

Active sampling for multiple output identification

Shai Fine · Yishay Mansour

Received: 28 September 2006 / Revised: 10 September 2007 / Accepted: 11 September 2007 /
Published online: 2 October 2007
Springer Science+Business Media, LLC 2007

Abstract We study functions with multiple output values, and use active sampling to identify an example for each of the possible output values. Our results for this setting include: (1) Efficient active sampling algorithms for simple geometric concepts, such as intervals on a line and axis parallel boxes. (2) A characterization for the case of binary output value in a transductive setting. (3) An analysis of active sampling with uniform distribution in the plane. (4) An efficient algorithm for the Boolean hypercube when each output value is a monomial.

Keywords Active learning · Active sampling · Hitting · VC dimension · Transductive learning · Output identification · Separation dimension

1 Introduction

Active sampling is much about “hitting” low probability events. In active learning, the active sampling is used to guide the learning process to learn a high accuracy hypothesis while using a limited number of examples (Seung et al. 1992; Kulkarni et al. 1993; Freund et al. 1997; Liere and Tadepalli 1997; Fine et al. 2002; Dasgupta et al. 2005;

Editors: Hans Ulrich Simon, Gabor Lugosi, Avrim Blum.

A preliminary version of this work was presented in the 19th Annual Conference on Learning Theory (COLT), 2006.

This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778, by a grant No. 1079/04 from the Israel Science Foundation, by a grant from BSF and an IBM faculty award. This publication only reflects the authors’ views.

S. Fine (✉)

IBM Research Laboratory in Haifa, Mount Carmel, Haifa 31905, Israel
e-mail: fshai@il.ibm.com

Y. Mansour

School of Computer Science, Tel Aviv University, Tel Aviv, Israel
e-mail: mansour@post.tau.ac.il

Dasgupta 2004; Dasgupta 2005). While in many applications the goal of an accurate hypothesis is the most natural one, there are other applications which require only examples of those low probability events. Examples of such application domains include hardware and software verification, fault tolerance, network security, data mining, etc. The usage of such examples in each of the applications can be very different: in fault tolerance one would like to simulate the performance of the system in extreme conditions (e.g., very high load), in network security one would like to have examples of potential intruders, while in data mining one would like to find new interesting relationships, which are not explained by the existing ones. Our original motivation, though, stem from dynamic hardware verification and from software testing. In both domains, the main industrial vehicles are simulation-based methods which aimed at exciting (and impacting) the occurrence of events and scenarios of desired functional behaviors that need to be verified (Wile et al. 2005; Edelstein et al. 2002). Coverage (Piziali 2004) is an information collection mechanism that is often used to monitor the progress of the verification process, and point to areas in the design that have not been properly tested.

The analysis of coverage reports, and their translation to a set of directives that guides the implementation of the test plan, result in major manual bottlenecks in the otherwise highly automated verification process. A verification methodology called *coverage directed test generation* (CDG) aims to resolve this problem, either by utilizing a mechanism that can directly translate a verification task into a simulation test (CDG by Construction, cf. Ur and Yadin 1999), or by extracting useful information from the observed events, and use it to bias future simulation runs (CDG by Feedback, cf. Fine and Ziv 2003 and references therein). However, too often, neither an accurate translation mechanism nor well structured coverage model can be provided, and we end up with the following naive, yet difficult, scenario: The verification team is given a list of events that should be covered, and the goal is to provide multiple sets of directives (inputs) that will tune the test generator to produce patterns that hit all items in the list. In these situations the common practice is “trial & error”.

We abstracted the above motivation in the following learning model. There is an unknown target function f which maps every input in \mathcal{X} to one of m output values. (For example, in the verification setting the output values would be desired scenarios for coverage.) The *output identification task* is to find m inputs, one for each output value.

The output identification algorithm is given some information about the target function. First, like much of the computational learning literature, it knows that the target function f is in a given function class \mathcal{F} . Second, it is given the number of output values, i.e., m (hence it knows when to terminate).

We assume that there is an unknown distribution D over the inputs. The algorithm has an access to an *induced distribution example oracle* which allows it to sample from sub-regions of the domain. Namely, the algorithm specifies a subset $Y \subset \mathcal{X}$, and the oracle returns an example from Y , distributed according to the induced distribution of D over Y . The goal of the output identification algorithm is to minimize the number of oracle samples it requires until a representative for each output value is found. The performance is measured as a function of the number of output values m and ϵ , a lower bound on the probability of each output value. (We note in passing that only for simplicity we assume that m and ϵ are known to the output identification algorithm, both of the requirements can be easily relaxed, and similar results hold.)

It is worth while pointing out that any m valued function can be output identified with an expected active sample of $O((\log m)/\epsilon)$ by simply sampling from the distribution, and this bound holds regardless of the complexity of the target function. (The bound follows immediately from the *coupon collector problem*, see Feller (1966).) This gives a trivial expected active sample bound, and clearly we would like our algorithms to outperform it.

Conceptually, we would refer to active sample bounds as being efficient when they are polylogarithmic in $1/\epsilon$, and inefficient when they are linear (or polynomial) in $1/\epsilon$.

We present efficient algorithms for many classes of functions. We start by presenting efficient output identification algorithms for a few simple geometric classes. For the function class of m intervals on a line we show an expected active sample bound of $O(m \log 1/\epsilon)$. For m axis parallel boxes in \mathbb{R}^d we give an expected active sample bound of $O(md \log d/\epsilon)$. We also derive lower bounds that exhibit classes with a constant VC dimension, such as a linear separator in the plane, which require $\Omega(1/\epsilon)$ active samples.

Our main result is a characterization for the case of binary outputs, i.e., $m = 2$. We define a separation dimension and show that if the separation dimension is d then the function class can be output identified in $O(d^2 \log^2 s)$ queries in a transductive setting, where the output identification algorithm is given a set of s unlabeled examples in advance. In addition, we show that for any function class with separation dimension d , there is an unlabeled sample for which no algorithm can output identify it in less than $\Omega(\min\{d, s\})$ queries.

The separation dimension is similar in flavor to the VC dimension (Vapnik and Chervonenkis 1971). It requires d points such that the function class induced on them has all the d singleton functions (rather than all 2^d possible functions in the VC dimension). We show that when a class has a separation dimension d , then some input in the unlabeled sample can be queried and guarantee that either we terminate (finding a representative for each output value) or we reduce the space of consistent functions considerably. Using this property we derive an efficient output identification algorithm.

We also study the case of specific distributions, namely the uniform distribution in the plane. Using classical results from computational geometry we can show that many classes are efficiently learnable under the uniform distribution. Specifically, we show that the class of linear separators in the plane can be output identified with expected $O(m^2 \log^2 1/\epsilon)$ active samples with respect to the uniform distribution over the unit square.

We conclude with a concept class defined over the Boolean hypercube $\{0, 1\}^n$. We show that the class where each output value is represented by a monomial can be output identify in mn active samples.

