

Washington University in St. Louis

Washington University Open Scholarship

All Computer Science and Engineering
Research

Computer Science and Engineering

Report Number: WUCS-93-2

1993-01-01

Teaching a Smarter Learner

Sally A. Goldman and H. David Mathias

We introduce a formal model of teaching in which the teacher is tailored to a particular learner, yet the teaching protocol is designed so that no collusion is possible. Not surprisingly, such a model remedies the non-intuitive aspects of other models in which the teacher must successfully teach any consistent learner. We prove that any class that can be exactly identified by a deterministic polynomial-time algorithm with access to a very rich set of example-based queries is teachable by a computationally unbounded teacher and a polynomial-time learner. In addition, we present other general results relating this model of teaching... **Read complete abstract on page 2.**

Follow this and additional works at: https://openscholarship.wustl.edu/cse_research



Part of the [Computer Engineering Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Goldman, Sally A. and Mathias, H. David, "Teaching a Smarter Learner" Report Number: WUCS-93-2 (1993). *All Computer Science and Engineering Research*.
https://openscholarship.wustl.edu/cse_research/307

Department of Computer Science & Engineering - Washington University in St. Louis
Campus Box 1045 - St. Louis, MO - 63130 - ph: (314) 935-6160.

Teaching a Smarter Learner

Sally A. Goldman and H. David Mathias

Complete Abstract:

We introduce a formal model of teaching in which the teacher is tailored to a particular learner, yet the teaching protocol is designed so that no collusion is possible. Not surprisingly, such a model remedies the non-intuitive aspects of other models in which the teacher must successfully teach any consistent learner. We prove that any class that can be exactly identified by a deterministic polynomial-time algorithm with access to a very rich set of example-based queries is teachable by a computationally unbounded teacher and a polynomial-time learner. In addition, we present other general results relating this model of teaching to various previous results. We also consider the problem of designing teacher/learner pairs in which both the teacher and learner are polynomial-time algorithms and describe teacher/learner pairs for the classes of 1-decision lists and Horn sentences.

Teaching a Smarter Learner

Sally A. Goldman and H. David Mathias

WUCS-93-02

September 1993

**Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
St. Louis MO 63130-4899**

This research was supported in part by NSF grant CCR-9110108 and by a GE Foundation Junior Faculty Grant. An earlier version appears in the Proceedings of the Sixth Annual Workshop on Computational Learning Theory, pages 67-86, July 1993.

Teaching a Smarter Learner *

Sally A. Goldman
Department of Computer Science
Washington University
St. Louis, MO 63130
sg@cs.wustl.edu

H. David Mathias
Department of Computer Science
Washington University
St. Louis, MO 63130
dmath@cs.wustl.edu

September 1993

WUCS-93-02

Abstract

We introduce a formal model of teaching in which the teacher is tailored to a particular learner, yet the teaching protocol is designed so that no collusion is possible. Not surprisingly, such a model remedies the non-intuitive aspects of other models in which the teacher must successfully teach any consistent learner. We prove that any class that can be exactly identified by a deterministic polynomial-time algorithm with access to a very rich set of example-based queries is teachable by a computationally unbounded teacher and a polynomial-time learner. In addition, we present other general results relating this model of teaching to various previous results. We also consider the problem of designing teacher/learner pairs in which both the teacher and learner are polynomial-time algorithms and describe teacher/learner pairs for the classes of 1-decision lists and Horn sentences.

1 Introduction

Recently, there has been interest in developing formal models of *teaching* [5, 10, 11, 16, 27] through which we can develop a better understanding of how a teacher can most effectively aid a learner in accomplishing a learning task. A weakness of the formal models of teaching that have been introduced in the learning theory community is that

*This research was supported in part by NSF grant CCR-9110108 and by a GE Foundation Junior Faculty Grant. An earlier version appears in the *Proceedings of the Sixth Annual Workshop on Computational Learning Theory*, pages 67–86, July 1993.

they place stringent restrictions on the learner to ensure that the teacher is not just providing the learner with an encoding of the target. In particular, previous models require the teacher to present a set of examples for which only the target function is consistent. Thus, teaching under these models is made unnecessarily difficult since the problem reduces to teaching an obstinate learner that tries as hard as possible not to learn while always outputting a hypothesis consistent with all previous examples. In fact, if a teacher is required to teach any consistent learner ¹, there are many examples for which an exponential length teaching set is required to teach even those classes for which efficient learning algorithms are known. For many of the classes that can be taught efficiently, it is necessary in previous models to allow the teacher and learner to share a small amount of “trusted information”, such as the size of the target function, since there is no other way to eliminate concepts from the given class that are more “complicated” than the target. Jackson and Tompkins [16] are able to show that anything learnable with membership and equivalence queries is teachable with trusted information. However, their method for preventing collusion still requires that the teacher successfully teaches any consistent learner.

As a concrete example of a consequence of requiring that the teaching set is sufficient for any consistent learner, consider the class \mathcal{C} consisting of all singletons plus the empty set. While, this class is quite simple and thus should be easy to teach, Goldman and Kearns show that $|\mathcal{C}| - 1$ examples are required to teach it in their model. Furthermore, this hardness result has been embedded into several other hardness results such as that for teaching full decision lists [16] and teaching linearly separable Boolean functions [5]. Thus these hardness results appear to be due to a defect in the model rather than an intrinsic difficulty in teaching these classes. Similar problems to the ones discussed above emerge when comparing these teaching models to the self-directed learning model ². In particular, for many classes the self-directed learning complexity is asymptotically less than the teaching complexity. Again, because the teacher must

¹A learner is consistent if its hypotheses are consistent with all previously seen examples.

²A self-directed learner [12, 13] is a learner that selects the presentation order for the instances. In this model, the learning complexity is measured according to the number of *incorrect* predictions made.

successfully teach all consistent learners, a “smart” self-directed learner can perform better on its own than with the teacher’s guidance. Yet, if the teacher and learner could cooperate (which should be the case if the teacher is working with the single learner), intuitively this phenomenon should not occur.

There are many applications in which it is desirable to have teacher/learner pairs. For example, such results provide lower bounds on the number of examples required by a learning algorithm and such bounds could then be used to direct efforts in collecting data. Another application area is the problem of training a neural network. Currently, most training algorithms work by adjusting the weights in the network based on randomly selected labeled examples. By having a teacher carefully select the set of labeled examples, the training time might be drastically reduced. In addition, there are potential applications of the research on formal models of teaching to improving automated manufacturing environments. For example, consider the problem of training sensor-referenced intelligent robot controllers. In this problem the goal is to map a general purpose robot to specific application domains. The task of directly programming the robot to perform the given task is extremely difficult for two reasons: (1) the operator thinks about the robot’s motions in terms of Cartesian space whereas the robot’s effectors are described in terms of rotations or movements of the joints, and (2) there are inherent calibration errors in the movement of the robot’s effectors. In such an application it would be very desirable to design a teacher/learner pair in which a learning algorithm is selected for the robot, and then a teacher (i.e. the operator) guides the robot through a “representative” set of actions to enhance the robot’s speed of learning.

