

# Interactive Algorithms: from Pool to Stream

**Sivan Sabato**

**Tom Hess**

*Department of Computer Science*

*Ben-Gurion University of the Negev, Beer Sheva 8410501, Israel.*

SABATOS@CS.BGU.AC.IL

TOMHE@POST.BGU.AC.IL

## Abstract

We consider interactive algorithms in the pool-based setting, and in the stream-based setting. Interactive algorithms observe suggested elements (representing actions or queries), and interactively select some of them and receive responses. Pool-based algorithms can select elements at any order, while stream-based algorithms observe elements in sequence, and can only select elements immediately after observing them. We assume that the suggested elements are generated independently from some source distribution, and ask what is the stream size required for emulating a pool algorithm with a given pool size. We provide algorithms and matching lower bounds for general pool algorithms, and for utility-based pool algorithms. We further show that a maximal gap between the two settings exists also in the special case of active learning for binary classification.

**Keywords:** Interactive algorithms, active learning, pool-based, stream-based

## 1. Introduction

*Interactive algorithms* are algorithms which are presented with input in the form of suggested elements (representing actions or queries), and iteratively select elements, getting a response for each selected element. The reward of the algorithm, which is application-specific, is a function of the final set of selected elements along with their responses. Interactive algorithms are used in many application domains, including, for instance, active learning (McCallum and Nigam, 1998), interactive sensor placement (Golovin and Krause, 2011), summarization (Singla et al., 2016) and promotion in social networks (Guillory and Bilmes, 2010). As a specific motivating example, consider an application in which elements represent web users, and the algorithm should select up to  $q$  users to present with a free promotional item. For each selected user, the response is the observed behavior of the user after having received the promotion, such as the next link that the user clicked on. The final reward of the algorithm depends on the total amount of promotional impact it obtained, as measured by some function of the set of selected users and their observed responses. Note that the algorithm can use responses from previous selected users when deciding on the next user to select.

We consider two interaction settings for interactive algorithms: The *pool-based* setting and the *stream-based* setting. In the pool-based setting, the entire set of suggested elements is provided in advance to the algorithm, which can then select any of the elements at any order. For instance, in the web promotion example, there might be a set of users who use the website for an extended period of time, and any of them can be approached with a promotion. In the stream-based setting, elements are presented to the algorithm in sequence, and the algorithm must decide immediately after observing an element, whether to select it or not. In the web promotion example, this is consistent with

a setting where users access the website for single-page sessions, and so any promotion must be decided on immediately when the user is observed.

The stream-based setting is in general weaker than the pool-based setting. Nonetheless, it is important and useful: In many real-life scenarios, it is not possible to postpone selection of elements, for instance due to storage and retrieval constraints, or because of timing constraints. This is especially pertinent when the data stream is real-time in nature, such as in streaming document classification (Bouguelia et al., 2013), in spam filtering (Chu et al., 2011), in web streams such as Twitter (Smailović et al., 2014), in video surveillance (Loy et al., 2012) and with active sensors (Krishnamurthy, 2002).

In this work, our goal is to study the relationship between these two important settings. Both of these settings have been widely studied in many contexts. In active learning, both settings have been studied in classic works (Cohn et al., 1994; Lewis and Gale, 1994). Works that address mainly the stream-based setting include, for instance, Balcan et al. (2009); Hanneke (2011); Dasgupta (2012); Balcan and Long (2013); Sabato and Munos (2014). Some theoretical results hold equally for the stream-based and the pool-based settings (e.g., Balcan and Long, 2013; Hanneke and Yang, 2015).

Several near-optimal algorithms have been developed for the pool-based setting (Dasgupta, 2005; Golovin and Krause, 2011; Golovin et al., 2010b; Hanneke, 2007; Sabato et al., 2013; Gonen et al., 2013; Cuong et al., 2014). The pool-based setting is also heavily studied in various active learning applications (e.g., Tong and Koller, 2002; Tong and Chang, 2001; Mitra et al., 2004; Gosselin and Cord, 2008; Cebon and Berthold, 2009; Guo et al., 2013). General interactive algorithms have also been studied in both a pool-based setting (e.g., Golovin and Krause, 2011; Guillory and Bilmes, 2010; Deshpande et al., 2014) and in stream-based settings (e.g., Demaine et al., 2014; Arlotto et al., 2014; Streeter and Golovin, 2009; Golovin et al., 2010a). Note that unlike some works on interactive algorithms, in our stream-based setting, the only direct restriction is on the timing of selecting elements. We do not place restrictions on storage space or any other resources.

To study the relationship between the pool-based setting and the stream-based setting, we assume that in both settings the suggested elements, along with their hidden responses, are drawn i.i.d. from some unknown source distribution. We then ask under what conditions, and at what cost, can a stream-based algorithm obtain the same output distribution as a given black-box pool algorithm. Such an exact emulation is advantageous, as it allows direct application of methods and results developed for the pool-based setting, in the stream-based setting. Especially, if a pool-based algorithm succeeds in practice, but its analysis is unknown or limited, exact emulation guarantees that success is transferred to the stream setting as well.

For discrete source distributions, any pool-based algorithm can be emulated in a stream-based setting, simply by waiting long enough, until the desired element shows up again. The challenge for stream-based interactive algorithms is thus to achieve the same output distribution as a pool-based algorithm, while observing as few suggested elements as possible. Clearly, there are many cases in which it is desired to require less suggested elements: this could result in saving of resources such as time, money, and communication. In active learning as well, while examples are usually assumed cheap, they are not usually completely free in all respects.

We study emulation of pool-based algorithm in two settings. First, we consider the fully general case. We provide a stream algorithm that can emulate any given black-box pool algorithm, and uses a uniformly bounded expected number of observed elements. The bound on the expected number of observed elements is exponential in the number of selected elements. We further prove a lower bound which indicates that this exponential dependence is necessary. Second, we consider *utility-*

based interactive algorithm for the pool setting. We provide a stream algorithm that emulates such pool algorithms, using repeated careful solutions of the well known ‘‘Secretary Problem’’ (Dykin, 1963; Gilbert and Mosteller, 1966; Ferguson, 1989). The expected number of observed elements for this algorithm is only linear in the number of selected elements. In this case too we prove a matching lower bound.

Finally, we show a lower bound that applies to active learning for binary classification. We conclude that even in this well-studied setting, there are cases in which there exists a significant gap between the best pool-based algorithm and the best stream-based algorithm. This result generalizes a previous observation of Gonen et al. (2013) on the sub-optimality of CAL (Cohn et al., 1994), the classical stream-based active learning algorithm, compared to pool algorithms.

This paper is structured as follows: In Section 2 formal definitions and notations are provided. Section 3 discusses natural but suboptimal solutions. Section 4 provides an algorithm and a lower bound for the general case, and Section 5 addresses the case of utility-based pool algorithms. In Section 6 we provide a lower bound that holds for active learning for binary classification. We conclude in Section 7. Some of the proofs are provided in Appendix A.

## 2. Definitions

For a predicate  $p$ , denote by  $\mathbb{I}[p]$  the indicator function which is 1 if  $p$  holds and zero otherwise. For an integer  $k$ , denote  $[k] := \{1, \dots, k\}$ . For a sequence  $S$ ,  $S(i)$  is the  $i$ ’th member of the sequence. Denote concatenation of sequences by  $\circ$ . For  $A, B$  which are both sequences, or one is a set and one a sequence, we use  $A =_{\pi} B$  and  $A \subseteq_{\pi} B$  to denote equality or inclusion on the unordered sets of elements in  $B$  and in  $A$ .

Let  $\mathcal{X}$  be a measurable domain of elements, and let  $\mathcal{Y}$  be a measurable domain of responses. A pool-based (or just pool) interactive algorithm  $\mathcal{A}_p$  receives as input an integer  $q \leq m$ , and a pool of elements  $(x_1, \dots, x_m) \in \mathcal{X}^m$ . We assume that for each  $x_i$  there is a response  $y_i \in \mathcal{Y}$ , which is initially hidden from  $\mathcal{A}_p$ . Denote  $S = ((x_i, y_i))_{i \in [m]}$ . For a given  $S$ ,  $S_X$  denotes the pool  $(x_1, \dots, x_m)$ . At each round,  $\mathcal{A}_p$  selects one of the elements  $i_t$  that have not been selected yet, and receives its response  $y_{i_t}$ . After  $q$  rounds,  $\mathcal{A}_p$  terminates. Its output is the set  $\{(x_{i_1}, y_{i_1}), \dots, (x_{i_q}, y_{i_q})\}$ . For a pool algorithm  $\mathcal{A}_p$ , denote by  $\text{sel}_p(S, t)$  the element that  $\mathcal{A}_p$  selects at round  $t$ , if  $S$  is the pool it interacts with.  $\text{sel}_p(S, t)$ , which can be random, can depend on  $S_X$  and on  $y_{i_k}$  for  $k < t$ . Denote by  $\text{sel}_p(S, [t])$  the sequence of elements selected by  $\mathcal{A}_p$  in the first  $t$  rounds.  $\text{pairs}_p(S, t)$  and  $\text{pairs}_p(S, [t])$  similarly denote the selected elements along with their responses. The final output of  $\mathcal{A}_p$  is the set of pairs in the sequence  $\text{pairs}_p(S, [q])$ . We assume that  $S \mapsto \text{pairs}_p(S, [q])$  is measurable.

We assume that the pool algorithm is permutation invariant. That is, for any  $S, S' \subseteq (\mathcal{X} \times \mathcal{Y})^m$ , if  $S'$  is a permutation of  $S$  then  $\text{sel}_p(S, [q]) = \text{sel}_p(S', [q])$ , or if  $\mathcal{A}_p$  is randomized then the output distributions are the same. When the pool  $S$  is drawn i.i.d. this does not lose generality.

A stream-based (or just stream) interactive algorithm  $\mathcal{A}_s$  receives as input an integer  $q$ . We assume an infinite stream  $S \subseteq (\mathcal{X} \times \mathcal{Y})^{\infty}$ , where  $S(t) = (x_t, y_t)$ . At iteration  $t$ ,  $\mathcal{A}_s$  observes  $x_t$ , and may select one of the following actions:

- Do nothing
- Select  $x_t$  and observe  $y_t$
- Terminate.