Our model has obvious connections to active learning (Seung et al. 1992; Kulkarni et al. 1993; Freund et al. 1997; Liere and Tadepalli 1997; Fine et al. 2002; Dasgupta et al. 2005; Dasgupta 2004, 2005). In some sense the output identification task is much simpler than the usual learning task, since we do not need to find an accurate hypothesis, only to target one example for each output value. Still it seems that the techniques we present here and the active learning techniques share much in common. In both cases the goal is to reduce uncertainty, however since the tasks are different (learning vs. identification), so does the choice which of the samples to query—while an active learner will choose to query for the label of the sample which maximizes disagreement between the consistent hypotheses in the version space (cf. QBC, Seung et al. 1992; Freund et al. 1997), an output identifier will select to query for the label of the sample which maximize the probability of an unseen output value (ideally have the probability be almost one). Another major difference is that the research in active learning has been focused on binary classification while the main motivation for output identification are cases with a large number of possible output values.

There are simple cases in which active learning fails to achieve a significant sample improvement, for example linear separators in the plane (Dasgupta 2004). In such cases one should expect the output identification task to suffer from similar drawbacks (and indeed some of our lower bounds are much in that spirit).

A somewhat related question was discuss in Linial et al. (1997) where efficient deterministic constructions for combinatorial hitting sets are given. A hitting set for a domain

$\{1, \dots, m\}^d$ guarantees to “hit” any combinatorial rectangle of volume at least ϵ , i.e., any combinatorial rectangle that includes at least ϵm^d points would intersect the hitting set in at least one point. The main contribution of Linial et al. (1997) is to construct such a hitting set *deterministically* (i.e., without any randomization) and have its size and computation time be polynomial in m , d and $1/\epsilon$. (Our setting can be viewed as *active hitting* where the goal is to hit each output value at least once.)

Following the presentation of the preliminary version of this work in the 19th Annual Conference on Learning Theory (COLT), significant progress has been made on some of our open problems. Srebro (2006) has extended our definition of a separation dimension to multi-valued function, deriving similar results for his characterization. Gilad-Bachrach (2006) was able to improve our upper bound for the active sample complexity of a separation dimension by both simplifying the proof and reducing the active sample complexity to match our lower bound.

2 The model

A function f is m -valued if f maps inputs from a domain \mathcal{X} to an m valued range $\{1, \dots, m\}$. An m -valued function class \mathcal{F} is a set of m valued functions.

The *output identification task* for an m valued function class \mathcal{F} is as follows. There is an unknown m valued *target function* $f \in \mathcal{F}$ and the goal of the output identification algorithm is to identify an input for each of the m output values, i.e., find examples, $x^1, \dots, x^m \in \mathcal{X}$ such that $f(x^i) = i$.

The output identification algorithm has access to examples of the target function as follows. There is some unknown distribution D over the domain \mathcal{X} . Given a subset of the domain $Y \subset \mathcal{X}$ let D_Y be the distribution D induces over Y . An *induced distribution example oracle* receives as an input a subset $Y \subset \mathcal{X}$ and returns a pair $\langle x, f(x) \rangle$, where x is distributed according to D_Y and f is the target function. (We assume that Y has a non-empty intersection with the support of D , otherwise the oracle would generate an error.)

At each time step t , the output identification algorithm specifies a subset $Y_t \subset \mathcal{X}$ to the induced distribution example oracle and receives an example $\langle x_t, f(x_t) \rangle$. The process terminates when the algorithm has an example for each of the m output values, i.e. $x^1, \dots, x^m \in \mathcal{X}$ such that $f(x^i) = i$.

In order to measure the complexity of an output identification algorithm we assume that each output value has a probability of at least ϵ under the (unknown) distribution D .

The *active sample complexity* of an output identification algorithm with respect to a distribution D and function $f \in \mathcal{F}$ is the number of examples it requests, i.e., the number of times it accesses the induced distribution example oracle. The active sample complexity of an output identification algorithm for a function class \mathcal{F} is the worst case over all $f \in \mathcal{F}$ and distributions D of its active sample complexity. The active sample complexity of a function class \mathcal{F} is the least active sample complexity of any output identification algorithm for \mathcal{F} .

3 Simple geometric concepts

In this section we consider simple geometric concepts where the domain is $\mathcal{X} = \mathbb{R}^d$. We describe the *generic consistency algorithm* and analyze its expected sample complexity when applied for active sampling two function classes—intervals on a line ($d = 1$), and axis parallel boxes (for arbitrary d). For the first function class, where each of the m output values

is represented by an interval, the expected sample complexity is $O(m \log 1/\epsilon)$. For the case of axis parallel boxes, our result extends to an algorithm whose expected active sample complexity is $O(dm \log d/\epsilon)$. We conclude the section with a few simple lower bounds, showing examples of concept classes with a finite VC dimensions which require $\Omega(1/\epsilon)$ active sample complexity.

3.1 Generic consistency algorithm

We assume that the points mapped to a specific output value belong to some function class \mathcal{A} . (For example, for intervals $A \in \mathcal{A}$ would be a single interval and for axis parallel boxes it would be a single axis parallel box.)

The *generic consistency algorithm* works as follows. Initially we have $C_0 = \emptyset$. At phase $t \geq 1$ we sample the induced distribution example oracle with the subset $\mathcal{X} - C_{t-1}$ and receive an example $\langle x_t, k_t \rangle$. For each output value k let S_t^k be the set of examples sampled with output value k until time t . Let C_t^k be the *minimal concept* in A which is consistent with the examples in S_t^k (in our applications such a concept always exists). Let C_t be the union of all the C_t^k and proceed to phase $t + 1$.

The generic consistency algorithm starts at phase $t = 1$ and terminate at the first phase where for every output value k we have $S_t^k \neq \emptyset$ (note that in such a case we have for each of the m output values at least one example).

For clarity, we explicitly define the minimal consistent concept, for the cases of intervals and axis parallel boxes. In the case that A is an interval, then C_t^k is an interval $[\lambda_-^k, \lambda_+^k]$ such that $\lambda_-^k = \min\{x \in S_t^k\}$ and $\lambda_+^k = \max\{x \in S_t^k\}$. In the case that A is an axis parallel box then C_t^k is $([c_-^1, c_+^1], \dots, [c_-^d, c_+^d])$ where $c_-^i = \min\{x_i : x \in S_t^k, x = x_1, \dots, x_d\}$ and $c_+^i = \max\{x_i : x \in S_t^k, x = x_1, \dots, x_d\}$.

The correctness of the generic consistency algorithm is obvious from its termination condition, the main interest in the analysis would be on the expected number of examples until termination, i.e., the expected active sample complexity.

3.2 Intervals on a line

In this function class the domain is $\mathcal{X} = [0, 1]$, the examples corresponding to an output value k are in an interval A^k (which can be either open or closed interval), and the intervals are a partition of the domain $\mathcal{X} = [0, 1]$, i.e., $\bigcup_{k=1}^m A^k = [0, 1]$ and $A^i \cap A^j = \emptyset$ for $i \neq j$.

