In combination with (21), it easily follows now that $L'' \in \text{NE-CF-LINSET}[N]$. Putting everything together, we obtain again a contradiction to the assumption (A).

∎

For the sake of brevity, let $\mathcal{L} = \text{NE-CF-LINSET}'$. Assume by way of contradiction that there exists a positive mapping T of order k that is admissible for $\mathcal{L}_k$. We will pursue the following strategy:

1. We define a set $L \in \mathcal{L}_k$ of the form $L = N + p + \langle 1 \rangle$.
2. We define a second set $L' = N + \langle G \rangle \in \mathcal{L}$ that is (k, N)-special and consistent with $T^+(L)$. Moreover, $L' \setminus L = \{N\}$.

If this can be achieved, then the proof will be accomplished as follows:

- According to Lemma 41, $T^+(L')$ must contain at least k examples (all of which are different from N) for distinguishing L' from all its proper subsets in $\mathcal{L}_k[N]$.
- Since L' is consistent with $T^+(L)$, the set $T^+(L')$ must contain an example which distinguishes L' from L. But the only example which fits this purpose is $(N, +)$.
- The discussion shows that $T^+(L')$ must contain k examples in order to distinguish L' from all its proper subsets in $\mathcal{L}_k$ plus one additional example, N, needed to distinguish L' from L.
- We obtain a contradiction to our initial assumption that T^+ is of order k.

We still have to describe how our proof strategy can actually be implemented. We start with the definition of L. Pick the smallest prime $p \geq k+1$. Then $\{p, p+1, \ldots, p+k\}$ is independent. Let $M = F(\{p, p+1\}) \stackrel{(18)}{=} p(p+1) - p - (p+1)$. An easy calculation shows that $k \geq 2$ and $p \geq k+1$ imply that $M \geq p+k$. Let $I = \{p, p+1, \ldots, M\}$. Choose N large enough so that all concepts of the form

$$N + \langle P \rangle \text{ where } |P| = k,\ p = \min P \text{ and } P \subseteq I$$

are (k, N)-special. With these choices of p and N, let $L = N+p+\langle 1 \rangle$. Note that $N+p, N+p+1 \in T^+(L)$ because, otherwise, one of the concepts $N+p+1+\langle 1 \rangle, N+p+\langle 2, 3 \rangle \subset L$ would be consistent with $T^+(L)$ whereas $T^+(L)$ must distinguish L from all its proper subsets in $\mathcal{L}_k$. Setting $A = \{x : N + x \in T^+(L)\}$, it follows that $|A| = |T^+(L)| \leq k$ and $p, p+1 \in A$. The set A is not necessarily independent but it contains an independent subset B such that $p, p+1 \in B$ and $\langle A \rangle = \langle B \rangle$. Since $M = F(\{p, p+1\})$, it follows that any integer greater than M is contained in $\langle p, p+1 \rangle$. Since B is an independent extension of $\{p, p+1\}$, it cannot contain any integer greater than M. It follows that $B \subseteq I$. Clearly, $|B| \leq k$ and $\gcd(B) = 1$. We would like to transform B into another generating system $G \subseteq I$ such that

$$\langle B \rangle \subseteq \langle G \rangle, \gcd(G) = 1 \text{ and } |G| = k \ .$$

If $|B| = k$, we can simply set $G = B$. If $|B| < k$, then we make use of the elements in the independent set $\{p, p+1, \ldots, p+k\} \subseteq I$ and add them, one after the other, to B (thereby removing other elements from B whenever their removal leaves $\langle B \rangle$ invariant) until the resulting set G contains k elements. We now define the set L' by setting $L' = N + \langle G \rangle$. Since $G \subseteq I = \{p, p+1, \ldots, M\}$, and $p, p+1 \in G$, it follows that $p = \min G$, $\gcd(G) = 1$ and $\min(L' \setminus \{N\})$ is $N + p$. Thus, $L' \setminus L = \{N\}$, as desired. Moreover, since N had been chosen large enough, the set L' is (k, N)-special. Thus L and L' have all properties that are required by our proof strategy and the proof of (13) is complete.