At termination, the algorithm outputs a subset of size  $q$  of the set of pairs  $(x_t, y_t)$  it observed. Denote by  $\text{sel}_s(S, t)$  the  $t$ 'th element that  $\mathcal{A}_s$  selects and is also in the output set. Denote by  $\text{sel}_s(S, [t])$  the sequence of first  $t$  elements selects and are also in the output set. Use  $\text{pairs}_s$  to denote the elements along with their responses. The output of  $\mathcal{A}_s$  when interacting with  $S$  is the set of the pairs in the sequence  $\text{pairs}_s(S, [q])$ . We assume  $S \mapsto \text{pairs}_s(S, [q])$  is measurable. The total number of elements selected by  $\mathcal{A}_s$  when interacting with  $S$  (including discarded elements) is denoted  $N_{\text{sel}}(\mathcal{A}_s, S, q)$ . The number of iterations (observed elements) until  $\mathcal{A}_s$  terminates is denoted  $N_{\text{iter}}(\mathcal{A}_s, S, q)$ .

We look for stream algorithms that emulate pool algorithms. We define an equivalence between a stream algorithm and a pool algorithm as follows.

**Definition 1** *Let  $\mathcal{D}$  be a distribution over  $\mathcal{X} \times \mathcal{Y}$  and let  $q$  be an integer. Let  $S \sim \mathcal{D}^m, S' \sim \mathcal{D}^\infty$ . A pool algorithm  $\mathcal{A}_p$  and a stream algorithm  $\mathcal{A}_s$  are  $(q, \mathcal{D})$ -equivalent, if the total variation distance between the distributions of  $\text{pairs}_p(S, [q])$  and  $\text{pairs}_s(S', [q])$  is zero.*

Denote by  $\mathcal{D}_X$  the marginal of  $\mathcal{D}$  on  $\mathcal{X}$ . Below, unless specified otherwise, we assume that the probability under  $\mathcal{D}_X$  of observing any single  $x \in \mathcal{X}$  is zero. This does not lose generality, since if this is not the case,  $\mathcal{D}_X$  can be replaced by the distribution  $\mathcal{D}_X \times \text{Unif}[0, 1]$ , with the interactive algorithms ignoring the second element in the pair.

### 3. Simple equivalent stream algorithms

Let  $\mathcal{A}_p$  be a pool algorithm. For any discrete distribution  $\mathcal{D}$  over  $\mathcal{X} \times \mathcal{Y}$ , and any  $q$ , it is easy to define a stream algorithm which is  $(q, \mathcal{D})$ -equivalent to  $\mathcal{A}_p$ . Let “ $\star$ ” be some value not in  $\mathcal{Y}$ , and define  $\mathcal{A}_{\text{wait}}$  as in Alg. 1.

---

**Algorithm 1** Algorithm  $\mathcal{A}_{\text{wait}}$

---

- 1: In the first  $m$  iterations, observe  $x_1, \dots, x_m$  and do nothing.
  - 2:  $S \leftarrow ((x_1, \star), \dots, (x_m, \star))$
  - 3:  $j \leftarrow 1$
  - 4: **repeat**
  - 5:   In iteration  $t$ , observe element  $x_t$
  - 6:   **if**  $x_t = \text{sel}_p(S, j)$  **then**
  - 7:     Select  $x_t$  and observe  $y_t$
  - 8:      $S(i) \leftarrow (x_t, y_t)$
  - 9:      $j \leftarrow j + 1$ .
  - 10:   **end if**
  - 11: **until**  $j = q + 1$
  - 12: Return the set of all the pairs  $(x, y)$  in  $S$  with  $y \neq \star$ .
- 

This stream algorithm is  $(q, \mathcal{D})$  equivalent to  $\mathcal{A}_p$  for any discrete distribution  $\mathcal{D}$ , and it has  $N_{\text{sel}}(\mathcal{A}_{\text{wait}}, S', q) = q$  for all  $S' \in (\mathcal{X} \times \mathcal{Y})^\infty$ . However,  $\mathbb{E}_{S' \sim \mathcal{D}^\infty} [N_{\text{iter}}(\mathcal{A}_{\text{wait}}, S', q)]$  is not bounded for the class of discrete distributions.

On the other hand, the stream algorithm  $\mathcal{A}_{\text{nowait}}$  defined in Alg. 2 is also  $(q, \mathcal{D})$  equivalent to  $\mathcal{A}_p$ . We have  $N_{\text{iter}}(\mathcal{A}_{\text{nowait}}, S', q) = m$  for all  $S' \in (\mathcal{X} \times \mathcal{Y})^\infty$ , the same as the pool algorithm. However, also  $N_{\text{sel}}(\mathcal{A}_{\text{nowait}}, S', q) = m > q$ . These two simple approaches demonstrate a possible tradeoff between the number of selected elements and the number of iterations when emulating a pool algorithm.

---

**Algorithm 2** Algorithm  $\mathcal{A}_{\text{nowait}}$

---

**input** Pool size  $m$ , Black-box pool algorithm  $\mathcal{A}_p$ .

- 1: In each iteration  $t \in [m]$ , select  $x_t$  and observe  $y_t$ .
  - 2: Return the pairs in  $\text{pairs}_p(S, q)$ .
- 

**4. An equivalent algorithm with a uniform bound on expected iterations**

We present the stream algorithm  $\mathcal{A}_{\text{gen}}$  (see Alg. 3), which can emulate any pool based algorithm  $\mathcal{A}_p$  using only black-box access to  $\mathcal{A}_p$ . The algorithm emulates a general pool algorithm, by making sure that in each iteration, its probability of selecting an element is identical to the conditional probability of the pool algorithm selecting the same element, conditioned on the history of elements and responses selected and observed so far. This is achieved by repeatedly drawing the remaining part of the pool, and keeping it only if it is consistent with the elements that were already selected. We further can use the partial pool draw only if the element to be selected happens to have been observed last.

---

**Algorithm 3** Algorithm  $\mathcal{A}_{\text{gen}}$

---

**input** Original pool size  $m$ , label budget  $q < m$ , black-box pool algorithm  $\mathcal{A}_p$ .

- 1:  $S_0 \leftarrow ()$
  - 2: **for**  $i = 1 : q$  **do**
  - 3:     **repeat**
  - 4:         Draw  $m - i + 1$  elements, denote them  $\bar{x}_{i,i}, \dots, \bar{x}_{i,m}$ .
  - 5:          $S'_i \leftarrow ((\bar{x}_{i,i}, \star), \dots, (\bar{x}_{i,m}, \star))$ .
  - 6:         **until**  $\text{pairs}_p(S_{i-1} \circ S'_i, [i - 1]) =_{\pi} S_{i-1}$  and  $\text{sel}_p(S_{i-1} \circ S'_i, i) = \bar{x}_{i,m}$ .
  - 7:         Select  $\bar{x}_{i,m}$ , get the response  $\bar{y}_{i,m}$ .
  - 8:          $S_i \leftarrow S_{i-1} \circ ((\bar{x}_{i,m}, \bar{y}_{i,m}))$ .
  - 9:     **end for**
  - 10: Output  $S_q$ .
- 

Below we show that  $\mathcal{A}_{\text{gen}}$  improves over the two stream algorithms presented above, in that it selects exactly  $q$  elements, and has a uniform upper bound on the expected number of iterations, for any source distribution. First, we prove that  $\mathcal{A}_{\text{gen}}$  indeed emulates any pool-based algorithm. The proof is provided in Appendix A.

**Theorem 2** For any pool algorithm  $\mathcal{A}_p$ , any distribution  $\mathcal{D}$  over  $\mathcal{X} \times \mathcal{Y}$ , any integer  $m$  and  $q \leq m$ ,  $\mathcal{A}_s := \mathcal{A}_{\text{gen}}(\mathcal{A}_p)$  is  $(q, \mathcal{D})$ -equivalent to  $\mathcal{A}_p$ .

The next theorem provides an upper bound on the expected number of elements observed by  $\mathcal{A}_{\text{gen}}$ . Unlike  $\mathcal{A}_{\text{wait}}$ , this upper bound holds uniformly for all source distributions.

**Theorem 3** For any pool algorithm  $\mathcal{A}_p$ , any distribution  $\mathcal{D}$  over  $\mathcal{X} \times \mathcal{Y}$ , any integer  $m$  and  $q \leq m$ , if  $\mathcal{A}_s := \mathcal{A}_{\text{gen}}(\mathcal{A}_p)$ ,  $N_{\text{sel}}(\mathcal{A}_s, S, q) = q$  for any  $S \in (\mathcal{X} \times \mathcal{Y})^\infty$ , and

$$\mathbb{E}_{S \sim \mathcal{D}^\infty}[N_{\text{iter}}(\mathcal{A}_s, S, q)] \leq m^2 \left( \frac{em}{q-1} \right)^{q-1}.$$

**Proof** First, clearly  $N_{\text{sel}}(\mathcal{A}_s, S, q) = q$  for any  $S \sim \mathcal{D}^\infty$ . We now prove the upper bound on the expected number of iterations of  $\mathcal{A}_s$ . Let  $S \sim \mathcal{D}^m$ . For  $i \geq 1$ ,  $z_1, \dots, z_{i-1} \in \mathcal{X}$ , denote  $Z_j = \{z_1, \dots, z_j\}$ , and let

$$p_i(z_1, \dots, z_i) := \mathbb{P}[\text{sel}_p(S, [i]) =_\pi Z_i \mid Z_i \subseteq_\pi S_X].$$

Suppose that  $(S_{i-1})_X =_\pi Z_{i-1}$ . The expected number of times that steps 3 to 6 are repeated for index  $i$  is the inverse of the probability that the condition in 6 holds. This condition, in our notation, is that  $\text{sel}_p(S_{i-1} \circ S'_i, [i-1]) =_\pi Z_{i-1}$  and  $\text{sel}_p(S_{i-1} \circ S'_i, i) = \bar{x}_{i,m}$ . We have, from the permutation invariance of  $\mathcal{A}_p$ ,

$$\mathbb{P}[\text{sel}_p(S_{i-1} \circ S'_i, [i-1]) =_\pi Z_{i-1} \mid (S_{i-1})_X =_\pi Z_{i-1}] = p_{i-1}(z_1, \dots, z_{i-1}).$$

In addition, for every draw of  $S'_i$ ,

$$\mathbb{P}[\text{sel}_p(S_{i-1} \circ S'_i, i) = \bar{x}_{i,m} \mid \text{sel}_p(S_{i-1} \circ S'_i, [i-1]) =_\pi Z_{i-1} \wedge (S_{i-1})_X =_\pi Z_{i-1}] = \frac{1}{m-i+1}.$$

This is since under the conditional, one of the elements in  $S'_i$  must be selected by  $\mathcal{A}_p$  in round  $i$ . Therefore, the probability that the condition in step 6 holds is  $p_{i-1}(z_1, \dots, z_{i-1})/(m-i+1)$ . The expected number of times that steps 3 to 6 are repeated for index  $i$  is the inverse of that, and in each round  $m-i+1$  elements are observed. Therefore the expected number of elements observed until selection  $i$  is made conditioned on  $z_1, \dots, z_{i-1}$  is  $(m-i+1)^2/p_{i-1}(z_1, \dots, z_{i-1})$ . The unconditional expected number of elements observed until selection  $i$  is  $(m-i+1)^2 \cdot \mathbb{E}[1/p_{i-1}(\text{sel}_s(S', [i-1]))]$ .