For the analysis of the generic consistency algorithm we introduce some additional notation. At time t , we have a set KN_t of output values which have already been sampled, and a set UKN_t of output values which have not been sampled. For each output $k \in KN_t$ let B_-^k and B_+^k be the points in A^k below and above C^k (respectively). I.e., if $C_t^k = [\lambda_-^k, \lambda_+^k]$ and $A^k = [\rho_-, \rho_+]$ then $B_-^k = [\rho_-, \lambda_-]$ and $B_+^k = (\lambda_+, \rho_+]$. Let $C_t = \bigcup_{k \in KN_t} C_t^k$.

Given a distribution D over $[0, 1]$ let $D(I)$ be the probability of the interval I . At time t let the probability of the unobserved output values be $\beta_t = \sum_{k \in UKN_t} D(A^k)$ and let $\gamma_t = \sum_{k \in KN_t} D(B_-^k) + D(B_+^k)$ be the probability of the observed values which are still in $\mathcal{X} - C_t$. Let the probability of the sampled region be $\alpha_t = D(\mathcal{X} - C_t) = \beta_t + \gamma_t$. Let \mathcal{H}_t include all the history of the execution of the algorithm until and including time t .

Our analysis uses a potential function $\Phi_t = \alpha_t + \beta_t = \gamma_t + 2\beta_t$, and shows that Φ_t decreases, in expectation, by a certain factor each time step. Specifically we show that,

$$E[\Phi_t - \Phi_{t+1} | \mathcal{H}_t] = \frac{\beta_t}{\alpha_t} \sum_{k \in UKN_t} \frac{D(A^k)}{\beta_t} D(A^k) + \frac{\gamma_t}{\alpha_t} \sum_{k \in KN_t, b \in \{+, -\}} \frac{D(B_b^k)}{\gamma_t} \cdot \frac{D(B_b^k)}{2}.$$

The first part follows from the fact that with probability $\frac{\beta_t}{\alpha_t}$ we sample an example with output value in UKN_t . Given that we sample such a point, the probability that the interval is A^k is $\frac{D(A^k)}{\beta_t}$. Given that we sample from A^k we have that $\beta_{t+1} = \beta_t - D(A^k)$ and $\gamma_{t+1} = \gamma_t + D(A^k)$, so the net reduction in the potential is $D(A^k)$, i.e., $\Phi_{t+1} = \Phi_t - D(A^k)$.

The second part follows from the fact that with probability $\frac{\gamma_t}{\alpha_t}$ we sample an example with output value in KN_t . Given that we sample such a point, the probability that it is in the interval B_b^k is $\frac{D(B_b^k)}{\gamma_t}$, where $b \in \{-, +\}$. Given that we sample from B_b^k the expected reduction in γ is $D(B_b^k)/2$, i.e., $E[\gamma_{t+1}] = \gamma_t - D(B_b^k)/2$. Therefore,

$$E[\Phi_t - \Phi_{t+1} | \mathcal{H}_t] = \frac{1}{\alpha_t} \left[\sum_{k \in UKN_t} D(A^k)^2 + \sum_{k \in KN_t, b \in \{+, -\}} \frac{1}{2} D(B_b^k)^2 \right].$$

Recall that

$$\alpha_t = \beta_t + \gamma_t = \sum_{k \in UKN_t} D(A^k) + \sum_{k \in KN_t, b \in \{+, -\}} D(B_b^k).$$

Using the general inequality $\sum_{i=1}^n X_i^2 \geq (1/n)(\sum_{i=1}^n X_i)^2$, and since there are at most $2m$ elements in the summation (each output value appears only in one of the two summations), we have that

$$\sum_{k \in UKN} D(A^k)^2 + \sum_{k \in KN, b \in \{+, -\}} \frac{1}{2} D(B_b^k)^2 \geq \frac{1}{4m} \alpha_t^2.$$

Since by definition $\alpha_t \geq \beta_t$, this implies that

$$E[\Phi_t - \Phi_{t+1} | \mathcal{H}_t] \geq \frac{1}{4m} \alpha_t \geq \frac{1}{8m} [\alpha_t + \beta_t] = \frac{1}{8m} \Phi_t.$$

By averaging over \mathcal{H}_t we have that

$$E[\Phi_{t+1}] \leq \left(1 - \frac{1}{8m}\right) E[\Phi_t] \leq \left(1 - \frac{1}{8m}\right)^t \Phi_1.$$

Initially we have $\gamma_1 = 0$ and $\beta_1 = 1$, therefore the initial potential is $\Phi_1 = \gamma_1 + 2\beta_1 = 2$. After $t = O(m \log(1/\epsilon))$ samples, the expected value of the potential is less than $\epsilon/2$. This implies that with probability at least $1/2$ its value is less than ϵ . Once the value of the potential is less than ϵ we are guaranteed to hit each output value (since each output value has probability at least ϵ). This establishes the following theorem.

Theorem 1 *The class of m intervals can be output identified in expected active sample complexity of $O(m \log 1/\epsilon)$.*

3.3 Axis parallel boxes

We extend the results from intervals to axis parallel boxes, where $\mathcal{X} = [0, 1]^d$, each output value k is represented by an axis parallel box A^k , and the collection of A^k are a partition of \mathcal{X} . Again, we use the generic consistency algorithm.

The analysis is similar in spirit to that of the intervals on a line and it appears in the appendix, where we establish the following theorem:

Theorem 2 *The class of m axis parallel boxes can be output identified in expected active sample complexity of $O(dm \log d/\epsilon)$.*

3.4 Lower bounds

In this section we derive two simple lower bounds.

Example 1 Let $\mathcal{X} = [0, 1]$ and U be the uniform distribution on \mathcal{X} . Consider the following function class $\mathcal{F}_{\text{seg}}^\eta$ that includes functions $f_z(x) = 1$ if $x \in (z, z + \eta)$ and otherwise $f_z(x) = 0$ (where $z \in [0, 1 - \eta]$ and $\eta < 1$). Note that $\mathcal{F}_{\text{seg}}^\eta$ has a VC dimension equals to 2, and in addition, the points with output value 0 are not a convex set. We show the following lower bound.