The key contribution of this work is the introduction of a formal teaching model that allows the teacher and learner to cooperate, yet the teacher cannot simply encode the target function. We start with a teacher/learner pair as in the model introduced by Jackson and Tompkins [16]. However, unlike in their work, we only require that if the teacher is replaced by an adversarial “substitute” that embeds the teaching set of the true teacher within his teaching set, then the learner will still output a hypothesis that is logically equivalent to the target function. While the adversarial substitute

has the strength to prevent collusion, it does not require that the teacher successfully teaches any consistent learner. We show that any class for which there is an efficient deterministic learning algorithm (even when provided with sophisticated queries) can be taught (without trusted information) under our model. Also the number of examples required by the teacher is at most the maximum number of mistakes made by any self-directed learning algorithm. Furthermore, using our model there is an interesting relationship between teaching and data compression. Applying the results of Floyd [9] we obtain that for any *maximum class* \mathcal{C} there is a teacher/learner pair for which at most $\text{VCD}(\mathcal{C})$ examples are presented. Likewise, from the results of Helmbold, Sloan and Warmuth [15], it follows that for any intersection-closed class \mathcal{C} , the nested difference of p functions from \mathcal{C} can be taught in our model with at most $p \cdot \text{VCD}(\mathcal{C})$ examples. It is clear that in this paper we only scratch the surface of such results that follow from previous work of others.

In addition to the more general results, we apply our model to the representation classes of 1-decision lists and Horn sentences to demonstrate the design of teacher/learner pairs in which both the teacher and learner require only polynomial computation time. For both classes, the sample complexity is asymptotically less than that for the best known learning algorithm.

2 Previous Work

We now briefly review the theoretical work studying the complexity of teaching. Goldman, Rivest and Schapire [12] introduced the model of teacher-directed learning, a variant of the on-line learning model in which a helpful teacher selects the instances, and applied it to the problem of learning binary relations and total orders. Building upon this framework, Goldman and Kearns [11] defined a formal model of teaching in which they measured the complexity of teaching by the minimum number of examples that must be presented to *any* consistent learner so that the learner outputs a hypothesis logically equivalent to the target function. Independently, Shinohara and Miyano [27] introduced an equivalent notion of teachability in which a class is *teachable by examples* if there exists a polynomial size sample under which all consistent learners

will exactly identify the target.

In other related work, Anthony, Brightwell, Cohen, and Shawe-Taylor [5] compute bounds on the size of the smallest sample with which only the target function is consistent for subclasses of linearly separable Boolean functions. Natarajan [21] defines a dimension measure for classes of Boolean functions that measures the complexity of a class \mathcal{C} by the length of the shortest example sequence for which the target function is the unique, most specific function from \mathcal{C} consistent with the sample. Salzberg, Delcher, Heath and Kasif [25] have also considered a model of learning with a helpful teacher. Their model requires the teacher to present the shortest example sequence so that any learner using a particular algorithm (namely, the nearest-neighbor algorithm) learns the target. However, their work does not address the issue of preventing the teacher and learner from colluding. Romanik and Smith [23, 24] propose a testing problem that involves specifying, for a given target function, a set of test points that can be used to determine if a tested object is equivalent to the target. However, their primary concern is to determine for which classes there exists a finite set of instances such that any representation in the class that is consistent on the test set is “close” to the target function in a probabilistic sense.

Within the inductive inference paradigm, Freivalds, Kinber and Wiehagen [10] and Lange and Wiehagen [17] have examined inference from “good examples”. Good examples are chosen by a helpful teacher to reduce the number of examples required. In both, encoding is avoided by requiring that the inference task is accomplished even when the learner is presented with any superset of the set of teacher-chosen examples. Neither of these results, however, offer careful proof that this method actually prevents collusion between the teacher and learner. Lange and Wiehagen [17] examine learning pattern languages and show that this can be achieved with good examples.

Our new teaching model is most closely related to the model introduced by Jackson and Tomkins [16]. In their model there are teacher/learner pairs in which the teacher chooses examples tailored to a particular learner. To avoid collusion between the teacher and learner, they consider the interaction between the teacher and learner as a modified prover-verifier session [14] in which the learner and teacher can collude

but no adversarial substitute teacher can cause the learner to output a hypothesis inconsistent with the sample. While it appears that the teacher’s knowledge of the learner in this model is powerful, they showed that under their model the teacher must still produce a teaching set that eliminates all but the target function from the representation class. They also introduced the notion of a small amount of *trusted information* that the teacher can provide the learner. This trusted information is used by the teacher to provide the learner with the size complexity of the target function or a stopping condition.

3 Our Model

We now formally define our model. The teacher’s goal is to teach the learner the target function ³ f chosen from some known *representation class* \mathcal{C} , which is a set of representations of functions mapping some domain \mathcal{X} into $\{0, 1\}$. In addition, $\mathcal{C} = \bigcup_{n \geq 1} \mathcal{C}_n$ is often parameterized by some natural dimension measure n . Let \mathcal{X}_n denote the set of instances to be classified for each problem of size n , and let $\mathcal{X} = \bigcup_{n \geq 1} \mathcal{X}_n$ denote the *instance space*. For $f \in \mathcal{C}_n$ and $x \in \mathcal{X}_n$, $f(x)$ denotes the classification of the function represented by f when evaluated on instance x . Each representation $f \in \mathcal{C}_n$ has a *size* denoted by $|f|$. Typically, this is the number of symbols needed to write the representation of f as a member of the representation class \mathcal{C}_n from which it is chosen. Finally, a *hypothesis* h is any polynomially-evaluatable function that, given any $x \in \mathcal{X}_n$, outputs a prediction for $f(x)$.

A *teaching set* for $f \in \mathcal{C}$ is an unordered set of labeled instances where each instance is selected from \mathcal{X} and labeled according to f . We define the teacher T to be an algorithm that when given a representation $f \in \mathcal{C}$ outputs a teaching set $T(f)$ for f . Similarly, we define the learner L to be an algorithm that takes as an argument any teaching set S and outputs a representation f' from \mathcal{C} . We use $L(S)$ to denote the representation output by L . (Observe that this definition can easily be extended to allow the learner to output a representation from some class $\mathcal{C}' \supseteq \mathcal{C}$.) If the learner

³Technically, the teacher is given a representation $f \in \mathcal{C}$. However, we shall equate f with the logical function it represents.




Figure 1: An overview of a teaching session. The adversary chooses the target function, f , giving it to the teacher. The teacher then generates $T(f)$. Given $T(f)$ the adversary generates $S_A \supseteq T(f)$ and gives this to the learner. The learner (possibly randomized) outputs a representation from \mathcal{C} thus defining a probability distribution $P_L(S_A)$ over \mathcal{C} .

is deterministic then $L(S)$ is well-defined, however, in the case that the learner uses a randomized algorithm, instead $L(S)$ induces a probability distribution over \mathcal{C} . We shall denote this distribution by $P_L(S)$.