For a set of indices  $J$ , denote  $S|_J = \{S(j) \mid j \in J\}$ .

$$\begin{aligned} \mathbb{E}[1/p_i(\text{sel}_s(S', [i]))] &= \mathbb{E}[1/p_i(\text{sel}_p(S, [i]))] \\ &= \int_{\{z_1, \dots, z_i\} \subseteq \mathcal{X} \times \mathcal{Y}} d\mathbb{P}[\text{sel}_p(S, [i]) =_\pi Z_i] \cdot \frac{1}{p_i(z_1, \dots, z_i)} \\ &= \int_{\{z_1, \dots, z_i\} \subseteq \mathcal{X} \times \mathcal{Y}} d\mathbb{P}[Z_i \subseteq_\pi S_X], \end{aligned}$$

Hence

$$\begin{aligned} \mathbb{E}[1/p_i(\text{sel}_s(S', [i]))] &\leq \int_{\{z_1, \dots, z_i\} \subseteq \mathcal{X} \times \mathcal{Y}} \sum_{J \subseteq [m], |J|=i} d\mathbb{P}[(S|_J)_X = Z_i] \\ &= \sum_{J \subseteq [m], |J|=i} \int_{\{z_1, \dots, z_i\} \subseteq \mathcal{X} \times \mathcal{Y}} d\mathbb{P}[(S|_J)_X = Z_i] \\ &= \sum_{J \subseteq [m], |J|=i} 1 = \binom{m}{i}. \end{aligned}$$

It follows that the expected number of elements observed after the  $i-1$ 'th selection and until selection  $i$  is at most  $(m-i+1)^2 \binom{m}{i-1}$ . We conclude that

$$\mathbb{E}[N_{\text{iter}}(\mathcal{A}_s, S, q)] \leq \sum_{i=0}^{q-1} (m-i)^2 \binom{m}{i} \leq m^2 \left( \frac{em}{q-1} \right)^{q-1}.$$

This completes the proof. ■

From the existence of  $\mathcal{A}_{\text{gen}}$  we can conclude that the pool-based and the stream-based setting are essentially equivalent, up to the number of observed elements. However, the expected number of observed elements is exponential in  $q$ . In the next section we show that this exponential dependence cannot be avoided for general pool algorithms.

#### 4.1. A lower bound for expected number of iterations

We provide a lower bound, which shows that for some pool algorithm, any equivalent stream algorithm has an expected number of observed elements which is at least exponential in  $q$ . This indicates that not much improvement can be achieved over  $\mathcal{A}_{\text{gen}}$  for the class of all pool-based algorithms. The proof involves constructing a pool-based algorithm in which the last selected element determines the identity of the previously selected elements. This is easy in a pool setting, since the algorithm has advance knowledge of all the available elements. In a stream setting, however, this requires a possibly long wait to obtain the matching last element. Because the stream algorithm is allowed to select elements in a different order than the pool algorithm, additional care is taken to make sure that in this case, it is not possible to circumvent the problem this way. The proof of Theorem 4 is provided in Appendix A.

**Theorem 4** *There is an integer  $q_0$  and a constant  $C > 0$ , such that for  $q \geq q_0$ , if  $4q^2 \log(4q) \leq m$ , then there exist a pool algorithm  $\mathcal{A}_p$  and a marginal  $\mathcal{D}_X$ , such that any stream algorithm  $\mathcal{A}_s$  which is  $(q, \mathcal{D})$  equivalent to  $\mathcal{A}_p$  for all  $\mathcal{D} \in \text{DS}(\mathcal{D}_X)$ , and selects only  $q$  elements, has*

$$\exists \mathcal{D} \in \text{DS}(\mathcal{D}_X), \mathbb{E}_{S \sim \mathcal{D}^\infty} [N_{\text{iter}}(\mathcal{A}, S, q)] \geq C \left( \frac{m}{q^2 \log(4q)} \right)^{\frac{q-1}{2}}.$$

### 5. Utility-based pool algorithms

$\mathcal{A}_{\text{gen}}$  gives a uniform guarantee on expected the number of iterations, however this guarantee is exponential  $q$ . We now consider a more restricted class of pool algorithms, and show that it allows emulation with an expected number of iterations linear in  $q$ .

A common approach for designing pool-based interactive algorithms, employed, e.g., in [Seung et al. \(1992\)](#); [Lewis and Gale \(1994\)](#); [Tong and Koller \(2002\)](#); [Guo and Greiner \(2007\)](#); [Golovin et al. \(2010b\)](#); [Guillory and Bilmes \(2010\)](#); [Golovin and Krause \(2011\)](#); [Gonen et al. \(2013\)](#); [Cuong et al. \(2014\)](#), is to define a utility function, that scores each element depending on the history of selected elements and their responses so far. In each round, the algorithm selects the element that maximizes the current utility function. We consider black-box emulation for this class of pool-based algorithms.

Formally, a utility-based interactive pool algorithm is defined by a utility function  $\mathcal{U}$ , of the form  $\mathcal{U} : \cup_{n=0}^{\infty} (\mathcal{X} \times \mathcal{Y})^n \times \mathcal{X} \rightarrow \mathbb{R}$ .  $\mathcal{U}(x, S_{t-1})$  is the score of element  $x$  given history  $S_{t-1}$ . The pool algorithm selects, in each round, the element that is assigned the maximal score by the utility function given the history. We assume for simplicity that there are no ties in  $\mathcal{U}$ . The utility-based interactive pool algorithm for  $\mathcal{U}$ , denoted  $\mathcal{A}_p^{\mathcal{U}}$ , is defined in Alg. 4.

**Algorithm 4**  $\mathcal{A}_p^{\mathcal{U}}$ 


---

**input** Elements  $x_1, \dots, x_m$ , budget  $q < m$ .

- 1:  $S_0 \leftarrow ()$
  - 2:  $M_0 \leftarrow [m]$
  - 3: **for**  $t = 1 : q$  **do**
  - 4:    $i_t \leftarrow \operatorname{argmax}_{j \in M_{t-1}} \mathcal{U}(x_j, S_{t-1})$ .
  - 5:   Select  $x_{i_t}$ , get  $y_{i_t}$ .
  - 6:    $S_t \leftarrow S_{t-1} \circ (x_{i_t}, y_{i_t})$ .
  - 7:    $M_t \leftarrow M_{t-1} \setminus \{i_t\}$ .
  - 8: **end for**
  - 9: Output the set of all pairs in  $S_q$ .
- 

**5.1. An stream algorithm for utility-based pool algorithms**

We propose a stream algorithm  $\mathcal{A}_s^{\mathcal{U}}$  that emulates utility-based pool algorithms  $\mathcal{A}_p^{\mathcal{U}}$ . We stress that we do not attempt to maximize the value of  $\mathcal{U}$  on selected elements, but to emulate the behavior of the pool algorithm that uses  $\mathcal{U}$ . This is because we do not assume any specific relationship between the value of the utility function and the reward of the algorithm. For instance, the utility-based pool algorithm might be empirically successful although its analysis is not fully understood (e.g. [Tong and Koller, 2002](#)).

The definition of  $\mathcal{A}_s^{\mathcal{U}}$  uses the solution to the well-known *secretary problem* ([Dynkin, 1963](#); [Gilbert and Mosteller, 1966](#); [Ferguson, 1989](#)). In the classical formulation of this problem, an algorithm sequentially observes a stream of  $n$  real numbers, and selects a single number. The goal of the algorithm is to select the maximal number out of the  $n$ , but it can only select a number immediately after it is observed, before observing more numbers. It is assumed that the  $n$  numbers in the stream are unknown and selected by an adversary, but their order of appearance is uniformly random. The goal is to select the maximal number with a maximal probability, where  $n$  is known to the algorithm.

This task can be optimally solved by a simple deterministic algorithm, achieving a success probability  $p_{\text{sp}}(n)$ , which satisfies  $\lim_{n \rightarrow \infty} p_{\text{sp}}(n) = 1/e$ . The optimal algorithm observes the first  $t(n)$  numbers, then selects the next observed number which is at least as large as the first  $t(n)$ . The limit of  $t(n)/n$  for  $n \rightarrow \infty$  is  $1/e$ .

Given a stream of size  $k$  of real values  $R = (r_1, \dots, r_k)$ , we say that  $\text{SecPr}(n, R)$  holds if the optimal solution to the secretary problem for size  $n$  selects  $r_k$  after observing the stream prefix  $R$ .  $\mathcal{A}_s^{\mathcal{U}}$  is given in [Alg. 5](#). It uses repeated applications of the solution to the secretary problem to retrieve each of the selected elements. Because the solution succeeds with a probability less than 1, its application might fail. This can be identified in retrospect. In this case, a new solution is selected. This trial-and-error approach means that  $\mathcal{A}_s^{\mathcal{U}}$  usually selects more than  $q$  elements. However the expected number of selected elements is a constant factor over  $q$ .