Claim *Any output identification algorithm for $\mathcal{F}_{\text{seg}}^\eta$ with the uniform distribution U requires an expected active sample complexity of $\Omega(1/\epsilon)$, when $\epsilon = \eta$.*

Proof Consider a subset \mathcal{F}' of $\mathcal{F}_{\text{seg}}^\eta$ which include the following $1/\epsilon$ functions: $f_{k\epsilon}$ for integers $0 \leq k \leq \lfloor 1/\epsilon \rfloor - 1$. Assume that one of the functions in \mathcal{F}' is selected at random. For $\eta = \epsilon$ the functions in \mathcal{F}' are not overlapping, i.e., there is no point x such that both $f_{k\epsilon}(x) = 1$ and $f_{k'\epsilon}(x) = 1$ for $k \neq k'$. Therefore, any point x for which the target function $f \in \mathcal{F}'$ has $f(x) = 1$, uniquely identifies the function in \mathcal{F}' . This implies that any output identification algorithm would require expected active sample complexity of $\Omega(1/\epsilon)$ to identify the randomly selected function. \square

Example 2 Consider a linear separator in the plane (there are only two output values). Let $\mathcal{X} = [-1, 1]^2$ and let \mathcal{F}_{ls} include all linear separators. Namely, for each $f_{\alpha,\beta} \in \mathcal{F}$ we have $f_{\alpha,\beta}(x) = 1$ if $\alpha x_1 + \beta < x_2$ and otherwise $f_{\alpha,\beta}(x) = 0$ (where $\alpha, \beta \in \mathbb{R}$).

Following Dasgupta (2004), we consider a distribution U_o whose support is the unit circle, e.g., $(x_1)^2 + (x_2)^2 = 1$ and it is uniform over it. Similar to the lower bound for active learning (Dasgupta 2004), the following lower bound for output identification holds.

Claim *Any output identification algorithm for \mathcal{F}_{ls} with distribution U_o requires an expected active sample complexity of at least $\Omega(1/\epsilon)$.*

4 Transductive setting: binary output values

In the transductive setting the algorithm is given in advance a set of unlabeled examples $S = \{x_1, \dots, x_s\}$. The goal of the output identification algorithm is to find a subset $S' \subset S$ of size at most m , such that each output value that appears in S has an example in S' . I.e., let $S^k = \{x \in S : f(x) = k\}$, we require that if $S^k \neq \emptyset$ then $S^k \cap S' \neq \emptyset$. The active sample complexity in the transductive setting is the number of queries the algorithm makes (i.e., the number of unlabeled examples from S for which it asks a label).

We give a characterization for the case of binary output values, i.e., $m = 2$. We first define the notion of a *separation dimension* of a function class \mathcal{F} . Then we show that if a function class \mathcal{F} has separation dimension d then there is an algorithm that queries only $O(d^2 \log^2 s)$ examples. In addition we show that if a function class has separation dimension d then the expected number of examples queried is $\Omega(\min\{d, s\})$. This implies that the for the binary case we have a complete characterization when can a function class be output identified with a poly-logarithmic number of queries.

4.1 Separation dimension: definition

We start by defining the notion of separation dimension. Let the separation dimension of a function class \mathcal{F} be the following. A function class \mathcal{F} is said to b -separate the set $\{x_1, \dots, x_d\} \subset \mathcal{X}$ if there are functions $f_1, \dots, f_d \in \mathcal{F}$ such that $f_i(x_i) = b$ and $f_i(x_j) = 1 - b$, for $j \neq i$, where $b \in \{0, 1\}$.

The b -separation dimension of \mathcal{F} is the size of the largest set that \mathcal{F} b -separates (or infinity, if it can b -separate sets of arbitrary size). The separation dimension of \mathcal{F} is the maximum of the 0-separation dimension and the 1-separation dimension.

Separation dimension—examples Consider the function class \mathcal{F}_{pre} over $[0, 1]$ such that $f_z(x) = 1$ if $z \leq x$ and $f_z(x) = 0$ otherwise (where $z \in [0, 1]$). The separation dimension of \mathcal{F}_{pre} is 1 since for any two points $x_1 < x_2$ no function f_z can have $f_z(x_1) = 1$ and $f_z(x_2) = 0$.

A simple extension of \mathcal{F}_{pre} is the function class $\mathcal{F}_{\text{pre+suf}}$ where $f_{z,b}(x) = b$ if $z \leq x$ and $f_{z,b}(x) = 1 - b$ otherwise (where $z \in [0, 1]$ and $b \in \{0, 1\}$). The separation dimension of $\mathcal{F}_{\text{pre+suf}}$ is 2. (Given, for example, $x_1 = 1/3$ and $x_2 = 2/3$, the functions $f_{1/2,1}$ and $f_{1/2,0}$ achieve a b -separation of $\{x_1, x_2\}$ for both $b = 1$ and $b = 0$. However, for any three points $x_1 < x_2 < x_3$ no function $f_{z,b}$ can have $f_{z,b}(x_2) = b$ and $f_{z,b}(x_1) = f_{z,b}(x_3) = 1 - b$, neither for $b = 1$ nor for $b = 0$.)

Recall the function class $\mathcal{F}_{\text{seg}}^\eta$ that includes functions $f_z(x) = 1$ if $x \in (z, z + \eta)$ and otherwise $f_z(x) = 0$ (where $z \in [0, 1 - \eta]$ and $\eta < 1$). The function class $\mathcal{F}_{\text{seg}}^\eta$ has a separation dimension of approximately $\lfloor 2/\eta \rfloor$.

Claim *The function class $\mathcal{F}_{\text{seg}}^\eta$ has separation dimension at least $\lfloor 2/\eta - 1 \rfloor$ and at most $\lfloor 2/\eta + 2 \rfloor$.*

Proof First we show that the separation dimension is at least $s = \lfloor 2/\eta - 1 \rfloor$. Consider the s points $x_i = i(\eta/2)$, $i \in \{1, \dots, \lfloor 2/\eta - 1 \rfloor\}$. The set of functions $f_{k\eta/2}$ for $k \in \{0, \dots, \lfloor 2/\eta \rfloor - 2\}$, 1-separates the s points. (Note that $f_{(i-1)\eta/2}(i\eta/2) = 1$ and $f_{(i-1)\eta/2}(j\eta/2) = 0$, for $i \neq j$.) This implies that the separation dimension is at least $s = \lfloor 2/\eta - 1 \rfloor$.

On the other hand, for any odd integer $s \geq 3$, and any set of s points in $[0, 1]$, $x_1 < \dots < x_s$, for some i we have $x_{i+1} - x_{i-1} \leq 2/(s - 1)$. If $\eta > 2/(s - 1)$ then for any function $f_z \in \mathcal{F}_{\text{seg}}^\eta$, we can not have $f_z(x_i) = 1$ and $f_z(x_{i+1}) = f_z(x_{i-1}) = 0$. This implies that the separation dimension is at most $\lfloor 2/\eta + 1 \rfloor$. Since we assumed that s is odd, the separation dimension is at most $\lfloor 2/\eta + 2 \rfloor$. □

An example of a function class with an infinite separation dimension is \mathcal{F}_{ind} , where for every $z \in \mathcal{X}$ we have an indicator function $f_z \in \mathcal{F}_{\text{ind}}$ (i.e., $f_z(x) = 1$ for $x = z$ and otherwise $f_z(x) = 0$), and \mathcal{X} is infinite. Since for any set of s distinct points $x_1, \dots, x_s \in \mathcal{X}$ the indicator functions f_{x_1}, \dots, f_{x_s} are a 1-separation, this implies that \mathcal{F}_{ind} has an infinite separation dimension.