A.3.3 PROOF OF (14), (15), AND (16)

We make use of some well known (and trivial) lower bounds on TD_{min}:

Example 5 *For every $k \in \mathbb{N}$, let $[k] = \{1, 2, \ldots, k\}$, let $2^{[k]}$ denote the powerset of $[k]$ and, for all $\ell = 0, 1, \ldots, k$, let*

$$\binom{[k]}{\ell} = \{S \subseteq [k] :\ |S| = \ell\}$$

denote the class of those subsets of $[k]$ that have exactly ℓ elements. It is trivial to verify that

$$\mathrm{TD}_{min}\left(2^{[k]}\right) = k \ \text{ and } \ \mathrm{TD}_{min}\left(\binom{[k]}{\ell}\right) = \min\{\ell, k - \ell\} \ .$$

In view of $\mathrm{PBTD}^+(\mathrm{LINSET}_k) = k$, the next results show that negative examples are of limited help only as far as preference-based teaching of concepts from LINSET_k is concerned:

Lemma 42 *For every $k \geq 1$ and for all $\ell = 0, \ldots, k-1$, let*

$$\begin{aligned}\mathcal{L}_k &= \{\langle k, p_1, \ldots, p_{k-1}\rangle : \ p_i \in \{k+i, 2k+i\}\} \ , \\ \mathcal{L}_{k,\ell} &= \{\{\langle k, p_1, \ldots, p_{k-1}\rangle \in \mathcal{L}_k : \ |\{i : p_i = k+i\}| = \ell\} \ .\end{aligned}$$

With this notation, the following holds:

$$\mathrm{TD}_{min}(\mathcal{L}_k) \geq k-1 \ \textit{and} \ \mathrm{TD}_{min}(\mathcal{L}_{k,\ell}) \geq \min\{\ell, k-1-\ell\} \ .$$

Proof For $k = 1$, the assertion in the lemma is vacuous. Suppose therefore that $k \geq 2$. An inspection of the generators $k, p_1, \ldots, p_{k-1}$ with $p_i \in \{k+i, 2k+i\}$ shows that

$$\begin{aligned}\mathcal{L}_k &= \{L_{k,S} : \ S \subseteq \{k+1, k+2, \ldots, 2k-1\}\} \\ \mathcal{L}_{k,\ell} &= \{L_{k,S} : \ (S \subseteq \{k+1, k+2, \ldots, 2k-1\}) \wedge (|S| = \ell)\}\end{aligned}$$

where

$$L_{k,S} = \{0, k\} \cup \{2k, 2k+1, \ldots\} \cup S \ .$$

Note that the examples in $\{0, 1, \ldots, k\} \cup \{2k, 2k+1, \ldots, \}$ are redundant because they do not distinguish between distinct concepts from $\mathcal{L}_k$. The only useful examples are therefore contained in the interval $\{k+1, k+2, \ldots, 2k-1\}$. From this discussion, it follows that teaching the concepts of $\mathcal{L}_k$ (resp. of $\mathcal{L}_{k,\ell}$) is not essentially different from teaching the concepts of $2^{[k-1]}$ $\left(\text{resp. of } \binom{[k-1]}{\ell}\right)$. This completes the proof of the lemma because we know from Example 5 that $\mathrm{TD}_{min}(2^{[k-1]}) = k-1$ and $\mathrm{TD}_{min}\left(\binom{[k-1]}{\ell}\right) = \min\{\ell, k-1-\ell\}$. ∎

We claim now that the inequalities (14), (15) and (16) are valid, i.e., we claim that the following holds:

1. $\mathrm{PBTD}(\text{CF-LINSET}_k) \geq k-1$.
2. $\mathrm{PBTD}(\text{NE-CF-LINSET}_k) \geq \lfloor (k-1)/2 \rfloor$.
3. $\mathrm{PBTD}(\text{NE-CF-LINSET}'_k) \geq k-1$.

Proof For $k = 1$, the inequalities are obviously valid. Suppose therefore that $k \geq 2$.