To make sure the equivalence holds,  $\mathcal{A}_s^{\mathcal{U}}$  never selects an element that could not have been in a pool in which the previous elements have been selected. This is achieved by discarding such elements in each round. The upper bound on the expected number of observed elements bounds the expected number of elements discarded in this way.

First, we show that  $\mathcal{A}_s^{\mathcal{U}}$  is indeed equivalent to  $\mathcal{A}_p^{\mathcal{U}}$ . The proof is provided in [Appendix A](#).



---

**Algorithm 5**  $\mathcal{A}_{\mathcal{U}}^s$

---

```

1:  $L_0 \leftarrow ()$ 
2:  $\mathcal{X}_1 = \mathcal{X}$ 
3: for  $i = 1 : q$  do
4:   repeat
5:     for  $j = 1 : m - i + 1$  do
6:       Repeatedly draw elements from  $\mathcal{D}_X$ , until drawing an element in  $\mathcal{X}_i$ .
       Denote it  $x_{i,j}$ , and let  $r_{i,j} \leftarrow \mathcal{U}(x_{i,j}, L_{i-1})$ .
7:       if  $\text{SecPr}(m - i + 1, (r_{i,1}, \dots, r_{i,j}))$  then
8:          $k \leftarrow j$ 
9:         Select  $x_{i,k}$ , get its response  $y_{i,k}$ .
10:      end if
11:    end for
12:    until  $r_{i,k} = \max\{r_{i,1}, \dots, r_{i,m-i+1}\}$ .
13:     $k_i \leftarrow k$ 
14:     $L_i \leftarrow L_{i-1} \circ (x_{i,k_i}, y_{i,k_i})$ .
15:     $\mathcal{X}_{i+1} \leftarrow \{x \in \mathcal{X}_i \mid \mathcal{U}(x, L_{i-1}) < \mathcal{U}(x_{i,k_i}, L_{i-1})\}$ 
16:  end for
17: Output the set of pairs in  $L_q$ .
```

---

**Theorem 5** For any utility function  $\mathcal{U}$ , any distribution  $\mathcal{D}$  over  $\mathcal{X} \times \mathcal{Y}$ , any integer  $m$  and  $q \leq m$ ,  $\mathcal{A}_s^{\mathcal{U}}$  is  $(q, \mathcal{D})$ -equivalent to  $\mathcal{A}_p^{\mathcal{U}}$ .

The following theorem give an upper bound on the expected number of selected elements and the expected number of observed elements used by  $\mathcal{A}_s^{\mathcal{U}}$ .

**Theorem 6** For any utility function  $\mathcal{U}$ , any distribution  $\mathcal{D}$  over  $\mathcal{X} \times \mathcal{Y}$ , any integer  $m$  and  $q \leq m$ ,

$$\mathbb{E}_{S \sim \mathcal{D}^\infty} [N_{\text{sel}}(\mathcal{A}_s^{\mathcal{U}}, S, q)] = p_{\text{sp}}^{-1}(m)q,$$

and

$$\mathbb{E}_{S \sim \mathcal{D}^\infty} [N_{\text{iter}}(\mathcal{A}_s^{\mathcal{U}}, S, q)] \leq p_{\text{sp}}^{-1}(m) \exp\left(\frac{q}{m-q}\right) \cdot qm.$$

For  $q \leq m/2$ , and  $m \rightarrow \infty$ , it follows from Theorem 6 that the expected number of selected elements is  $eq$ , and the expected number of observed elements is at most  $e^2qm$ .

**Proof** [of Theorem 6] Call a full run of the loop starting at step 5 an attempt for the  $i$ 'th element. In each attempt for the  $i$ 'th element,  $m - i + 1$  elements from  $\mathcal{X}_i$  are observed. The expected number of attempts for each element  $i$  is  $e$ , since each attempt is a run of the secretary problem, with a success probability of  $p_{\text{sp}}(m)$ . Therefore, the expected number of elements from  $\mathcal{X}_i$  observed until  $x_i$  is selected is  $p_{\text{sp}}^{-1}(m) \cdot (m - i + 1)$ .

Denote by  $f_i$  the utility function  $\mathcal{U}(\cdot, L_{i-1})$ . Let  $x_i := x_{i,k_i}$ , be the  $i$ 'th element added to  $L_i$ . Then  $\mathcal{X}_i = \{x \in \mathcal{X}_{i-1} \mid f_{i-1}(x) \leq f_{i-1}(x_{i-1})\}$ .

Consider the probability space defined by the input to the stream algorithm  $S \sim \mathcal{D}^\infty$ , and let  $Z_i, Z'_i \sim \mathcal{D}_X$  for  $i \in [q]$  such that these random variables and  $S$  are all independent. Denote

$$p(\alpha, i) := \mathbb{P}[f_i(Z_i) \leq \alpha \mid Z_i \in \mathcal{X}_i].$$

$p(\alpha, i)$  is a random variable since  $\mathcal{X}_i$  depends on  $S$ . Let  $U_i := p(f_i(Z'_i), i)$ . Since we assume no ties in  $\mathcal{U}$ , and no single  $x$  has a positive probability in  $\mathcal{D}_X$ , then conditioned on  $\mathcal{X}_i$ ,  $U_i$  is distributed uniformly in  $[0, 1]$ . Hence  $U_1, \dots, U_q$  are statistically independent.

For  $i > 1$ , define the random variable  $M_i := p(f_{i-1}(x_{i-1}), i-1)$ . Then  $M_i = \mathbb{P}[\mathcal{X}_i]/\mathbb{P}[\mathcal{X}_{i-1}]$ . The expected number of elements that need to be drawn from  $\mathcal{D}$  to get a single element from  $\mathcal{X}_i$  is  $1/\mathbb{P}[\mathcal{X}_i] = (\prod_{j=1}^i M_j)^{-1}$ . Therefore,

$$\mathbb{E}[N_{\text{iter}}(\mathcal{A}_s^{\mathcal{U}}, S, q) \mid M_2, \dots, M_q] = \sum_{i=1}^q \frac{p_{\text{sp}}^{-1}(m) \cdot (m - i + 1)}{\prod_{j=1}^i M_j}.$$

The element  $x_i$  maximizes the function  $x \mapsto f_i(x)$  over  $m - i + 1$  independent draws of elements  $x$  from  $\mathcal{D}_X$  conditioned on  $x \in \mathcal{X}_i$ , hence it also maximizes  $x \mapsto p(f_i(x), i)$ . Therefore, for  $i > 1$ ,  $M_i$  is the maximum of  $m - i + 2$  independent copies of  $U_i$ , hence  $P[M_i \leq p] = p^{m-i+2}$ . Hence

$$dP[M_2, \dots, M_q](p_2, \dots, p_q)/dp_2 \cdots dp_q = \prod_{i=2}^q dP[M_i \leq p_i]/dp_i = \prod_{i=2}^q (m - i + 2)p_i^{m-i+1}.$$

We have

$$\begin{aligned} \mathbb{E}[N_{\text{iter}}(\mathcal{A}_s^{\mathcal{U}}, S, q)] &= \int_{M_2=0}^1 \cdots \int_{M_q=0}^1 \mathbb{E}[N_{\text{iter}}(\mathcal{A}_s^{\mathcal{U}}, S, q) \mid M_1, \dots, M_q] dP[M_1, \dots, M_q] \\ &= \int_{M_2=0}^1 \cdots \int_{M_q=0}^1 \sum_{i=1}^q \frac{p_{\text{sp}}^{-1}(m) \cdot (m - i + 1)}{\prod_{j=1}^i M_j} \prod_{l=2}^q (m - l + 2) M_l^{m-l+1} dM_l \\ &= \sum_{i=1}^q p_{\text{sp}}^{-1}(m) \cdot (m - i + 1) \int_{M_2=0}^1 \cdots \int_{M_q=0}^1 \prod_{l=2}^i (m - l + 2) M_l^{m-l} dM_l \\ &\quad \cdot \prod_{l=i+1}^q (m - l + 2) M_l^{m-l+1} dM_l, \end{aligned}$$

Therefore

$$\begin{aligned} \mathbb{E}[N_{\text{iter}}(\mathcal{A}_s^{\mathcal{U}}, S, q)] &= \sum_{i=1}^q p_{\text{sp}}^{-1}(m) \cdot (m - i + 1) \prod_{l=2}^i \frac{m - l + 2}{m - l + 1} \\ &= \sum_{i=1}^q p_{\text{sp}}^{-1}(m) \cdot (m - i + 1) \prod_{l=2}^i \left(1 + \frac{1}{m - l + 1}\right) \\ &\leq p_{\text{sp}}^{-1}(m) \cdot qm \left(1 + \frac{1}{m - q}\right)^q \leq p_{\text{sp}}^{-1}(m) \cdot \exp\left(\frac{q}{m - q}\right) \cdot mq. \end{aligned}$$

This concludes the proof. ■

## 5.2. A lower bound for expected number of iterations

The following lower bound shows that the expected number of observed elements required by Alg. 5 cannot be significantly improved by any emulation of general utility-based pool algorithms. This

theorem holds for stream algorithms that select exactly  $q$  elements, while Alg. 5 selects approximately  $eq$  elements. We conjecture that even if allowing a constant factor more element selections, one can achieve at most a constant factor improvement in the expected number of observed elements.

The proof of the lower bound follows by constructing a utility function which in effect allows only one set of selected elements, and has an interaction pattern that forces the stream algorithm to select them in the same order as the pool algorithm. For a given distribution  $\mathcal{D}_X$  over  $\mathcal{X}$ , let  $\text{DS}(\mathcal{D}_X)$  be the set of distributions over  $\mathcal{X} \times \mathcal{Y}$  such that their marginal over  $\mathcal{X}$  is equal to  $\mathcal{D}_X$ . The proof of Theorem 7 is provided in Appendix A.