We now describe our teaching protocol. The learner L and teacher T both have prior knowledge of the representation class \mathcal{C} from which the target function will be selected. Furthermore, they can cooperate to develop coordinated teaching and learning strategies that best enable the teacher to teach the learner some unknown function from the class. In addition to the teacher and learner, there is an adversary A who has unlimited computing power and complete knowledge of T and L . The teaching session, illustrated in Figure 1, proceeds as follows:

- The adversary selects a target function $f \in \mathcal{C}$ and gives f to T .
- The teacher computes $T(f)$ and gives it to A .
- Next the adversary (with knowledge of \mathcal{C}, f, T and L) adds properly labeled examples to $T(f)$ with the goal of causing the learner to fail. The teaching set obtained ($S_A \supseteq T(f)$) is then given to the learner.
- Finally, the learner outputs the representation given by $L(S_A)$.

The goal of the teacher is to teach the learner to perfectly predict whether any given instance is a positive or negative instance of the target function. Thus, the learner must achieve exact (logical) identification of the target. Of course, the teacher

would like to help the learner achieve this goal with the fewest number of examples possible. However, as discussed above, we must preclude unnatural “collusion” between the teacher and the learner (such as agreed-upon coding schemes to communicate the representation of the target via the instances selected without regard for the labels) which could trivialize our model. Informally, we define collusion as the passing of information, by the teacher to the learner, about the representation of the target rather than about the function represented.

We define a *valid T/L pair* for \mathcal{C} to consist of a teacher T and learner L such that: For any $f \in \mathcal{C}$ the teaching set $T(f)$ output has the property that if L is provided with any teaching set $S_A \supseteq T(f)$ where all added examples are properly labeled according to f , then $P_L(S_A)$ has the property that for all $f' \in \mathcal{C}$, if f' has non-zero weight in the distribution $P_L(S_A)$ then f' is logically equivalent to f . In other words, any representation output by L will be logically equivalent to f .

Given a valid T/L pair, we say that the teacher T is a polynomial-time teacher if, given any $f \in \mathcal{C}_n$, it outputs $T(f)$ in time polynomial in n and $|f|$. Likewise, the learner L is a polynomial-time learner if it runs in time polynomial in n , $|f|$ and $|S_A|$ for any $S_A \supseteq T(f)$. We say that a representation class \mathcal{C} is *T/L-teachable* if, for all $f \in \mathcal{C}_n$, there exists a valid T/L pair for which $|T(f)|$ is polynomial in $|f|$ and n . We say that \mathcal{C} is *polynomially T/L-teachable* if it is T/L-teachable by a pair for which T is a polynomial-time teacher and L is a polynomial-time learner. Finally, we say that \mathcal{C} is *semi-poly T/L-teachable* if it is T/L-teachable with a polynomial-time learner but a teacher that may be computationally unbounded.

In the next section we argue that the adversarial substitute is sufficient to prevent collusion. Again, intuitively one view of collusion is for the teacher to use the teaching set to pass information about the representation of the target versus the function it represents. We formalize this notation in the following manner ⁴. We define a *colluding T/L pair* for \mathcal{C} to consist of a teacher T and learner L such that: There exist logically equivalent f_0, f_1 in \mathcal{C} (possibly with $f_0 = f_1$) such that for all $S_0 \supseteq T(f_0)$ and all $S_1 \supseteq T(f_1)$, $P_L(S_0) \neq P_L(S_1)$. In other words, if the teacher is able to influence

⁴Clearly there are other ways in which one could formalize the notion of collusion. However, throughout the remainder of this paper, we shall use the given definition.

the distribution on \mathcal{C} returned by L when presented with a teaching set for logically equivalent representations then there is collusion between T and L .

4 General Results

In this section we explore the properties of our new teaching model. We begin by proving that the adversarial substitute is sufficient to prevent collusion.

Theorem 1 *There is no colluding T/L pair.*

Proof: We use a proof by contradiction. Suppose there exists a colluding T/L pair. Since the T/L pair is colluding, by definition there exist logically equivalent f_0, f_1 in \mathcal{C} (possibly with $f_0 = f_1$) such that for all $S_0 \supseteq T(f_0)$ and all $S_1 \supseteq T(f_1)$, $P_L(S_0) \neq P_L(S_1)$. However, let $S_0 = S_1 = T(f_0) \cup T(f_1)$. Then, since $S_0 = S_1$ it follows that $P_L(S_0) = P_L(S_1)$ giving the desired contradiction. \square

It immediately follows from this theorem, that any valid T/L pair must not collude. Notice that unlike the Jackson and Tompkins model in which the teacher uses “trusted bits” (in addition to labeled examples) to successfully teach many classes, our teacher must teach the target function entirely through a careful selection of labeled examples.

While technically not collusion as we have defined it, there is a type of cooperation between the teacher and learner, allowing the transmission of information about the target, that merits discussion. Let \mathcal{C} be a representation class over the domain $X = \{0, 1\}^n$ with no more than $2^{O(n)}$ representations and let there be a total order on \mathcal{C} . Thus, each representation f can be identified by a constant number of instances encoding the place of f in the total order. To teach such a class, the teacher simply places the instances $\{x_1, \dots, x_k\}$ identifying f , along with their labels according to f , in its teaching set. Additionally, the teacher must add to the teaching set a set of examples that uniquely identify the function represented by f (i.e. a teaching sequence as defined by Goldman and Kearns [11]). The learner chooses an ordered k -tuple of instances from the teaching set and identifies its associated $f \in \mathcal{C}$. If every instance in the teaching set is labeled according to f then f is the target function and is output

by the learner. Otherwise, the learner continues choosing ordered k -tuples until the target is identified in this manner.

We now demonstrate that the above method is not collusion according to the definition we have given. Let $f_1, f_2 \in \mathcal{C}$ such that for all $x \in X$, $f_1(x) = f_2(x)$. Without loss of generality, if the teacher attempts to teach f_1 then the adversarial substitute can simply add the k -tuple for f_2 , with their labels, to the teaching set. Then both f_1 and f_2 are consistent with the entire teaching set and either may be output by the learner. Thus, no collusion is occurring. Intuitively, we do not feel that collusion is occurring since the teaching set contains examples to *uniquely identify* the function represented by the target representation. However, the learner may have to solve a hard consistency problem to identify the target without using the above scheme. So, this T/L pair relies on the power of the teacher’s knowledge of the learner to solve a hard problem.