Separation dimension—number of consistent functions Given a set of points $S = \{x_1, \dots, x_s\}$ let \mathcal{F}^S be the function class \mathcal{F} restricted to the set S . We would like to bound the number of consistent functions \mathcal{F}^S as a function of $s = |S|$ and the separation dimension of the function class \mathcal{F} .

It is obvious that if a function class \mathcal{F} has a separation dimension of d then it has a VC dimension (Vapnik and Chervonenkis 1971) of at most d . Therefore, using Sauer Lemma (Sauer 1972), we can bound $|\mathcal{F}^S|$. We can also show a function class for which this bound is almost tight also for separation dimension.

Lemma 1 *Let \mathcal{F} be a function class with separation dimension d , then $|\mathcal{F}^S| \leq \sum_{i=0}^d \binom{|S|}{i}$. In addition, there is a function class \mathcal{F} of separation dimension d such that $|\mathcal{F}^S| \geq (|S|/d)^d$.*

Proof The upper bound follows from Sauer Lemma (Sauer 1972). We show that the bound is close to being tight for the separation dimension. Consider s points on $[0, 1]$ where $x_1 < \dots < x_s$ (we assume that s is divisible by d). Split the s points to d sets where the i -th set J_i includes the points $\{x_{(i-1)s/d+1}, \dots, x_{is/d}\}$. For the set J_i of s/d points, define the function class G_i which are the $s/d + 1$ function induced by \mathcal{F}_{pre} (namely a prefix of the points is labeled 1 while the suffix is labeled 0). The function class \mathcal{F} will be composed from functions f_{j_1, \dots, j_d} which for the inputs in J_i selects function $g_{j_i} \in G$. Note that $|\mathcal{F}| = (s/d + 1)^d$. In addition any $d + 1$ points have two points $x, x' \in J_i$, for some i . Wlog, $x < x'$, then our function class does not have any function f such that $f(x) = 0$ and $f(x') = 1$ (since they are both in the same J_i). □

4.2 Separation dimension: lower bound

In this section we give a lower bound based on the separation dimension.

Theorem 3 *Let \mathcal{F} be a function class with separation dimension d , then its expected active sample complexity is $\Omega(\min\{d, s\})$, where s is the transductive sample size. In addition, if \mathcal{F} has an infinite separation dimension then its expected active sample complexity is $\Omega(s)$.*

Proof Assume that \mathcal{F} has 1-separation dimension of d (the other case is identical). Then there are $k = \min\{d, s\}$ inputs $x_1, \dots, x_k \in \mathcal{X}$ such that for any i there is an $f_i \in \mathcal{F}$ that $f_i(x_i) = 1$ and $f_i(x_j) = 0$, for $j \neq i$. (When $k < s$, we extend the set of k points to s points by duplicating point x_1 for $s - k$ times.)

Assume that we select the target function f to be f_i with probability $1/k$. Clearly any output identification algorithm would have to make, in expectation, at least $\Omega(k)$ queries before hitting the input which is labeled 1.

If \mathcal{F} has an infinite separation dimension, then for every s , there are s inputs $x_1, \dots, x_s \in \mathcal{X}$ and functions $f_1, \dots, f_s \in \mathcal{F}$ such that $f_i(x_i) = 1$ and $f_i(x_j) = 0$, for $j \neq i$. Again, this implies that the expected number of queries is $\Omega(s)$. □

4.3 Separation dimension: upper bound

In this section we derive an upper bound on the sample complexity based on the separation dimension. Assume that the first point we sampled has label 0, and therefore the output identification task reduces to finding an input with label 1. (Hence, we will also concentrate on the 1-separation dimension.)

We start with a few notations. Let $\text{sup}(f, S) = \{x_i \in S : f(x_i) = 1\}$, i.e., the set of points in S for which f is 1. Given a set of functions \mathcal{F} , let $\text{deg}(x_i, \mathcal{F}^S) = \{f \in \mathcal{F}^S : f(x_i) = 1\}$, i.e., the set of functions which classify x_i as 1. We partition \mathcal{F}^S according to the size of $\text{sup}(f, S)$, and define $\mathcal{F}_k^S = \{f \in \mathcal{F}^S : k \leq |\text{sup}(f, S)| < 2k\}$. Let $\ell_k = |\mathcal{F}_k^S|$.

Lemma 2 *Let \mathcal{F} be a function class with 1-separation dimension d . Given an unlabeled sample S , for every $k \geq 1$, there exists an input $x_i \in S$ such that $|\text{deg}(x_i, \mathcal{F}_k^S)| \geq \lfloor \ell_k / 8d \rfloor$.*

Proof For contradiction, assume that for some $k \geq 1$ no such $x_i \in S$ exists. We will show that the 1-separation dimension is at least $d + 1$, which would be a contradiction.

Since we assume that no such x_i exists for k , then for every $x_i \in S$ we have $|\text{deg}(x_i, \mathcal{F}_k^S)| < \ell_k/8d$. By definition of \mathcal{F}_k^S , for every $f \in \mathcal{F}_k^S$ we have $|\text{sup}(f, S)| \in [k, 2k]$. Consider the set of pairs Z that includes all the pairs (f, x) such that $x \in S, f \in \mathcal{F}_k^S$ and $f(x) = 1$. There are at least $k\ell_k$ such pairs in Z .

We would like to find a subset of pairs $(f_1, x_1) \dots, (f_{d+1}, x_{d+1})$ in Z such that for any $i \neq j$ we have $f_j(x_i) = 0$. Recall that by the definition of the pairs in Z we have $f_i(x_i) = 1$. Therefore, such a subset would imply that the separation dimension of \mathcal{F} is at least $d + 1$.

For the first pair we pick any (f, x) in Z . We would like to delete some of the pairs in Z such that any remaining pair (h, z) has the property that $f(z) = 0$ and $h(x) = 0$. This would guarantee that the subset that we select would have the required property.

Formally, we delete from Z both the set $\{(g, y) \mid g \in \mathcal{F}_k^S, y \in \text{sup}(f, S)\}$ and the set $\{(g, y) \mid g \in \mathcal{F}_k^S, g \in \text{deg}(x, \mathcal{F}_k^S)\}$. The deletion of the first set guarantees that any remaining pair (h, z) would have $f(z) = 0$ while the deletion of the second set guarantees that for any remaining pair (h, z) we have $h(x) = 0$. The size of the first set is at most $\sum_{y \in \text{sup}(f, S)} \text{deg}(y, \mathcal{F}_k^S) \leq k\ell_k/4d$ while the size of the second set is at most $\sum_{h \in \text{deg}(x, \mathcal{F}_k^S)} \text{sup}(h, S) \leq k\ell_k/4d$. This implies that we delete at most $k\ell_k/2d$ pairs.