**Theorem 7** *For any  $m \geq 8$ ,  $q \leq m/2$ , there exists a utility-based pool algorithm, and a marginal  $\mathcal{D}_X$ , such that any stream algorithm  $\mathcal{A}_s$  which is  $(q, \mathcal{D})$  equivalent to the pool algorithm for all  $\mathcal{D} \in \text{DS}(\mathcal{D}_X)$ , and selects only  $q$  elements, has*

$$\exists \mathcal{D} \in \text{DS}(\mathcal{D}_X), \mathbb{E}_{S \sim \mathcal{D}^\infty} [N_{\text{iter}}(\mathcal{A}_s, S, q)] \geq \frac{q}{8} \left\lfloor \frac{m}{2 \log(2q)} \right\rfloor.$$

## 6. Active Learning for Binary Classification

In active learning for binary classification, recent works provide relatively tight label complexity bounds, that hold for both the stream-based and the pool-based settings. In Balcan and Long (2013), tight upper and lower bounds for active learning of homogeneous linear separators under log-concave distributions are provided. The bounds hold for both the stream-based and the pool-based setting, and with the same bound on the number of unlabeled examples. In Hanneke and Yang (2015), tight minimax label complexity bounds for active learning are provided for several classes of distributions. These bounds also hold for both the stream-based and the pool-based setting. In that work no restriction is placed on the number of unlabeled examples.

These results leave open the possibility that for some distributions, a pool-based algorithm with the same label complexity as a stream-based algorithm might require significantly fewer unlabeled examples. In Example 1 and Theorem 8 we show that this is indeed the case.

**Example 1** *For given integers  $m$  and  $q \leq m$ , and  $T \leq q$ , define  $\mathcal{X} = \{a_{k,j} \mid k \in [q], j \in \{0, \dots, 2^{\min(k,T)} - 1\}\} \cup \mathcal{X}'$ , where  $\mathcal{X}'$  includes arbitrary elements so that  $|\mathcal{X}| = n$ , for some  $n \geq q2^T/2$ . Define the following hypothesis class  $\mathcal{H} \subseteq \mathcal{Y}^{\mathcal{X}}$ .*

$$\mathcal{H} := \{h_i \mid i \in \{0, \dots, 2^q - 1\}\}, \text{ where } h_i(a_{k,j}) = \begin{cases} \mathbb{I}[i \bmod 2^k = j] & k \leq T, \\ \mathbb{I}[\lfloor i/2^{T-k} \rfloor \bmod 2^T = j]. & k > T. \end{cases} \quad (1)$$

*Essentially, for  $k \leq T$ ,  $h_i(a_{k,j}) = 1$  if the  $k$  least significant bits in the binary expansion of  $i$  are equal to the binary expansion of  $j$  to  $T$  bits. For  $k \geq T$ ,  $h_i(a_{k,j}) = 1$  if  $T$  consecutive bits in  $i$ , starting from bit  $T - k$ , are equal to the binary expansion of  $j$ .*

**Theorem 8** *Let  $q \geq 22$  and  $m \geq 8 \log(2q)q^2$  be integers. Consider Example 1 with  $m, q$ , setting  $T = \lceil \log_2(q) \rceil$  and  $n = \lfloor m/7 \log(2q) \rfloor$ . Consider  $\mathcal{H}$  as defined in Eq. (1). There exist  $\delta, \epsilon \in (0, 1)$  such that there is a pool-based active learning algorithm that uses a pool of  $m$  unlabeled examples and  $q$  labels, such that for any distribution  $\mathcal{D}$  which is consistent with some  $h^* \in \mathcal{H}$  and has a uniform marginal over  $\mathcal{X}$ , with a probability of at least  $1 - \delta$ ,  $\mathbb{P}[\hat{h}(X) \neq h^*(X)] \leq \epsilon$ . On the other*

hand, for  $q > 22$ , any stream-based active learning algorithm with the same guarantee requires at least  $\frac{q}{32} \left\lceil \frac{m}{7 \log(2q)} \right\rceil$  unlabeled examples in expectation.

The proof is provided in Appendix A. This result shows that a gap between the stream-based and the pool-based settings exists not only for general interactive algorithms, but also specifically for active learning for binary classification.

The gap is more significant when  $q = \tilde{\Theta}(\sqrt{m})$ , and can be as large as  $\tilde{\Omega}(m^{3/2})$  unlabeled examples in a stream, versus  $m$  that are required in a pool. It has been previously observed (Gonen et al., 2013) that in some cases, a specific pool-based active learning algorithm for halfspaces is superior to the classical stream-based algorithm CAL (Cohn et al., 1994). Theorem 8 shows that this is not a limitation specifically of CAL, but of any stream-based active learning algorithm.

The upper bound in Theorem 6 for utility-based pool algorithms can be applied for several deterministic pool-based active-learning algorithms which use a utility function (e.g., Golovin and Krause, 2011; Gonen et al., 2013; Cuong et al., 2014). The upper bound shows that when the label budget  $q$  is relatively small, the gap between the stream and the pool settings is not significant. For instance, consider an active learning problem in which a utility-based pool active learner achieves a label complexity close to the information-theoretic lower bound for the realizable setting (Kulkarni et al., 1993), so that  $q \in \Theta(\log(1/\epsilon))$ . The passive learning sample complexity is at most  $m \in \Theta(1/\epsilon)$ . Therefore, a stream-based active learner with the same properties needs at most  $O(\log(1/\epsilon)/\epsilon)$  unlabeled examples. Therefore, in this case the difference between the pool-based setting and the stream-based setting can be seen as negligible.

## 7. Conclusions

In this work we studied the relationship between the stream-based and the pool-based interactive settings, by designing algorithms that emulate pool-based behavior in a stream-based setting, and proving upper and lower bounds on the stream sizes required for such emulation. Our results concern mostly the case where the label budget of the stream algorithm is similar or identical to that of the pool algorithm. We expect that as the label budget grows, there should be a smooth improvement in the expected stream length, which should approach  $m$  as the label budget approaches  $m$ . There are many open problems left for further work. Among them, whether it is possible to emulate utility based pool algorithms with a linear stream size in  $q$  and exactly  $q$  labels, and a relaxation of the requirement for exact equivalence, which would perhaps allow using smaller streams.

## Acknowledgments

This work was supported in part by the Israel Science Foundation (grant No. 555/15).

## References

- A. Arlotto, E. Mossel, and J. Michael Steele. Quickest online selection of an increasing subsequence of specified size. *arXiv preprint arXiv:1412.7985*, 2014.
- M.-F. Balcan and Phil Long. Active and passive learning of linear separators under log-concave distributions. In *Proceedings of the Twenty-Sixth Annual Conference on Computational Learning Theory (COLT)*, pages 288–316, 2013.

- M.-F. Balcan, A. Beygelzimer, and J. Langford. Agnostic active learning. *Journal of Computer and System Sciences*, 75(1):78–89, 2009.
- M. Bouguelia, Y. Belaid, and A. Belaïd. A stream-based semi-supervised active learning approach for document classification. In *Document Analysis and Recognition (ICDAR), 2013 12th International Conference on*, pages 611–615. IEEE, 2013.
- N. Cebron and M. R. Berthold. Active learning for object classification: from exploration to exploitation. *Data Mining and Knowledge Discovery*, 18(2):283–299, 2009.
- W. Chu, M. Zinkevich, L. Li, A. Thomas, and B. Tseng. Unbiased online active learning in data streams. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 195–203. ACM, 2011.
- D. Cohn, L. Atlas, and R. Ladner. Improving generalization with active learning. *Machine Learning*, 15:201–221, 1994.
- N.V. Cuong, W.S. Lee, and N. Ye. Near-optimal adaptive pool-based active learning with general loss. In *30th conference on Uncertainty in Artificial Intelligence*, 2014.
- S. Dasgupta. Analysis of a greedy active learning strategy. *Advances in Neural Information Processing Systems 18 (NIPS)*, 17:337–344, 2005.
- S. Dasgupta. Consistency of nearest neighbor classification under selective sampling. In *COLT*, pages 18–1, 2012.
- E. D. Demaine, P. Indyk, S. Mahabadi, and A. Vakilian. On streaming and communication complexity of the set cover problem. In *Distributed Computing*, pages 484–498. Springer, 2014.
- A. Deshpande, L. Hellerstein, and D. Kletenik. Approximation algorithms for stochastic boolean function evaluation and stochastic submodular set cover. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1453–1467. SIAM, 2014.
- E. B. Dynkin. The optimum choice of the instant for stopping a markov process. In *Soviet Math. Dokl*, volume 4, pages 627–629, 1963.
- T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.
- J. P. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61(313):35–73, 1966.
- D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *Journal of Artificial Intelligence Research*, 42:427–486, 2011.
- D. Golovin, M. Faulkner, and A. Krause. Online distributed sensor selection. In *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*, pages 220–231. ACM, 2010a.
- D. Golovin, A. Krause, and D. Ray. Near-optimal bayesian active learning with noisy observations. In *Advances in Neural Information Processing Systems 23 (NIPS)*, pages 766–774, 2010b.

- A. Gonen, S. Sabato, and S. Shalev-Shwartz. Efficient active learning of halfspaces: an aggressive approach. *Journal of Machine Learning Research*, 14:2487–2519, 2013.
- P. H. Gosselin and M. Cord. Active learning methods for interactive image retrieval. *Image Processing, IEEE Transactions on*, 17(7):1200–1211, 2008.
- A. Guillory and J. A. Bilmes. Interactive submodular set cover. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pages 415–422, 2010.
- Y. Guo and R. Greiner. Optimistic active-learning using mutual information. In *IJCAI*, volume 7, pages 823–829, 2007.
- Y. Guo, I. Silins, U. Stenius, and A. Korhonen. Active learning-based information structure analysis of full scientific articles and two applications for biomedical literature review. *Bioinformatics*, 29(11):1440–1447, 2013.
- S. Hanneke. Teaching dimension and the complexity of active learning. In *Proceedings of the Twentieth Annual Conference on Computational Learning Theory (COLT)*, 2007.
- S. Hanneke. Rates of convergence in active learning. *The Annals of Statistics*, 39(1):333–361, 2011.
- S. Hanneke and L. Yang. Minimax analysis of active learning. *Journal of Machine Learning Research*, 16:3487–3602, 2015.
- V. Krishnamurthy. Algorithms for optimal scheduling and management of hidden markov model sensors. *Signal Processing, IEEE Transactions on*, 50(6):1382–1397, 2002.
- S. R. Kulkarni, S. K. Mitter, and J. N. Tsitsiklis. Active learning using arbitrary binary valued queries. *Machine Learning*, 11(1):23–35, 1993.
- D. D. Lewis and W. A. Gale. A sequential algorithm for training text classifiers. In *Proceedings of the 17th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 3–12. Springer-Verlag New York, Inc., 1994.
- C.C. Loy, T.M. Hospedales, T. Xiang, and S. Gong. Stream-based joint exploration-exploitation active learning. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 1560–1567, June 2012.
- A. K. McCallum and K. Nigam. Employing em and pool-based active learning for text classification. In *Proceedings of the Fifteenth International Conference on Machine Learning (ICML)*, 1998.
- P. Mitra, C.A. Murthy, and S. K. Pal. A probabilistic active support vector learning algorithm. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(3):413–418, 2004.
- S. Sabato and R. Munos. Active regression by stratification. In Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 469–477, 2014.
- S. Sabato, A. D. Sarwate, and N. Srebro. Auditing: Active learning with outcome-dependent query costs. In *Advances in Neural Information Processing Systems 26 (NIPS)*, 2013.