We now show that any representation class learnable in deterministic polynomial time from “example-based” queries (including equivalence, membership, subset, superset, disjointness, exhaustiveness, justifying assignments, partial equivalence) is teachable in our model by a computationally unbounded teacher and a polynomial-time learner. We define an *example-based query* to be any query of the form:

$$\forall (x_1, x_2, \dots, x_k) \in \mathcal{X}^k, \text{ does } \varphi_f(x_1, x_2, \dots, x_k) = 1?$$

where k is constant, and $\varphi_f(x_1, x_2, \dots, x_k)$ is any poly-time computable predicate with membership-query access to the target f . Observe that the predicate φ may use the x_i ’s to compute other instances on which to perform membership queries. The answer provided to the example-based query is either “yes” or a counterexample consisting of $(x_1, x_2, \dots, x_k) \in \mathcal{X}^k$ (with their labels) for which $\varphi_f(x_1, x_2, \dots, x_k) = 0$ and the examples (and their labels) for which membership queries were made to evaluate the predicate. In Figure 2 we give the definition for the predicate φ_f corresponding to queries in standard use. Observe that all reasonable, and some fairly bizarre, queries can be formulated in this manner.

Theorem 2 *Any representation class \mathcal{C} learnable in deterministic polynomial-time using example-based queries is semi-poly T/L-teachable.*

Type of Query	k	$\varphi_f(x_1, \dots, x_k)$
equivalence(h)	1	$h(x_1) = f(x_1)$
membership(x)	0	$f(x)$
subset(h)	1	$h(x_1) = 1 \Rightarrow f(x_1) = 1$
superset(h)	1	$f(x_1) = 1 \Rightarrow h(x_1) = 1$
disjointness(h)	1	$(f(x_1) = 1 \Rightarrow h(x_1) = 0) \wedge (h(x_1) = 1 \Rightarrow f(x_1) = 0)$
exhaustiveness(h)	1	$(f(x_1) = 0 \Rightarrow h(x_1) = 1) \wedge (h(x_1) = 0 \Rightarrow f(x_1) = 1)$
justifying-assign(v_i)	1	$f_{v_i \rightarrow 0}(x_1) = f_{v_i \rightarrow 1}(x_1)$
partial-equivalence(h)	1	$(h(x_1) = f(x_1)) \vee (h(x_1) \neq *)$

Figure 2: We show how to represent various queries as example-based queries. Equivalence, membership, subset, superset, disjointness and exhaustiveness queries are defined by Angluin [2]. A *justifying assignment* for an input variable is an instance whose classification changes if the value of the variable is changed. Thus, for Boolean domains, a justifying assignment query on v_i returns “yes” if there is no justifying assignment, or as a counterexample it returns two instances that provide a justifying assignment for v_i . (The notation $f_{v_i \rightarrow 0}$ denotes the function obtained from f by fixing $v_i = 0$.) Finally, in a partial equivalence query (as defined by Maass and Turán [19]) the learner can present a hypothesis $h : \mathcal{X} \rightarrow \{0, 1, *\}$, and is either told that all specified instances are correct or is given an $x \in \mathcal{X}$ such that $h(x) \in \{0, 1\}$ and x is misclassified by h .

Proof: We prove this result by demonstrating a valid T/L pair for any representation class \mathcal{C} learnable in deterministic polynomial-time by an algorithm \mathcal{A} that uses only example-based queries. We assume there is some total ordering π on \mathcal{X} (such as a lexicographical order) upon which the learner and teacher have agreed. Given any two sets of k instances from \mathcal{X} we define the following ordering among them. Let $x = (x_1, \dots, x_k)$ for $x_1 < x_2 < \dots < x_k$ be one set and let $y = (y_1, \dots, y_k)$ for $y_1 < y_2 < \dots < y_k$ be the other set. Then we say that $x < y$ (according to π) if there exists a $1 \leq j \leq k$ such that $x_j < y_j$ and for all $i < j$, $x_i = y_i$.

The teacher T constructs its teaching set $T(f)$ for f as follows. Initially, let $T(f) = \emptyset$. Now T simulates \mathcal{A} 's execution until the point at which the first example-based query q is performed. Let k be the number of instances over which q is quantified. The teacher now goes through all $(x_1, x_2, \dots, x_k) \in \mathcal{X}^k$ from smallest to largest, evaluating $\varphi(x_1, \dots, x_k)$. If there are no counterexamples to q then the teacher replies "yes" to \mathcal{A} 's query. Otherwise, let (x_1, x_2, \dots, x_k) be the smallest counterexample and let q_1, \dots, q_ℓ be the instances on which membership queries were made to evaluate $\varphi(x_1, \dots, x_k)$. Note that the number of membership queries made by φ_f (and thus ℓ) is polynomial since φ_f is poly-time computable. The teacher lets $S = \{x_1, f(x_1)\} \cup \dots \cup \{x_k, f(x_k)\} \cup \{q_1, f(q_1)\} \cup \dots \cup \{q_\ell, f(q_\ell)\}$, replies to \mathcal{A} with S , and updates $T(f)$ to be $T(f) \cup S$. The teacher continues in this manner until \mathcal{A} halts.

We now create the learner L from \mathcal{A} as follows. Let $S_A \supseteq T(f)$ be the teaching set that the learner receives. Whenever \mathcal{A} makes a query q , the learner will proceed as follows. The learner will consider all k -tuples in S_A from smallest to largest. For each such tuple (x_1, \dots, x_k) the learner attempts to evaluate $\varphi(x_1, \dots, x_k)$. In order to evaluate φ , recall that the learner may need to perform some additional membership queries. If these instances appear in S_A then the learner computes $\varphi(x_1, \dots, x_k)$. If $\varphi(x_1, \dots, x_k) = 0$, then L gives \mathcal{A} the k -tuple (x_1, \dots, x_k) along with the labeled examples corresponding to the membership queries made in evaluating φ on this k -tuple. What if the learner is unable to evaluate φ ? Since $T(f)$ contained the minimum counterexample for φ (including all instances needed to evaluate φ) and $S_A \supseteq T(f)$, it follows that (x_1, \dots, x_k) must not be a counterexample. Thus, if the learner computes

that $\varphi(x_1, \dots, x_k) = 1$ or is unable to evaluate it, then the learner continues with the next k -tuple in the ordering. If for all k -tuples in S_A the predicate φ is true or unevaluatable (i.e. there is no counterexample to q in S_A) then L responds to \mathcal{A} with “yes”. Observe that since q is quantified over a constant number of instances, in time polynomial in $|S_A|$ the learner can consider all k -tuples from S_A . Furthermore, because the teacher evaluated k -tuples of instances from minimum to maximum when constructing $T(f)$, the membership queries needed to evaluate φ will be present for the minimum counterexample.

We now argue that L will halt in polynomial time and output f . The key observation here is that the teacher’s and learner’s simulations of \mathcal{A} always remain the same. Since \mathcal{A} is deterministic its execution is altered only by the responses given to its queries. While the adversarial substitute may add other counterexamples, the learner will always find the minimum one, which was included in $T(f)$ by T , and thus T and L both give \mathcal{A} exactly the same counterexamples. \square

The following corollary follows directly from Theorem 2.