By iteratively selecting a pair from Z and deleting from Z the two related sets, this implies that we can select $2d \geq d + 1$ such pairs. This is a contradiction to the assumption that separation dimension is d . □

We can now use Lemma 2 to derive an output identification algorithm in the transductive setting.

Theorem 4 *Let \mathcal{F} be a function class of separation dimension d . For any unlabeled sample S , it can be output identified in active sample complexity of $O(d \log |S| \log |\mathcal{F}^S|)$.*

Proof Initially we query an arbitrary point in S . W.l.o.g., assume that the first point has a label 0. This implies that we need to search for a point in S with a label of 1. Note that the 1-separation dimension of \mathcal{F} is at most d .

We run the algorithm in rounds, where in each round we select at most $\log |S|$ inputs (one for to each $\mathcal{F}_{2^i}^S$, where $i \in [0, \log |S|]$). In each round for each $\mathcal{F}_{2^i}^S$ we select the input x which maximizes $\text{deg}(x, \mathcal{F}_{2^i}^S)$. By Lemma 2, since \mathcal{F} has 1-separation dimension of at most d , there exists an input x_i such that $|\text{deg}(x, \mathcal{F}_{2^i}^S)| \geq \ell_k/8d$. Therefore, in each round, the number of possible target functions in each $\mathcal{F}_{2^i}^S$ shrink by a factor of $(1 - 1/8d)$. After at most $O(d \log |\mathcal{F}^S|)$ rounds we will either: (1) find an $x \in S$ with label 1, i.e., $f(x) = 1$, or, (2) the only remaining consistent function in \mathcal{F}^S is the all zero function, i.e., there are no points in S with a label of 1. □

Since, by Lemma 1, a d separation dimension implies that $|\mathcal{F}^S| = O(|S|^d)$, we have,

Corollary 1 *If \mathcal{F} has separation dimension d then any unlabeled sample S can be output identified with an active sample complexity of $O(d^2 \log^2 |S|)$.*

5 Uniform distribution

In this section we discuss active sample complexity for specific distributions. We will focus attention on the case where the input is uniformly distributed over the unit square, i.e.,

$\mathcal{X} = [0, 1]^2$. In this case we show that many natural geometric concepts, which for a general distribution require a sample of size $\Omega(1/\epsilon)$, are efficiently output identified, with respect to the uniform distribution. (For simplicity of the presentation, we will focus our discussion on geometric concepts in the plane. However, one can consider concepts in higher dimensions. In such a case, our basic methodology would still apply, but would require computational geometry results analogous to that of Theorem 5.)

Generic convex-hull algorithm Our algorithm would be a generic consistency algorithm (Sect. 3.1) for the case where the domain of each output value is convex. Initially we have $S_0 = C_0 = \emptyset$. At time t we sample the induced distribution example oracle with $\mathcal{X} - C_t$ and receive an example $\langle x_t, k_t \rangle$. For each output value k let S_t^k be the set of points sampled with output value k . Let C_t^k be the convex hull of the points in S_t^k , and let C_t be the union of all those sets. We terminate when for every output value k we have $S_t^k \neq \emptyset$.

Review from computational geometry Classical results in computational geometry regarding uniform sampling from the plane, relate the area of the convex-hull to the total area of the convex domain from which the points are sampled at uniform. For a set of points Z let $\text{ConvexHull}(Z)$ be their convex hull, and for a body Y let $\text{area}(Y)$ be its area. The following theorem summarizes the related results (see, cf., Preparata and Shamos 1985).

Theorem 5 (Rényi and Sulamke 1963; Efron 1965) *Let G be an r -gon in $[0, 1]^2$, and S_n be a sample of size n sampled uniformly from G . Then, $E[\text{area}(G - \text{ConvexHull}(S_n))] = \Theta(\frac{r \ln n}{n})$.*

Triangles in the plane Let $\mathcal{X} = [0, 1]^2$ be the unit square. Consider the case where there are m output values such that the domain of each output value is a triangle and their union is the unit square.

It would be more beneficial in this case to consider an alternative way of sampling. Assume that each time we access the induced distribution example oracle with $\mathcal{X} - C_t$, the oracle samples from \mathcal{X} until it hits a point $x \notin C_t$. In our analysis let us consider also the extra points that the oracle samples (but we will give them zero weight). Assume that in this process the total number of samples the oracle makes until it terminates, is T . (Recall that the generic convex hull algorithm terminates when each of the m output values is sampled at least once.) Let T_i be the number of times the oracle samples output value i (out of the T samples). Let X_j^i be a random variable which is 1 if the j -th point with output value i is outside the convex hull of the previous $j - 1$ points, and 0 otherwise. This implies that the expected number of samples of the generic convex hull algorithm is $E_T E_{T_1, \dots, T_m} E[\sum_{i=1}^m \sum_{j=1}^{T_i} X_j^i]$.¹

From Theorem 5, applied to triangles (i.e., $r = 3$), we can deduce that, $E[X_j^i] = O(\frac{\ln j}{j})$. Therefore, $\sum_{j=1}^{T_i} E[X_j^i] = \sum_{j=1}^{T_i} O(\frac{\ln j}{j}) \leq \alpha \log^2 T_i$, for some constant $\alpha > 0$. Summing over all possible output values we have

$$E_T E_{T_1, \dots, T_m} E \left[\sum_{i=1}^m \sum_{j=1}^{T_i} X_j^i \right] \leq E_T \left[\alpha \sum_{i=1}^m E_{T_i} [\log^2 T_i] \right] \leq E_T \left[\alpha m \log^2 \frac{T}{m} \right] \leq \alpha m \log^2 \frac{1}{\epsilon}$$

¹This follows since an equivalent way of active sampling is to sample points from the distribution and request a label (and charge) only for points not in C_t .

where the second inequality follows by maximizing over the T_i 's subject to $\sum_{i=1}^m T_i = T$, and the last inequality uses the fact that $E[T] \leq m/\epsilon$ and the concavity of the logarithm function. This implies the following theorem,

Theorem 6 *Let \mathcal{F}_Δ be a function class such that every $f \in \mathcal{F}_\Delta$ partitions the unit square to m triangles each of area at least ϵ . Then for the uniform distribution U , the expected active sample complexity of the generic convex hull algorithm is $O(m \log^2 1/\epsilon)$.*