- H. S. Seung, M. Opper, and H. Sompolinsky. Query by committee. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 287–294. ACM, 1992.
- A. Singla, S. Tschitschek, and A. Krause. Noisy submodular maximization via adaptive sampling with applications to crowdsourced image collection summarization. In *Conference on Artificial Intelligence (AAAI)*, 2016.
- J. Smailović, M. Grčar, N. Lavrač, and M. Žnidaržič. Stream-based active learning for sentiment analysis in the financial domain. *Information Sciences*, 285(0):181 – 203, 2014.
- M. Streeter and D. Golovin. An online algorithm for maximizing submodular functions. In *Advances in Neural Information Processing Systems*, pages 1577–1584, 2009.
- S. Tong and E. Chang. Support vector machine active learning for image retrieval. In *Proceedings of the Ninth ACM International Conference on Multimedia*, MULTIMEDIA '01, pages 107–118. ACM, 2001.
- S. Tong and D. Koller. Support vector machine active learning with applications to text classification. *The Journal of Machine Learning Research (JMLR)*, 2:45–66, 2002.

## Appendix A. Additional Proofs

Several proofs use the following lemma.

**Lemma 9** *Let  $\alpha \in (0, \frac{1}{2})$ ,  $p \in (0, \alpha^2/2)$ . Let  $X_1, X_2, \dots$  be independent Bernoulli random variables with  $\mathbb{P}[X_i = 1] \leq p$ . Let  $I$  be a random integer, which can be dependent on the entire sequence  $X_1, X_2, \dots$ . Suppose that  $\mathbb{P}[X_I = 1] \geq \alpha$ . Then  $\mathbb{E}[I] \geq \frac{\alpha^2}{2p}$ .*

**Proof**  $\mathbb{E}[I]$  is minimized under the constraint when  $\mathbb{P}[X_i = 1] = p$ . Therefore assume this equality holds. Let  $W$  be the random variable whose value is the smallest integer such that  $X_W = 1$ . Let  $T$  be the largest integer such that  $\mathbb{P}[W \leq T] \leq \alpha$ .

The expectation of  $I$  is lower bounded subject to  $\mathbb{P}[X_I = 1] \geq \alpha$  by  $I$  such that  $\mathbb{P}[I = W \mid W \leq T] = 1$ ,  $\mathbb{P}[I = W \mid W = T + 1] = \alpha - \mathbb{P}[W \leq T]$ , and in all other cases,  $I = 0$ . Therefore,

$$\mathbb{E}[I] \geq \mathbb{E}[W \cdot \mathbb{I}[W \leq T]].$$

We have

$$\begin{aligned} \frac{1}{p} &= \mathbb{E}[W] = \mathbb{E}[W \cdot \mathbb{I}[W \leq T]] + \mathbb{E}[W \cdot \mathbb{I}[W > T]] \\ &= \mathbb{E}[W \cdot \mathbb{I}[W \leq T]] + \left(\frac{1}{p} + T\right)(1 - p)^T. \end{aligned}$$

Therefore

$$\mathbb{E}[I] \geq \mathbb{E}[W \cdot \mathbb{I}[W \leq T]] = \frac{1}{p} - \left(\frac{1}{p} + T\right)(1 - p)^T.$$

From the definition of  $T$ ,  $T$  is the largest integer such that  $1 - (1 - p)^T \leq \alpha$ . Hence  $T \geq \frac{\log(1-\alpha)}{\log(1-p)}$  and  $(1 - p)^T \leq (1 - \alpha)/(1 - p)$ . Therefore

$$\mathbb{E}[I] \geq \frac{1}{p} - \left( \frac{1}{p} + \frac{\log(1-\alpha)}{\log(1-p)} \right) \frac{1-\alpha}{1-p} \geq \frac{1}{p} - \left( \frac{1}{p} - \frac{\log(1-\alpha)}{2p} \right) \frac{1-\alpha}{1-p}$$

Hence

$$p\mathbb{E}[I] \geq 1 + \frac{1-\alpha}{1-p}(\log(1-\alpha)/2 - 1)$$

For  $p \leq a^2/2$  and  $\alpha \in (0, 1/2)$ , elementary calculus shows that  $p\mathbb{E}[I] \geq \alpha^2/2$ . ■

**Proof** [of Theorem 2] Consider the probability space defined by the infinite sequence  $S' \sim \mathcal{D}^\infty$  which generates the input to the stream algorithm, and an independent sequence  $S \sim \mathcal{D}^m$  which is the input to the pool algorithm.

For  $z_1, \dots, z_q \in \mathcal{X} \times \mathcal{Y}$ , denote  $Z_j = \{z_1, \dots, z_j\}$ . We have, for every  $i \in [q]$ ,

$$\begin{aligned} d\mathbb{P}[\text{pairs}_p(S, [i]) =_\pi Z_i] &= \\ \sum_{j=1}^i d\mathbb{P}[\text{pairs}_p(S, i) = z_j \mid \text{pairs}_p(S, [i-1]) =_\pi Z_i \setminus \{z_j\}] \cdot d\mathbb{P}[\text{pairs}_p(S, [i-1]) =_\pi Z_i \setminus \{z_j\}]. \end{aligned}$$

The same holds for  $\text{pairs}_s(S', \cdot)$ . To show the equivalence it thus suffices to show that for all  $z_1, \dots, z_q \in \mathcal{X} \times \mathcal{Y}$ ,  $i \in [q]$ ,

$$d\mathbb{P}[\text{pairs}_s(S', i) = z_i \mid \text{pairs}_s(S', [i-1]) =_\pi Z_{i-1}] = d\mathbb{P}[\text{pairs}_p(S, i) = z_i \mid \text{pairs}_p(S, [i-1]) =_\pi Z_{i-1}].$$

From the definition of  $\mathcal{A}_s$  we have

$$\begin{aligned} d\mathbb{P}[\text{pairs}_s(S', i) = z_i \mid \text{pairs}_s(S', [i-1]) =_\pi Z_{i-1}] &= \\ = d\mathbb{P}[\text{pairs}_p(S_{i-1} \circ S'_i, i) = z_i \mid S_{i-1} =_\pi Z_{i-1} \wedge \text{pairs}_p(S_{i-1} \circ S'_i, [i-1]) =_\pi Z_{i-1}] &= \\ = d\mathbb{P}[\text{pairs}_p(S, i) = z_i \mid \text{pairs}_p(S, [i-1]) =_\pi Z_{i-1}]. \end{aligned}$$

The last equality follows since  $\mathcal{A}_p$  is permutation invariant and never selects the same index twice. This proves the equivalence. ■

**Proof** [of Theorem 4] Denote by  $\Pi_k$  the set of permutations over  $[k]$ . Let the domain of elements be  $\mathcal{X} = [0, 2]$  and assume responses in  $\mathcal{Y} = \{0, 1\}$ . We now define a pool algorithm as follows. Call a pool  $S_X$  in which exactly one element in the pool is in  $(1, 2]$  and the rest are in  $[0, 1]$  a “good pool”. On bad pools,  $\mathcal{A}_p$  always selects only elements in  $[0, 1]$  or only elements in  $(1, 2]$ .

For a good pool, denote for simplicity the single element in  $(1, 2]$  by  $x_m$ , and other elements by  $x_1, \dots, x_{m-1}$ , where  $x_{i-1} < x_i$  for  $i \in [m-1]$ . Define a mapping  $\psi : (1, 2] \rightarrow \Pi_{m-1}$ , such that if  $x_m$  is uniform over  $(1, 2]$ , then for  $\psi(x_m)$  all permutations in the range are equally likely.

$\mathcal{A}_p$  behaves as follows: Let  $\sigma = \psi(x_m)$ . The first  $q-1$  elements it selects are  $x_{\sigma(1)}, \dots, x_{\sigma(q-1)}$ . The last element it selects is  $x_m$  if the response for all previous elements was 0, and  $x_{\sigma(q)}$  otherwise.

Define the marginal  $\mathcal{D}_X$  over  $\mathcal{X}$  in which for  $X \sim \mathcal{D}_X$ ,  $\mathbb{P}[X \in [0, 1]] = 1 - 1/m$ ,  $\mathbb{P}[X \in (1, 2]] = 1/m$ , and in each range  $[0, 1]$ ,  $(1, 2]$ ,  $X$  is uniform. The probability of a good pool under



$\mathcal{D} \in \text{DS}(\mathcal{D}_X)$  is  $(1 - 1/m)^{m-1} \geq 1/e^2 =: p_g$ . We now show a lower bound on the expected number of iterations of a stream algorithm which is  $(q, \mathcal{D})$ -equivalent to any  $\mathcal{D} \in \text{DS}(\mathcal{D}_X)$ . Let  $\mathcal{D}_0$  be the distribution over  $\mathcal{X} \times \mathcal{Y}$  such that for  $(X, Y) \sim \mathcal{D}_0$ ,  $X \sim \mathcal{D}_X$  and  $Y = 0$  with probability 1. Let  $S \sim \mathcal{D}_0^m$  be the input to  $\mathcal{A}_p$ .

The proof will follow a series of claims:

1. The probability that, on a good pool,  $\psi(x_m)$  is in a given set of permutation  $\Phi(Z)$ , where  $Z$  is the set of first  $q - 1$  selected elements, is at least  $1/2$ .
2. When  $\mathcal{A}_s$  emulates a good pool, it selects an element from  $(1, 2]$  only after selecting  $q - 1$  elements from  $[0, 1]$ .
3. Therefore, when  $\mathcal{A}_s$  emulates a good pool, the expected number of observed elements until selecting the last element is lower bounded, and so the overall expected number is lower bounded.