Corollary 3 *If representation class \mathcal{C} is not semi-poly T/L-teachable then it is not deterministically learnable in polynomial time using example-based queries.*

Thus, negative results obtained for a class in our model give very strong negative results with regards to the learnability of the class. It is possible that this correspondence will provide new techniques to prove hardness results for learning. As an immediate consequence of this result we know that many classes (namely, all of those for which exact-identification is efficiently achieved with queries) are T/L-teachable with an efficient learner. In particular, this contrasts the negative result of Jackson and Tompkins [16] that the class of 1-decision lists is not teachable without trusted information, and the negative result of Anthony et. al [5] that linearly separable Boolean functions are not efficiently teachable. In fact, Bshouty’s [8] result that arbitrary decision *trees* are learnable with membership and equivalence queries implies that a much broader class than 1-decision lists is T/L-teachable with a polynomial-time learner.

Letting \mathcal{A} in the proof of Theorem 2 be the halving algorithm [6, 18], we immediately get the following corollary.

Corollary 4 *Any representation class \mathcal{C} is T/L-teachable (by a computationally unbounded teacher and learner) with a teaching set of length at most $\log |\mathcal{C}|$.*

Because our model incorporates a very powerful set of queries, classes that may not be learnable using membership and equivalence queries are T/L-teachable. In particular, from Angluin’s result [2] that pattern languages can be exactly identified in polynomial time using only restricted superset queries, we get that pattern languages are semi-poly T/L-teachable. Lange and Wiehagen independently presented an algorithm to learn pattern languages from good examples [17].

It is also easily shown that the self-directed learning model [13] can be simulated in our model. Namely, we can build a valid T/L pair by having T simulate the self-directed learning algorithm \mathcal{A} , and then include in $T(f)$ only those examples that \mathcal{A} misclassifies. Then L can simply simulate \mathcal{A} assuming that any example not in $S_{\mathcal{A}} \supseteq T(f)$ is properly labeled by its hypothesis.

Theorem 5 *If there is a deterministic self-directed learning algorithm for representation class \mathcal{C} that makes polynomially bounded mistakes, then \mathcal{C} is T/L-teachable. Furthermore, the number of examples in an optimal teaching set is at most the number of mistakes made by the self-directed learning algorithm.*

Thus, for example, it follows that the classes of monomials and axis-parallel rectangles in $\{0, 1, \dots, n - 1\}^d$ are both T/L-teachable using only 2 examples. Furthermore, for these classes it is easily shown that both the teacher and learner can be efficiently implemented.

5 Efficient Teaching Strategies

While we know from Theorem 2 that any representation class learnable by an algorithm using example-based queries is T/L-teachable, we now demonstrate that often both the learner and the teacher use polynomial time and that for many classes the number of examples in $T(f)$ is asymptotically less than the number of queries made by existing learning algorithms. While we have currently only designed good teachers for existing

learning algorithms, in the long run we expect to design new learning algorithms that may not be good against an adversarial environment, but work very well when paired with an appropriate teacher.

In this section we give polynomial T/L pairs for the classes of decision lists and Horn sentences. Jackson and Tomkins show that the class of 1-decision lists with no irrelevant variables is teachable in their model. They prove, however, that 1-decision lists are not teachable in their model without trusted information.

5.1 Decision Lists

As defined by Rivest [22], a 1-decision list (1-DL) over the set $V_n = \{v_1, v_2, \dots, v_n\}$ of n Boolean variables is an ordered list $f = \langle (\ell_1, b_1), \dots, (\ell_r, b_r) \rangle$ where each ℓ_i is v_i or \bar{v}_i for $v_i \in V_n$ and each $b_i \in \{0, 1\}$. For an instance $x \in \{0, 1\}^n$, we define $f(x) = b_j$ where $1 \leq j \leq r$ is the least value such that ℓ_j is 1 in x ; $f(x) = \bar{b}_r$ if there is no such j . We refer to each pair (ℓ_i, b_i) as a node in f . Rivest also defines the class of k -decision lists (k -DL) as the generalization of 1-decision lists in which ℓ_i is any conjunction of at most k literals from V_n . Rivest presents an algorithm to learn the class of k -decision lists in the PAC model, and later Nick Littlestone⁵ constructed an algorithm to exactly identify k -DLs using only equivalence queries. When applied to the class of 1-DLs, Littlestone's algorithm uses $O(rn)$ equivalence queries and $O(rn^2)$ time.

For a 1-DL, f , let (ℓ_i, b_i) and (ℓ_j, b_j) be two nodes in f for $i < j$. If $\ell_i = \bar{\ell}_j$ then a logically equivalent decision list is obtained by replacing node j by the constant b_j and removing all following nodes. Similarly, if $\ell_i = \ell_j$ then node j can just be removed. We say a *reduced* 1-DL is one to which these reductions have been applied. Note that the number of nodes r in a reduced decision list is at most $n + 1$ (this includes the constant node at the end). We now show that the class of 1-DLs is polynomially T/L-teachable.

Theorem 6 *The class of 1-DLs is polynomially T/L-teachable with a teaching set of length at most $2r$ where r is the number of nodes in the reduced target decision list. The teacher requires $O(rn)$ time and the learner requires $O(r^2n)$ time.*

⁵This result is unpublished and may have been discovered independently by others.

T/L pair for 1-DL:

Teacher:

$$T(f) \leftarrow \emptyset$$

For each node i in the reduced target DL

$$T(f) \leftarrow T(f) \cup \{X_0^{(\ell_i)}\} \cup \{X_1^{(\ell_i)}\}$$

Learner:

All candidate nodes begin at level 1

Repeat

For each example $(x, f(x))$ in S_A

Let k be the highest level for which $\ell_i = 1$ in x , for (ℓ_i, b_i) at level k

For each candidate node i at level k

If $\ell_i = 1$ in x and $b_i \neq f(x)$ then increment level of node i

Until no candidate changes level

Output the 1-DL obtained by arbitrarily ordering the nodes within each level,
stopping when reaching a constant node

Figure 3: Algorithms for the teacher and learner for the class of 1-DL.

Proof: We informally describe the teacher and learner for this class. We assume that the target decision list f is already reduced. Observe that if it were not then the teacher could reduce it in $O(n^2)$ time. Let $X_b^{(\ell_i)}$ be the example for which literals $\ell_1, \dots, \ell_{i-1}$ are 0, literal ℓ_i is 1 and all remaining variables are assigned b , for b either 0 or 1. For $1 \leq i \leq r$, the teacher simply includes $X_0^{(\ell_i)}$ and $X_1^{(\ell_i)}$. Thus, the teacher generates only $\Theta(r)$ examples, each requiring $O(n)$ time to generate.