Lines in the plane Consider the case where there are k lines in the plane that partition the unit square. Each of the m output values is one of the cells created by the intersection of the lines. The analysis would be the same as in the case of the triangles in the plane, where the only difference is that now we have $r = k$ (for triangles $r = 3$). Again, one can deduce from Theorem 5 that, $E[X_j^i] = O(\frac{k \ln j}{j})$. Therefore, $\sum_{j=1}^{T_i} E[X_j^i] = \sum_{j=1}^{T_i} O(\frac{k \ln j}{j}) \leq \alpha k \log^2 T_i \leq \alpha m \log^2 T_i$, for some constant $\alpha > 0$, where the last inequality follows from the fact that k lines in the plane create at least k cells, which implies that $m \geq k$. Summing over all possible output values we have

$$\begin{aligned}
 & E_T E_{T_1, \dots, T_m} E \left[\sum_{i=1}^m \sum_{j=1}^{T_i} X_j^i \right] \\
 & \leq E_T \left[\alpha m \sum_{i=1}^m E_{T_i} [\log^2 T_i] \right] \leq E_T \left[\alpha m^2 \log^2 \frac{T}{m} \right] \leq \alpha m^2 \log^2 \frac{1}{\epsilon}.
 \end{aligned}$$

This implies the following theorem,

Theorem 7 *Let $\mathcal{F}_{k\text{-line}}$ be a function class such that every $f \in \mathcal{F}_{k\text{-line}}$ is represented by at most k lines in the plane. If each of the m cells has area at least ϵ then for the uniform distribution U the expected active sample complexity of the generic convex hull algorithm is $O(m^2 \log^2 1/\epsilon)$.*

Note that \mathcal{F}_{ls} , for which we showed a lower bound of $\Omega(1/\epsilon)$ with respect to an arbitrary distribution, is simply $\mathcal{F}_{1\text{-line}}$ and $m = 2$. This shows an exponential gap between the uniform distribution and an arbitrary one.

6 Monomials in a hypercube

In this section we will concentrate on the Boolean cube and consider the case that each output value is a monomial. Formally, the domain \mathcal{X} is $\{0, 1\}^n$. The function class \mathcal{F}_{mon} includes functions f of the following type: For every output k there is a monomial M_k^f such that $f(x) = k$ iff $M_k^f(x) = 1$. (Note that for any $f \in \mathcal{F}_{\text{mon}}$ the monomials M_i^f are a partition of the hypercube.)²

For any i let x^i be the input $x \in \{0, 1\}^n$ with the i -th bit flipped. Note that if $M_k(x) = 1$ then $M_k(x^i) = 0$ iff M_k depends on attribute x_i . This would be the basic property that will allow us to efficiently output identify \mathcal{F}_{mon} .

²When clear from the context we drop the superscript f .

Monomial output identification algorithm The algorithm performs an exact identification using (essentially) membership queries. Let KN be the set of known output values and CH be a set of inputs we need to “check”. Initially both $KN = \emptyset$ and $CH = \emptyset$. In the first phase, we sample any x , get its value $f(x) = k$ and add k to KN and x to CH . In every phase we take an input x from CH , and for every i such that $f(x^i) \notin KN$ we add the output value $f(x^i)$ to KN , add the input x^i to CH , and continue to the next input from CH . We terminate either when we have processed all the values in CH or have already recovered m output values (i.e., $|KN| = m$). Since at most m inputs enter the set CH , and each is used at most n times, we derive the following.

Theorem 8 *For any distribution which has non-zero probability for every example $\{0, 1\}^n$, the class \mathcal{F}_{mon} can be output identified with active sample complexity of mn .*

Note that the monomial output identification algorithm does not need as an input parameter the value of m . Also, note that it can output a complete model of the target function f in addition to the m output values. This is in contrast to all our previous results where we never constructed an accurate model of the target function.

7 Conclusion and open problems

In this work we introduced the problem of output identification, which applies to functions with multiple output values. We used active sampling to identify an example for each of the possible output values. For simple geometric concepts we provided efficient active sampling algorithms. In addition, under the uniform distribution, we have shown that in many cases one can achieve much better active sampling bounds than with respect to an arbitrary distribution.

We derived a characterization for the case of binary output value in a transductive setting. This result has recently been extended in a few interesting directions (see Srebro 2006; Gilad-Bachrach 2006). An interesting additional extension would be to extend the results from the transductive setting to a distributional setting.

A more challenging (and abstract) question addresses the definition of dimension. The VC dimension considers all possible Boolean functions on d points, our separation dimension considers only the singleton functions. Is there an intuitive computational learning meaning of other similar definitions of a dimension, for example considering all functions with at most k ones.

Finally, we believe that the output identification problem has many applications in practice, and we look forward to seeing experimental works in this direction.

Appendix: Axis parallel boxes: analysis

The analysis of the axis parallel boxes would be similar to the analysis of intervals on a line. At time t , we have a set KN_t of output values which we have already sampled and UKN_t which are output values we have not sampled. For each output value $k \in KN_t$ we define $2d$ boxes which extend the current C_t^k to A^k , in each of the d dimensions both above and below. (Note that the boxes overlap unlike in the intervals case.)

Formally, let $C_t^k = ([c_-^1, c_+^1], \dots, [c_-^d, c_+^d])$ and $A^k = ([a_-^1, a_+^1], \dots, [a_-^d, a_+^d])$. Since $C_t^k \subset A^k$, then $a_-^i \leq c_-^i \leq c_+^i \leq a_+^i$. The $2d$ boxes are

$$B_-^{k,i} = ([a_-^1, a_+^1], \dots, [a_-^{i-1}, a_+^{i-1}], [a_-^i, c_-^i], [a_-^{i+1}, a_+^{i+1}], \dots, [a_-^d, a_+^d]).$$

and

$$B_+^{k,i} = ([a_-^1, a_+^1], \dots, [a_-^{i-1}, a_+^{i-1}], [c_+^i, a_+^i], [a_-^{i+1}, a_+^{i+1}], \dots, [a_-^d, a_+^d]).$$

Let, $\gamma_t = \sum_{k \in KN_t} \sum_{i=1}^d D(B_+^{k,i}) + D(B_-^{k,i})$, $\beta_t = \sum_{k \in UNKN_t} D(A^k)$, and $\alpha_t = D(\mathcal{X} - C_t)$ where $C_t = \bigcup_{k \in KN_t} C_t^k$. The potential function would be

$$\Phi_t = \gamma_t + 2d\beta_t.$$

For intuition, note that when we sample an output value $k \in UNKN_t$ for the first time, we decrease β_t by $D(A^k)$ and increase γ_t by at most $dD(A^k)$, so the net reduction in the potential is at least $dD(A^k)$. Let \mathcal{H}_t include all the history of the execution of the algorithm until and including time t .

Our analysis of the generic consistency algorithm for axis parallel boxes is very much in the spirit of the analysis in the case of intervals. We will show that the expected potential Φ_t decreases by a certain factor each time step. Specifically, we show that

$$\begin{aligned} E[\Phi_t - \Phi_{t+1} | \mathcal{H}_t] &\geq \frac{\beta_t}{\alpha_t} \sum_{k \in UNKN_t} \frac{D(A^k)}{\beta_t} dD(A^k) \\ &\quad + \frac{\alpha_t - \beta_t}{\alpha_t} \sum_{k \in KN_t, i \in [1, d], b \in \{+, -\}} \frac{D(B_b^{k,i})}{\alpha_t - \beta_t} \cdot \frac{D(B_b^{k,i})}{2}. \end{aligned}$$

The first part follows from the fact that with probability $\frac{\beta_t}{\alpha_t}$ we sample an example with output value in $UNKN_t$. Given that we sample such an example, the probability that the output value is k is $\frac{D(A^k)}{\beta_t}$. Given that we sample from A^k we reduce β_t by $D(A^k)$ and increase γ_t by at most $dD(A^k)$, so the net reduction in the potential is at least $dD(A^k)$.