We start with **claim 1**. For a given set  $Z = \{z_1, \dots, z_{q-1}\} \subseteq [0, 1]$ , define the set of permutations  $\Phi(Z)$  as follows. The expected number of elements that are smaller than  $z_i$  in  $S_X \sim \mathcal{D}_X^m$ , if  $Z \subseteq_\pi S_X$ , is  $n_i = (m - q)z_i + \sum_{j=1}^{q-1} \mathbb{I}[z_j < z_i]$ . Let  $\epsilon := \sqrt{(m - q) \log(4q)}/2$ , and define

$$\Phi(Z) := \{\sigma \in \Pi_{m-1} \mid \exists \sigma' \in \Pi_{q-1}, \forall i \in [q-1], |\sigma^{-1}(i) - n_{\sigma'(i)}| \leq \epsilon\}. \quad (2)$$

These are the permutations such that the first  $q-1$  elements according to the permutation are mapped from elements with ranks in  $[n_i - \epsilon, n_i + \epsilon]$ . For  $x \in S_X$ , denote by  $r_S(x)$  the rank of  $x$  in  $S_X$ , when the elements in  $S_X$  are ordered by value. Since  $\psi(\text{sel}_p(S, q))$  determines the choice of  $Z$  from  $S_X$ , we have

$$\begin{aligned} & \mathbb{P}[\psi(\text{sel}_p(S, q)) \in \Phi(Z) \mid \text{sel}_p(S, [q-1]) =_\pi Z \wedge S \text{ is good}] \\ & \geq \mathbb{P}[\forall i \in [q-1], |r_S(z_i) - n_i| \leq \epsilon \mid \text{sel}_p(S, [q-1]) =_\pi Z \wedge S \text{ is good}] \\ & = \mathbb{P}[\forall i \in [q-1], |r_S(z_i) - n_i| \leq \epsilon \mid Z \subseteq_\pi S \wedge S \text{ is good}]. \end{aligned}$$

The last inequality follows since  $\psi(\text{sel}_p(S, q))$  is uniform over all permutations. By Hoeffding's inequality, for any  $i \leq q - 1$ ,

$$\mathbb{P}[|r_S(z_i) - n_i| > \epsilon \mid Z \subseteq_\pi S \wedge S \text{ is good}] \leq 2 \exp(-2\epsilon^2/(m - q)).$$

Therefore, using the definition of  $\epsilon$  and applying the union bound, we get, for any  $Z \subseteq [0, 1]$  with  $|Z| = q - 1$ ,

$$\mathbb{P}[\psi(\text{sel}_p(S, q)) \in \Phi(Z) \mid \text{sel}_p(S, [q-1]) =_\pi Z \wedge S \text{ is good}] \geq \frac{1}{2}. \quad (3)$$

This completes the proof of **claim 1**.

We now turn to **claim 2**. Consider a stream algorithm which is  $(q, \mathcal{D})$ -equivalent to  $\mathcal{A}_p$  for any  $\mathcal{D} \in \text{DS}(\mathcal{D}_X)$ . Consider runs of  $\mathcal{A}_s$  with input  $S' \sim \mathcal{D}_0^\infty$ . Denote by  $E_g$  the event that the output of  $\mathcal{A}_s$  is equal to a possible output of  $\mathcal{A}_p$  on a good pool with  $S \sim \mathcal{D}_0^m$ . Then  $\mathbb{P}[E_g] \geq p_g$ . Claim 2 is that

$$\mathbb{P}[\text{sel}_s(S', [q-1]) \subseteq_\pi [0, 1] \mid E_g] = 1. \quad (4)$$

In other words, when simulating a good pool, the elements in  $[0, 1]$  are all selected before the element in  $(1, 2]$ .

To show claim 2, note that by the definition of  $\mathcal{A}_p$ , for any source distribution over  $\mathcal{X} \times \mathcal{Y}$ , if  $\mathcal{A}_p$  outputs a set with elements both in  $[0, 1]$  and in  $(1, 2]$ , then there is exactly one element in  $(1, 2]$  in the output, and all the responses in the output for elements in  $[0, 1]$  are 0 with probability 1.

Now, suppose that  $\mathbb{P}[\text{sel}_s(S', [q-1]) \subseteq_\pi [0, 1] \mid E_g] < 1$ . Then  $\mathbb{P}[\text{sel}_s(S', q) \in [0, 1] \mid E_g] > 0$ , since there can be only one element in  $(1, 2]$  in the output of a good pool. But, consider running  $\mathcal{A}_s$  with a source distribution  $\mathcal{D}' \in \text{DS}(\mathcal{D}_X)$  such that for  $(X, Y) \sim \mathcal{D}'$ ,  $X \sim \mathcal{D}_X$  and  $\mathbb{P}_{\mathcal{D}'}[Y = 0 \mid X = x] = \frac{1}{2}$  for all  $x$ . There is a positive probability that in the first  $q-1$  selected elements all the responses are 0, just as for  $\mathcal{D}_0$ . Therefore, also for  $S'' \sim \mathcal{D}'^\infty$ ,  $\mathbb{P}[\text{sel}_s(S'', q) \in [0, 1] \mid E_g] > 0$ . But then there is a positive probability that the response for the last element, which is in  $[0, 1]$ , is 1, contradicting the  $(q, \mathcal{D}')$ -equivalence of the pool and  $\mathcal{A}_s$ . This proves **claim 2**.

We now show **claim 3** which completes the proof. From claim 2 in Eq. (4), we conclude that  $\mathbb{P}[\text{sel}_s(S', q) \in (1, 2] \mid E_g] = 1$ . Therefore, from claim 1 in Eq. (3), for any  $Z \subseteq_\pi [0, 1]$  with  $|Z| = q-1$ ,

$$\mathbb{P}[\psi(\text{sel}_s(S', q)) \in \Phi(\text{sel}_s(S', [q-1])) \mid E_g] \geq 1/2.$$

Therefore

$$\mathbb{P}[\psi(\text{sel}_s(S', q)) \in \Phi(\text{sel}_s(S', [q-1]))] \geq \mathbb{P}[E_g]/2 \geq p_g/2.$$

Now, let  $X_i \sim \mathcal{D}_X$  be the  $i$ 'th element observed after selecting the first  $q-1$  elements, and let  $B_i = \mathbb{I}[\psi(X_i) \in \Phi(Z)]$ , where  $Z$  is the set of  $q-1$  selected elements.  $B_i$  are independent Bernoulli random variables, each with a probability of success at most  $p$ , where from the definition of  $\phi$  in Eq. (2),

$$p \leq \frac{|\Phi(Z)|}{|\Pi_{m-1}|} \leq \left( \frac{(q-1)(2\epsilon+1)}{m-1} \right)^{q-1} \leq \left( \frac{2q^2 \log(4q)}{m} \right)^{\frac{q-1}{2}}.$$

Let  $I$  be the number of elements  $\mathcal{A}_s$  observes after selecting  $Z$ , until selecting element  $q$ . We have  $\mathbb{P}[B_I = 1] \geq p_g/2$ . By Lemma 9, for  $p \leq p_g^2/8$ ,  $p\mathbb{E}[I] \geq p_g^2/8$ . From the assumption in the theorem statement,  $2q^2 \log(4q)/m \leq \frac{1}{2}$ , hence for a large enough  $q$ ,  $p \leq 2^{-(q-1)} \leq p_g^2/8$ , and so  $\mathbb{E}[I] \geq \frac{p_g^2}{8} p^{-1}$ . Hence there is a constant such that

$$\mathbb{E}[I] \geq C \left( \frac{m}{q^2 \log(4q)} \right)^{\frac{q-1}{2}}.$$

Since  $\mathbb{E}[N_{\text{iter}}(\mathcal{A}, S, q)] \geq \mathbb{E}[I]$ , this completes **claim 3** and finalizes the proof.  $\blacksquare$

**Proof** [of Theorem 5] Consider the probability space defined by  $S \sim \mathcal{D}^m$  and  $S' \sim \mathcal{D}^\infty$ , where  $S, S'$  are independent. We prove the equivalence by showing that for any  $j \in [q]$  and  $L_j = ((x_{i,k_i}, y_{i,k_i}))_{i \in [j]}$  that could have been selected by the pool algorithm,

$$d\mathbb{P}[\text{pairs}_p(S, j+1) \mid \text{pairs}_p(S, [j]) = L_j] = d\mathbb{P}[\text{pairs}_s(S', j+1) \mid \text{pairs}_s(S', [j]) = L_j].$$

For a given  $L_j$ , denote by  $\mathcal{D}_{j+1}$  the distribution generated by drawing  $(X, Y) \sim \mathcal{D}$  conditioned on  $X \in \mathcal{X}_{j+1}$ , where  $\mathcal{X}_{j+1}$  depends on  $L_j$ . Denote by  $\mathcal{G}$  all the finite sequences of pairs such that when

the optimal secretary problem solution is applied to the sequence, it succeeds. That is, the optimal value under the score  $(x, y) \rightarrow \mathcal{U}(x, L_j)$  is indeed selected. From the definition of  $\mathcal{A}_{\mathcal{U}}^s$ , we have

$$d\mathbb{P}[\text{pairs}_s(S', j+1) \mid \text{pairs}_s(S', [j]) = L_j] = d\mathbb{P}_{\bar{S} \sim \mathcal{D}_{j+1}^{m-j}}[\arg\max_{(x,y) \in \bar{S}} \mathcal{U}(x, L_j) \mid \bar{S} \in \mathcal{G}].$$