1. Since $\gcd(k, k+1) = \gcd(k, 2k+1) = 1$, it follows that $\mathcal{L}_k$ is a finite subclass of CF-LINSET$_k$. Thus $\mathrm{PBTD}(\text{CF-LINSET}_k) \geq \mathrm{PBTD}(\mathcal{L}_k) \geq \mathrm{TD}_{min}(\mathcal{L}_k) \geq k-1$.
2. Define $\mathcal{L}_k[N] = \{N + L : \ L \in \mathcal{L}_k\}$ and $\mathcal{L}_{k,\ell}[N] = \{N + L : \ L \in \mathcal{L}_{k,\ell}\}$. Clearly $\mathrm{TD}_{min}(\mathcal{L}_k[N]) = \mathrm{TD}_{min}(\mathcal{L}_k)$ and $\mathrm{TD}_{min}(\mathcal{L}_{k,\ell}[N]) = \mathrm{TD}_{min}(\mathcal{L}_{k,\ell})$ holds for every $N \geq 0$. It follows that the lower bounds in Lemma 42 are also valid for the classes $\mathcal{L}_k[N]$ and $\mathcal{L}_{k,\ell}[N]$ in place of $\mathcal{L}_k$ and $\mathcal{L}_{k,\ell}$, respectively. Let

$$N(k) = k^2 + (k-1-\lfloor (k-1)/2 \rfloor)k + \sum_{i=1}^{k-1} i = k^2 + (k-1-\lfloor (k-1)/2 \rfloor)k + \frac{1}{2}(k-1)k \ . \quad (22)$$

It suffices to show that $N(k) + \mathcal{L}_{k,\lfloor (k-1)/2 \rfloor}$ is a finite subclass of NE-CF-LINSET$_k$. To this end, first note that

$$\langle k, p_1, \ldots, p_{k-1} \rangle_+ = k + \sum_{i=1}^{k-1} p_i + \langle k, p_1, \ldots, p_{k-1} \rangle \ .$$

Call p_i "light" if $p_i = k + i$ and call it "heavy" if $p_i = 2k + i$. Note that a concept L from $N(k) + \mathcal{L}_{k,\ell}$ is of the general form

$$L = N(k) + \langle k, p_1, \ldots, p_{k-1} \rangle \tag{23}$$

with exactly ℓ light parameters among $p_1, \ldots, p_{k-1}$. A straightforward calculation shows that, for $\ell = \lfloor (k-1)/2 \rfloor$, the sum $k + \sum_{i=1}^{k-1} p_i$ equals the number $N(k)$ as defined in (22). Thus, the concept L from (23) with exactly $\lfloor (k-1)/2 \rfloor$ light parameters among $\{p_1, \ldots, p_{k-1}\}$ can be rewritten as follows:

$$L = N(k) + \langle k, p_1, \ldots, p_{k-1} \rangle = \langle k, p_1, \ldots, p_{k-1} \rangle_+ \ .$$

This shows that $L \in$ NE-CF-LINSET$_k$. As L is a concept from $N(k) + \mathcal{L}_{k,\lfloor (k-1)/2 \rfloor}$ in general form, we may conclude that $N(k) + \mathcal{L}_{k,\lfloor (k-1)/2 \rfloor}$ is a finite subclass of NE-CF-LINSET$_k$, as desired.

3. The proof of the third inequality is similar to the above proof of the second one. It suffices to show that, for every $k \geq 2$, there exists $N \in \mathbb{N}$ such that $N + \mathcal{L}_k$ is a subclass of NE-CF-LINSET$'_k$. To this end, we set $N = 3k^2$. A concept L from $3k^2 + \mathcal{L}_k$ is of the general form

$$L = 3k^2 + \langle k, p_1, \ldots, p_{k-1} \rangle$$

with $p_i \in \{k + i, 2k + i\}$ (but without control over the number of light parameters). It is easy to see that the constant $3k^2$ is large enough so that L can be rewritten as

$$L = 3k^2 - \left(k + \sum_{i=1}^{k-1} p_i \right) + \langle k, p_1, \ldots, p_{k-1} \rangle_+$$

where $3k^2 - \left(k + \sum_{i=1}^{k-1} p_i \right) \geq 0$. This shows that $L \in$ NE-CF-LINSET$'_k$. As L is a concept from $3k^2 + \mathcal{L}_k$ in general form, we may conclude that $3k^2 + \mathcal{L}_k$ is a finite subclass of NE-CF-LINSET$'_k$, as desired.

■