The second part follows from the fact that with probability $\frac{\alpha_t - \beta_t}{\alpha_t}$ we sample an example with output value in KN_t . Given that we sample such an example, the probability that the example is in box $B_b^{k,i}$ is $\frac{D(B_b^{k,i})}{\alpha_t - \beta_t}$. (Note that a point can be in more than one box. We used here the linearity of expectations, to be able to consider each box separately.) Given that we sample from $B_b^{k,i}$ the expected reduction is at least $D(B_b^{k,i})/2$. (The best way to see this is to project the distribution to dimension i , and there the reduction is obvious.) Therefore,

$$E[\Phi_t - \Phi_{t+1} | \mathcal{H}_t] \geq \frac{1}{\alpha_t} \left[\sum_{k \in UNKN_t} dD(A^k)^2 + \sum_{k \in KN_t, i \in [1, d], b \in \{+, -\}} \frac{1}{2} D(B_b^{k,i})^2 \right].$$

Again, we use the general inequality $\sum_{i=1}^n X_i^2 \geq (1/n)(\sum_{i=1}^n X_i)^2$, for each summation separately,

$$\begin{aligned} &\sum_{k \in UNKN_t} dD(A^k)^2 + \sum_{k \in KN_t, i \in [1, d], b \in \{+, -\}} \frac{1}{2} D(B_b^{k,i})^2 \\ &\geq \frac{1}{m} d\beta_t^2 + \frac{1}{4dm} \gamma_t^2 = \frac{1}{4dm} (4d^2\beta_t^2 + \gamma_t^2). \end{aligned}$$

Therefore,

$$E[\Phi_t - \Phi_{t+1} | \mathcal{H}_t] \geq \frac{1}{4dm} \frac{4d^2 \beta_t^2 + \gamma_t^2}{\alpha_t}.$$

Since each point can be counted at most d times in γ_t , we have that,

$$\beta_t + \gamma_t \geq \alpha_t = D(\mathcal{X} - C_t) \geq \beta_t + \gamma_t/d.$$

This implies that

$$E[\Phi_t - \Phi_{t+1} | \mathcal{H}_t] \geq \frac{1}{4dm} \frac{4d^2 \beta_t^2 + \gamma_t^2}{\beta_t + \gamma_t} \geq \frac{1}{8dm} [2d\beta_t + \gamma_t] = \frac{1}{8dm} \Phi_t,$$

where in the second inequality we use that $X^2 + Y^2 \geq \frac{1}{2}(X + Y)^2$. By averaging over \mathcal{H}_t we have that

$$E[\Phi_{t+1}] \leq \left(1 - \frac{1}{8dm}\right) E[\Phi_t] \leq \left(1 - \frac{1}{8dm}\right)^t \Phi_1.$$

Initially we have $\gamma_1 = 0$ and $\beta_1 = 1$, therefore the initial potential is $\Phi_1 = 2d$. After $t = O(dm \log(d/\epsilon))$ samples the expected potential is less than $\epsilon/2$. Therefore with probability at least $1/2$ the potential is less than ϵ . If the potential is less than ϵ , this implies that we hit every output value (since we assume that each output value has probability of at least ϵ). Hence, we establish Theorem 2.

References

- Dasgupta, S. (2004). Analysis of a greedy active learning strategy. In *Advances in neural information processing systems (NIPS)*.
- Dasgupta, S. (2005). Coarse sample complexity bounds for active learning. In *Advances in neural information processing systems (NIPS)*.
- Dasgupta, S., Kalai, A., & Monteleoni, C. (2005). Analysis of perceptron-based active learning. In *Eighteenth annual conference on learning theory (COLT)*.
- Edelstein, O., Farchi, E., Nir, Y., Ratzaby, G., & Ur, S. (2002). Multithreaded java program test generation. *IBM Systems Journal*, 41(3), 111–125.
- Efron, B. (1965). The convex hull of a random set of points. *Biometrika*, 52, 331–343.
- Feller, W. (1966). *An introduction to probability theory and its applications* (Vol. I). New York: Wiley.
- Fine, S., & Ziv, A. (2003). Coverage directed test generation for functional verification using Bayesian networks. In *Proceedings of the 40th design automation conference* (pp. 286–291), June 2003.
- Fine, S., Gilad-Bachrach, R., & Shamir, E. (2002). Query by committee, linear separation and random walks. *Theoretical Computer Science*, 284(1). (A preliminary version appeared in EuroColt 1999).
- Freund, Y., Seung, H., Shamir, E., & Tishby, N. (1997). Selective sampling using the query by committee algorithm. *Machine Learning*, 28(2/3), 133–168.
- Gilad-Bachrach, R. (2006). Personal communication.
- Kulkarni, S. R., Mitter, S. K., & Tsitsiklis, J. N. (1993). Active learning using arbitrary binary valued queries. *Machine Learning*, 11, 23–35.
- Liere, R., & Tadepalli, P. (1997). Active learning with committees for text categorization. In *AAAI-97*.
- Linial, N., Luby, M., Saks, M., & Zuckerman, D. (1997). Efficient construction of a small hitting set for combinatorial rectangles in high dimension. *Combinatorica*, 17(2), 215–234. (A preliminary version appeared in STOC 1993.)
- Piziali, A. (2004). *Functional verification coverage measurement and analysis*. Berlin: Springer.
- Preparata, F. P., & Shamos, M.I. (1985). *Computational geometry: an introduction*. Berlin: Springer.
- Rényi, A., & Sulamke, R. (1963). Über die konvexe hülle von n zufällig gewählten punkten. *Zeitschrift für Wahrscheinlichkeitstheorie*, 2, 75–84.

- Sauer, N. (1972). On the density of family of sets. *Journal of Combinatorial Theory, Series A*, 13, 145–147.
- Seung, H. S., Oppen, M., & Sompolinsky, H. (1992). Query by committee. In *Proceedings of the fifth workshop on computational learning theory* (pp. 287–294). San Mateo: Kaufman.
- Srebro, N. (2006). Personal communication.
- Ur, S., & Yadin, Y. (1999). Micro-architecture coverage directed generation of test programs. In *Proceedings of the 36th design automation conference* (pp. 175–180), June 1999.
- Vapnik, V. N., & Chervonenkis, A. Y. (1971). On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, XVI(2), 264–280.
- Wile, B., Goss, J. C., & Roesner, W. (2005). *Comprehensive functional verification the complete industry cycle*. Amsterdam: Elsevier.