The ideas in the design of our learner are based the algorithm given by Littlestone to exactly identify 1-DLs using only equivalence queries. It is convenient to view target decision list f as “leveled” where each level contains consecutive nodes that have the same associated bit and level 1 is the top level. Thus the number of levels in f is exactly one plus the number of times that $b_i \neq b_{i+1}$. The learner initially assumes that all of the $4n + 2$ possible nodes are the first (or top) node in f . Then on example x , any node that would have caused x to be mislabeled is moved to the set of candidates for the second level of f . So in general, L maintains a set of candidates for all $n + 1$ possible levels of f . For a given example x , the learner finds the highest level (least index) containing some node (ℓ_i, b_i) for which $\ell_i = 1$ in x . Each node (ℓ_j, b_j) in this level such that $b_j \neq f(x)$ is moved down to the next level by L . Finally, when no such counterexample exists, L outputs a 1-DL by arbitrarily ordering the nodes within each level, stopping at the earliest level in which the constant 1 or 0 appears. See Figure 3 for a complete description of the teaching and learning algorithms.

We now argue that this is a valid T/L pair and that L runs in $O(r^2n)$ time. We use induction to show that after the k th iteration of the repeat loop all candidates that belong on level $j \geq k + 1$ have been bumped to at least level $k + 1$. Note that any literal reaching level $r + 1$ is irrelevant. The inductive hypothesis holds for $k = 0$ (i.e. at the beginning all nodes are at level 1). Given that all nodes belonging at level $j \geq k$ are at level k or greater after the $(k - 1)$ st iteration of the loop, we must show that after the k th iteration the nodes that belong at level $j \geq k + 1$ will be at level $k + 1$ or greater. Let $(\ell_i, b_i), \dots, (\ell_j, b_j)$ be the nodes of f that are at level k and assume, for ease of exposition, that for $i \leq t \leq j$, $b_t = 1$. Observe that $X_0^{(\ell_t)}$ and $X_1^{(\ell_t)}$ for $i \leq t \leq j$ bump all candidates that do not belong in level k except for those with an

associated bit of 1 (including the constant 1). For now, assume that $k + 3 \leq r$, and for $s = k + 1$ or $k + 3$, let $(\ell_s, 0)$ be any node in f at level s . Of the nodes that still must be bumped from level k , observe that $X_0^{(\ell_{k+1})}$ and $X_1^{(\ell_{k+1})}$ bump all but $(\bar{\ell}_{k+1}, 1)$ and all nodes of the form $(\bar{\ell}_t, 0)$ for $i \leq t \leq j$. Finally, these remaining nodes are bumped by $X_b^{(\ell_{k+2})}$ and $X_b^{(\ell_{k+3})}$ for $b \in \{0, 1\}$ where $(\ell_{k+2}, 1)$ is a node at level $k + 2$ in f . In the case that $k + 3 > r$ it is easily shown that the candidates that do not belong at level k are bumped by the examples associated with node r . Finally, notice that instances placed in S_A by the adversary will not affect the correctness of the hypothesis returned by the learner. \square

This result can be generalized to k -DL but to do so the teacher must present examples for node i that turn off all nodes above node i , turn on node i and individually turn off nodes below node i (rather than en masse as is done for 1-DL). This increases the number of examples required to $O(n^k)$ and increases the running time of the learner correspondingly. The running time of the teacher is no longer polynomial, assuming $P \neq NP$, since implicit in turning of all nodes above node i is a satisfiability problem that the teacher must solve.

5.2 Conjunctions of Horn Clauses

A Horn clause is a disjunction of literals at most one of which is unnegated. A Horn sentence is a conjunction of Horn clauses. Angluin, Frazier and Pitt [4] gave a polynomial-time algorithm to exactly identify an m -clause Horn sentence using $O(mn)$ equivalence queries, $O(m^2n)$ membership queries and, $\tilde{O}(m^2n^2)$ ⁶, where n is the number of variables in the instance space. Note that each Horn clause can be viewed as a logical implication in which the consequent contains the, at most one, unnegated variable.

A clause C of the Horn sentence f is violated by x when all variables in the antecedent (denoted $ant(C)$) are 1 in x and the variable in the consequent (denoted $cons(C)$) is 0 in x . For clause C , let N_C denote the minimum negative example that violates clause C . Namely, N_C is the example for which the variables in $ant(C)$ are 1 and all other variables are 0. Likewise, let P_C be the minimum positive example for

⁶The “soft-oh” notation is like the standard “big-oh” except that log factors are also left out.

which all variables in $\text{ant}(C)$ and $\text{cons}(C)$ are 1 (each remaining variable is set to 0 unless this causes some clause in f to be violated).

Theorem 7 *The class of Horn sentences is polynomially T/L-teachable using $2m$ examples where m is the number of clauses in the target Horn sentence. The teacher and learner each require $O(m^2n)$ time.*

Proof: The teacher first modifies the target Horn sentence f so that the clauses can be individually violated (see Figure 4). The resultant target f' is logically equivalent to the original target f and has the same number of clauses. Next, the teacher lets $T(f) = \{N_C \cup P_C \mid C \in f\}$. The first step of T takes at most $O(m^2n)$ time since there are $O(m^2)$ pairs of clauses and each clause can be modified at most $O(n)$ times. The second step takes at most $O(mn)$ time outputting a teaching set of length $2m$. Finally, the learner runs what is essentially the standard algorithm for learning Horn sentences [4]. Figure 4 gives details of T and L .

We now show that this learner and teacher are a valid T/L-pair. Each negative example, N_C , added to $T(f)$ by the teacher causes the learner to add one of the clauses of the target to its hypothesis (as well as some clauses that do not belong). However, it is easily shown that because the target was modified by the teacher, P_C will violate these extra clauses, thus causing the learner to remove them. There are $O(n)$ extra clauses for each true clause added. Since there are m clauses in the target, a total of $O(mn)$ clauses are added to the learner's hypothesis. Thus, L is polynomial and, using only examples in $T(f)$, returns a hypothesis equivalent to f' . Finally, since L selects the minimum negative counterexample at each step, it follows that any negative example placed in S_A by the adversary will not modify the learner's course. Likewise, any positive counterexamples placed in S_A by the adversary are easily shown to cause no harm. \square

The class of k -quasi Horn sentences is defined similarly to Horn sentences except that each clause can have at most k unnegated literals (consequents of size at most k). Angluin, Frazier and Pitt note that learning 2-quasi Horn sentences with membership queries is as hard as learning CNF formulas with membership queries. With the added power of a cooperative teacher, however, it is unclear if there exists a valid T/L pair

T/L pair for Horn sentences:

Teacher:

$T(f) \leftarrow \emptyset$

Repeat until no clause is changed

For all pairs of clauses in the target

If $ant(C_i) \subset ant(C_j)$

then replace C_j by $(ant(C_j) \wedge cons(C_i) \Rightarrow cons(C_j))$

For $i \leftarrow 1$ to m

$T(f) \leftarrow T(f) \cup \{P_{C_i}\} \cup \{N_{C_i}\}$

Learner:

$h \leftarrow \text{true}$

Repeat

While S_A contains a positive counterexample

delete from h all clauses violated

If S_A contains a negative counterexample

then choose one, not previously used, with fewest 1's

let the 1's in the counterexample define $ant(C_i)$

$h \leftarrow h \wedge (\bigwedge_{(j|v_j \notin ant(C_i))} (ant(C_i) \Rightarrow v_j))$

Until S_A contains no counterexamples

Figure 4: Algorithms for the teacher and learner for the class of Horn sentences.

for this class. A straightforward adaptation of the above T/L pair cannot be used to teach 2-quasi Horn sentences. The problem lies in attempting to alter the target so that clauses may be violated individually. Because each consequent may have size two there can be exponential blowup in the number of clauses in the resultant target function. Attempts at different approaches encountered similar difficulties (such as an exponential number of examples required to teach a polynomial number of clauses in an effort to anticipate membership queries). The polynomial T/L-teachability of 2-quasi Horn sentences remains an interesting open problem.

6 Relation to Data Compression

In this section we uncover an interesting relationship between teaching in our model and data compression. A *data compression* scheme of size k for representation class \mathcal{C} consists of a *compression algorithm* \mathcal{F} and a *reconstruction algorithm* \mathcal{G} . Let S_m be any subset of $m \geq k$ examples from \mathcal{X} labeled according to some $f \in \mathcal{C}$. The compression algorithm \mathcal{F} takes as input S_m and outputs some subset S of S_m such that $|S| \leq k$. The reconstruction algorithm \mathcal{G} takes as input any possible subset of at most k labeled examples and outputs a hypothesis h on \mathcal{X} . A *valid data compression scheme* of size k for \mathcal{C} consists of a pair \mathcal{F} and \mathcal{G} such that for any $f \in \mathcal{C}$ and any set S_m of at least k examples labeled according to f , the hypothesis output by $\mathcal{G}(\mathcal{F}(S_m))$ must be consistent with all examples in S_m . Let S_f be the labeled sample consisting of all instances in \mathcal{X} labeled according to f . Observe that $\mathcal{G}(\mathcal{F}(S_f))$ must be logically equivalent to f . Thus, a computationally unbounded teacher could produce the examples in $\mathcal{F}(S_f)$ as a teaching set. We now give sufficient conditions under which the learner can simply use \mathcal{G} to obtain a hypothesis logically equivalent to f .

Theorem 8 *If there is a valid data compression scheme of size k for representation class \mathcal{C} and $\mathcal{F}(S_f)$ produces an example set for which f is the minimum consistent hypothesis (for any pre-defined ordering of \mathcal{C}), then \mathcal{C} is T/L-teachable with at most k examples.*

Proof: The teacher will use $\mathcal{F}(S_f)$ as the teaching set. The learner will output the minimum consistent hypothesis consistent with the teaching set. Since, by the

conditions of the theorem, the target f is the minimum hypothesis consistent with $\mathcal{F}(S_f)$, the hypothesis output by the learner cannot be affected by the additional examples added to the teaching set by the adversary. \square

Many of the space-bounded learning algorithms that Floyd [9] presents satisfy the conditions of Theorem 8 and thus we immediately obtain results for our teaching model. To state this results, we must define the Vapnik-Chervonenkis dimension [29]. Let X be any instance space, and C be a concept class over X . A finite set $Y \subseteq X$ is *shattered* by C if $\{c \cap Y \mid c \in C\} = 2^Y$. In other words, $Y \subseteq X$ is shattered by C if for each subset $Y' \subseteq Y$, there is a concept $c \in C$ which contains all of Y' , but none of the instances in $Y - Y'$. The *Vapnik-Chervonenkis dimension* of C , denoted $\text{VCD}(C)$, is defined to be the smallest d for which no set of $d + 1$ points is shattered by C . Blumer et al. [7] have shown that this combinatorial measure of a concept class characterizes the number of examples required for learning any concept in the class under the distribution-free or PAC model of Valiant [28].

Related to the VC-dimension are the notions of *maximal* and *maximum* concept classes [9, 30]. A concept class is *maximal* if adding any concept to the class increases the VC dimension of the class. Define

$$\Phi_d(m) = \begin{cases} \sum_{i=0}^d \binom{m}{i} & \text{for } m \geq d \\ 2^m & \text{for } m < d. \end{cases}$$

If C is a concept class of VC-dimension d on a finite set X with $|X| = m$, then the cardinality of C is at most $\Phi_d(m)$ [26, 29]. A concept class C over X is *maximum* if for every finite subset $Y \subseteq X$, the class C , when restricted to be a class over Y , contains $\Phi_d(|Y|)$ concepts.

Floyd [9] shows that if C is a maximum class of VC-dimension d on the set X , then there is a data compression scheme of size d for C . Furthermore, it is easily shown that for this compression scheme, $\mathcal{F}(S_f)$ produces an example set for which f is the only consistent hypothesis. Thus from Floyd's results and Theorem 8, we get the following corollary.

Corollary 9 *For any maximum class C , there is a valid T/L pair such that the optimal teaching set has length at most the VC-dimension of C .*

We can show the corresponding result for intersection-closed classes by applying results from Helmbold, Sloan and Warmuth [15]. They define a spanning set of a representation $c \in \mathcal{C}$ with respect to the class \mathcal{C} to be a set $I \subseteq c$ having the property that c is the unique, most specific representation consistent with the instances in I and show that for intersection closed classes all minimum spanning sets have size at most $\text{VCD}(\mathcal{C})$. They then give an algorithm to learn the nested-difference of a set of size p of functions from \mathcal{C} while saving at most $p \cdot \text{VCD}(\mathcal{C})$ examples, thus giving the following result.

Corollary 10 *For any intersection-closed representation class \mathcal{C} for which there is an efficient query algorithm, the nested difference of p functions from \mathcal{C} is T/L-teachable using at most $p \cdot \text{VCD}(\mathcal{C})$ examples with a polynomial-time learner.*

7 Alternate Models

In this section we briefly describe some variations of our model that we feel are worthy of study. The model we presented here places minimal restrictions on the T/L pair while ensuring that there is no collusion. For some applications, one may want to limit the assumptions the teacher may make about the learner without going to the extreme of only allowing the teacher to assume that the learner is consistent. For example, one class of learners that would be interesting to study is that of learners that always select a minimum (in terms of the number of instances classified as positive) consistent hypothesis. In fact, this type of learner was studied by Natarajan [21] in terms of one-sided learning. As another example, we could consider the class of learners that only select a minimal consistent hypothesis. This corresponds to requiring that the learner always selects an element from S in Mitchell's version space [20].

Another interesting variation is one in which the learner is not required to exactly identify the target, but rather needs only output an ϵ -good approximation to the target. (Romanik and Smith [23, 24] consider a PAC-style criterion in their work.)

Finally, another interesting model is one in which there are two distinct stages. The first stage operates as in our current model except that the teacher is not required to

provide examples sufficient for exact identification. In the second stage the teacher lists all instances that are exceptions to the rule taught during the first stage. During the second stage the learner just appends to its hypothesis the list of “exceptions” given by the teacher. The motivation behind this model is that it is often easiest to teach by first oversimplifying the truth and then making the needed corrections. For example, in teaching a child about spelling, you could first teach the rule that the letter “i” always comes before “e”. Then, you can add the needed exceptions. It is also possible that the list of exceptions could take the form of another rule. For example, “i” comes before “e” except after “c”.

8 Concluding Remarks

We now briefly discuss the two directions of future research that we find most interesting. First of all, an interesting open question is to determine whether there exist teacher/learner pairs in which both the teacher and learner require only polynomial computation time for such classes as 2-quasi Horn Sentences and read-thrice DNF that appear difficult to learn from queries [3, 1].

Another interesting research direction opened up by this new model is the following. As we did for 1-decision lists and Horn sentences, one can take known learning algorithms and design teachers that enable the sample and time complexity of the algorithm to be asymptotically reduced. However, the algorithms that have been designed to work against an adversarial environment are most likely not going to be the best algorithms when we allow the algorithm to be paired with a teacher. In other words, for some classes we expect that there should be teacher/learner pairs which work better than any known algorithm when paired with the best possible teacher, yet the algorithm used by the learner may be very poor against an adversarial environment. Such a study could lead to general techniques for designing teacher/learner pairs.

Acknowledgements

We are very grateful to Dana Angluin for many valuable discussions on this topic. In particular, Theorem 1 was suggested by her and the extension of Theorem 2 to

include all example-based queries (along with the definition of example-based queries) was due to her. In addition, the model of teaching with patching was developed in collaboration with her. We thank Rob Schapire for suggesting that we explore the relationship between teaching and data compression. The observation of T/L cooperation for classes of size at most $2^{O(n)}$ is due to Avrim Blum and we thank him for that. Finally, we thank Lenny Pitt for his help in formalizing the notion of collusion and for describing to us Nick Littlestone's algorithm for exactly identifying decision lists.

References

- [1] Howard Aizenstein, Lisa Hellerstein, and Leonard Pitt. Read-thrice DNF is hard to learn with membership and equivalence queries. In *33rd Annual Symposium on Foundations of Computer Science*, pages 523–532, October 1992.
- [2] Dana Angluin. Queries and concept learning. *Machine Learning*, 2(4):319–342, 1988.
- [3] Dana Angluin, Michael Frazier, and Leonard Pitt. Learning conjunctions of horn clauses. In *31st Annual Symposium on Foundations of Computer Science*, pages 186–192, October 1990.
- [4] Dana Angluin, Michael Frazier, and Leonard Pitt. Learning conjunctions of horn clauses. *Machine Learning*, 9:147–164, 1992. Special Issue for COLT 90.
- [5] Martin Anthony, Graham Brightwell, Dave Cohen, and John Shawe-Taylor. On exact specification by examples. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, pages 311–318. ACM Press, July 1992.
- [6] J. Barzdin and R. Freivald. On the prediction of general recursive functions. *Soviet Mathematics Doklady*, 13:1224–1228, 1972.
- [7] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the Association for Computing Machinery*, 36(4):929–965, October 1989.
- [8] Nader H. Bshouty. Exact learning via the monotone theory. In *34th Annual Symposium on Foundations of Computer Science*, November 1993. To appear.

- [9] Sally Floyd. *On Space-bounded Learning and the Vapnik-Chervonenkis Dimension*. PhD thesis, International Computer Science Institute, December 1989.
- [10] Rūsiņš Freivalds, Efim Kinber, and Rolf Wiehagen. Inference from good examples. *Theoretical Computer Science*, 110:131–144, 1993.
- [11] Sally A. Goldman and Michael J. Kearns. On the complexity of teaching. In *Proceedings of the Fourth Annual Workshop on Computational Learning Theory*, pages 303–315. Morgan Kaufmann, August 1991.
- [12] Sally A. Goldman, Ronald L. Rivest, and Robert E. Schapire. Learning binary relations and total orders. In *30th Annual Symposium on Foundations of Computer Science*, pages 46–51, October 1989. To appear in *SIAM Journal of Computing*.
- [13] Sally A. Goldman and Robert H. Sloan. The power of self-directed learning. Technical Report WUCS-92-49, Washington University, Department of Computer Science, November 1992. To appear in *Machine Learning*.
- [14] S. Goldwasser, S. Goldwasser, and C. Rackoff. The knowledge complexity of interactive proofs. In *26th Annual Symposium on Foundations of Computer Science*, pages 291–304, October 1985.
- [15] David Helmbold, Robert Sloan, and Manfred K. Warmuth. Learning nested differences of intersection-closed concept classes. *Machine Learning*, 5:165–196, 1990. Special issue for COLT 89.
- [16] Jeffrey Jackson and Andrew Tomkins. A computational model of teaching. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, pages 319–326. ACM Press, July 1992.
- [17] Stefan Lange and Rolf Wiehagen. Polynomial-time inference of arbitrary pattern languages. *New Generation Computing*, 8:361–370, 1991.
- [18] Nick Littlestone. Learning when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2:285–318, 1988.
- [19] Wolfgang Maass and György Turán. Lower bound methods and separation results for on-line learning models. *Machine Learning*, 9:107–145, 1992.

- [20] Tom M. Mitchell. Generalization as search. *Artificial Intelligence*, 18(2):203–226, 1982.
- [21] B. K. Natarajan. On learning Boolean functions. In *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, pages 296–304, May 1987.
- [22] Ronald L. Rivest. Learning decision lists. *Machine Learning*, 2(3):229–246, 1987.
- [23] Kathleen Romanik. Approximate testing and learnability. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, pages 327–332. ACM Press, July 1992.
- [24] Kathleen Romanik and Carl Smith. Testing geometric objects. Technical Report UMIACS-TR-90-69, University of Maryland College Park, Department of Computer Science, 1990.
- [25] Steven Salzberg, Arthur Delcher, David Heath, and Simon Kasif. Learning with a helpful teacher. In *12th International Joint Conference on Artificial Intelligence*, pages 705–711, August 1991.
- [26] N. Sauer. On the density of families of sets. *Journal of Combinatorial Theory (A)*, 13:145–147, 1972.
- [27] Ayumi Shinohara and Satoru Miyano. Teachability in computational learning. *New Generation Computing*, 8:337–347, 1991.
- [28] Leslie Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, November 1984.
- [29] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and Its Applications*, XVI(2):264–280, 1971.
- [30] E. Welzl. Complete range spaces. Unpublished manuscript, 1987.