For a given sequence  $\bar{S} = ((\bar{x}_i, y_i))_{i \in [m-j]}$ , let  $\sigma(\bar{S}) : [m-j] \rightarrow [m-j]$  be a permutation such that for all  $i \leq m-j$ ,  $\bar{x}_{\sigma(i)} \leq \bar{x}_{\sigma(i+1)}$ . The success of the optimal secretary problem algorithm depends only on the ordering of ranks in its input sequence, hence there is a set of permutations  $\mathcal{G}'$  such that  $\bar{S} \in \mathcal{G}$  if and only if  $\sigma(\bar{S}) \in \mathcal{G}'$ . Now,  $\arg\max_{(x,y) \in \bar{S}} \mathcal{U}(x, L_j)$  depends only on the identity of pairs in  $\bar{S}$ , while  $\sigma(\bar{S})$  depends only on their order. Since the elements in  $\bar{S}$  are i.i.d., these two properties are independent. Therefore

$$d\mathbb{P}_{\bar{S} \sim \mathcal{D}_{j+1}^{m-j}}[\arg\max_{(x,y) \in \bar{S}} \mathcal{U}(x, L_j) \mid \bar{S} \in \mathcal{G}] = d\mathbb{P}_{\bar{S} \sim \mathcal{D}_{j+1}^{m-j}}[\arg\max_{(x,y) \in \bar{S}} \mathcal{U}(x, L_j)].$$

Therefore

$$\begin{aligned} & d\mathbb{P}[\text{pairs}_s(S', j+1) \mid \text{pairs}_s(S', [j]) = L_j] \\ &= d\mathbb{P}_{\bar{S} \sim \mathcal{D}_{j+1}^{m-j}}[\arg\max_{(x,y) \in \bar{S}} \mathcal{U}(x, L_j)] \\ &= d\mathbb{P}_{\hat{S} \sim \mathcal{D}^{m-j}}[\arg\max_{(x,y) \in \hat{S}} \mathcal{U}(x, L_j) \mid \hat{S} \subseteq (\mathcal{X}_{j+1} \times \mathcal{Y})^{m-j}] \\ &= d\mathbb{P}_{\hat{S} \sim \mathcal{D}^{m-j}}[\arg\max_{(x,y) \in \hat{S}} \mathcal{U}(x, L_j) \mid \forall (x, y) \in \hat{S}, i \in [j], \mathcal{U}(x, L_{i-1}) < \mathcal{U}(x_{i,k_i}, L_{i-1})] \\ &= d\mathbb{P}_{\hat{S} \sim \mathcal{D}^{m-j}}[\arg\max_{(x,y) \in \hat{S}} \mathcal{U}(x, L_j) \mid \text{pairs}_p(L_j \circ \hat{S}, [j]) = L_j] \\ &= d\mathbb{P}_{S \sim \mathcal{D}^m}[\arg\max_{(x,y) \in S \setminus L_j} \mathcal{U}(x, L_j) \mid \text{pairs}_p(S, [j]) = L_j] \\ &= d\mathbb{P}_{S \sim \mathcal{D}^m}[\text{pairs}_p(S)(j+1) \mid \text{pairs}_p(S, [j]) = L_j]. \end{aligned}$$

Here  $L_i$  is the prefix of length  $i$  of  $L_j$ . Since this equality holds for all  $j \in [q-1]$ ,  $d\mathbb{P}[\text{pairs}_s(S', [q])] = d\mathbb{P}[\bar{O}_q(S, [q])]$ .  $\blacksquare$

**Proof** [of Theorem 7] Let  $n = \lfloor \frac{m}{2 \log(2q)} \rfloor$ , and let  $\mathcal{D}_X$  be a uniform distribution over  $\mathcal{X} = \{a_i \mid i \in [n]\}$ . Assume  $\mathcal{Y} = \{0, 1\}$ . A pool of size  $m$  then includes all elements in  $A := \{a_i \mid i \in [2q-1]\}$  with a probability of at least  $\alpha \geq 1 - (2q-1) \exp(-m/n) \geq 1 - \frac{1}{2q}$ .

Consider a utility function  $\mathcal{U}$  such that given a history of the form  $((a_1, 0), \dots, (a_t, 0))$  for  $t \in [q-1]$ , assigns a maximal score in  $\mathcal{X}$  to  $a_{t+1}$ , and given a history of the form  $((a_1, 0), \dots, (a_{t-1}, 0), (a_t, 1))$ , for  $t \in [q-1]$ , assigns a maximal score in  $\mathcal{X}$  to  $a_{q+t-1}$ . Then, in a pool that includes all elements  $a_1, \dots, a_{2q-1}$ , the pool algorithm based on  $\mathcal{U}$  behaves as follows: In every round, if all selected elements so far received the response 0, it selects at round  $t$  the element  $a_t$ . Otherwise, it selects the element  $a_{q+t-1}$ .

Let  $\mathcal{D}_0$  be a distribution in which the response is deterministically zero. If the distribution is  $\mathcal{D}_0$ ,  $\mathcal{A}_s$  selects  $Z_0 = \{a_1, \dots, a_q\}$  with a probability of at least  $\alpha$ . Denote  $\mathcal{D}_t$  for  $t \in [q]$ , in which the response is deterministically zero for  $X \in \{a_1, \dots, a_q\} \setminus \{a_t\}$  and 1 for  $a_t$ . For this distribution,

the algorithm must select the elements in  $Z_t = \{a_1, \dots, a_t, a_{q+t}, \dots, a_{2q-1}\}$  with a probability at least  $\alpha$ .

We show a lower bound on the probability that  $\mathcal{A}_s$  selects  $a_1, \dots, a_q$  in order when the input sequence is  $S \sim \mathcal{D}_0^\infty$ . Denote this probability  $\beta$ , and the event that this occurs  $E$ .

Consider the random process defined by the input sequence  $S \sim \mathcal{D}_0^\infty$  and the randomness of  $\mathcal{A}_s$ . Let  $T$  be a random variable, such that  $T$  is the smallest round in which the algorithm selects some  $a_{t'}$ , for  $t' > T$ , or  $T = 0$  if no such round exists. Since  $\mathbb{P}[T \in [q]] = 1 - \beta$ , there exists some  $t^* \in [q]$  such that  $\mathbb{P}[T = t^*] \geq (1 - \beta)/q$ . Now, consider the distribution  $\mathcal{D}_{t^*}$ . Define a sequence of pairs  $\gamma(S)$  such that  $S$  and  $\gamma(S)$  have the same elements in the same order, and the responses in  $\gamma(S)$  are determined by  $\mathcal{D}_{t^*}$  instead of by  $\mathcal{D}_0$ . Clearly,  $\gamma(S)$  is distributed according to  $\mathcal{D}_{t^*}^\infty$ . Consider a run of the algorithm on  $S$  and a parallel run (with the same random bits) on  $\gamma(S)$ . The algorithm selects the same elements for both sequences until the  $T$ 'th selection, inclusive. But the  $T$ 'th selection is some element in  $\{a_{T+1}, \dots, a_q\}$ . If  $T = t^*$ , then  $Z_{t^*}$  does not include the element selected in round  $T$ . Since  $\mathcal{A}_s$  selects exactly the set  $Z_{t^*}$  with a probability of at least  $\alpha$ , we have  $\mathbb{P}[T = t^*] \leq 1 - \alpha$ . Therefore  $(1 - \beta) = \mathbb{P}[T \in [q]] \leq q(1 - \alpha)$ , hence  $\beta \geq \frac{1}{2}$ .

Let  $W_i$  be the number of elements that  $\mathcal{A}_s$  observes after selecting element  $i - 1$ , until observing the next element. Let  $X_i \sim \mathcal{D}_X$  be the  $i$ 'th element observed after selecting the first  $i - 1$  elements, and let  $B_i = \mathbb{I}[X_i = a_i]$ .  $B_i$  are independent Bernoulli random variables with  $\mathbb{P}[B_i = 1] = 1/n$ , and  $\mathbb{P}[B_{W_i} = 1] \geq \mathbb{P}[E] = \beta \geq \frac{1}{2}$ . By Lemma 9, if  $\frac{1}{n} \leq \frac{1}{8}$ ,  $\mathbb{E}[W_i] \geq \frac{n}{8}$ .

It follows that the expected number of iterations over  $q$  selections is at least  $\frac{qn}{8} = \frac{q}{8} \left\lfloor \frac{m}{2 \log(2q)} \right\rfloor$ .

**Proof** [of Theorem 8] Let  $\mathcal{D}_X$  be uniform over  $\mathcal{X}$ . Let  $E$  be the event that  $\mathcal{X} \not\subseteq_{\pi} S_X$ , and define  $\delta := \mathbb{P}_{S \sim \mathcal{D}_X^m}[E]$ . Define  $\epsilon = 1/n$ , so that  $\mathbb{P}[\hat{h}(X) \neq h^*(X)] < \epsilon$  if and only if  $\hat{h} = h^*$ . Let  $i^*$  such that  $h^* = h_{i^*}$ .

First, a pool-based algorithm can achieve the required accuracy as follows: Let  $j_t := i^* \bmod 2^t$  for  $t \leq T$ , and  $j_t := \lfloor i^*/2^{T-t} \rfloor \bmod 2^T$  for  $t \geq T$ . If  $E$  holds, then  $t$ 'th element selected by the pool algorithm is  $a_{t,j}$ , where  $j$  is obtained as follows: If  $t \leq T$ ,  $j = j_{t-1}$ . If  $t > T$ ,  $j = \lfloor j_{t-1}/2 \rfloor$ . In round 1,  $j = 0$  and the selected element is  $a_{1,0}$ . Inductively, in this strategy the algorithm finds the  $t$ 'th least significant bit in the binary expansion of  $i^*$  in round  $t$ , thus it can use  $j_{t-1}$  to set  $j$  for round  $t$ . Under  $E$ , after  $q$  labels  $i^*$  is identified exactly. This happens with a probability of  $1 - \delta$  for any  $\mathcal{D}$  with the uniform marginal  $\mathcal{D}_X$ .

Now, let  $\mathcal{D}_h$  be a distribution with a uniform marginal over  $\mathcal{X}$  with labels consistent with  $h \in \mathcal{H}$ . Consider a stream-based algorithm  $\mathcal{A}_s$ , denote its output by  $\bar{h}$  and its input by  $S \sim \mathcal{D}_{h^*}^\infty$ .

Let  $I$  be a random variable drawn uniformly at random from  $\{0, \dots, 2^q - 1\}$ . Let  $H = h_I$  be a hypothesis chosen uniformly at random from  $\mathcal{H}$ . Consider the probability space defined by  $I, S \sim \mathcal{D}_H^\infty$ , and the run of  $\mathcal{A}_s$  on  $S$ . Let  $(Z_1, Y_1), \dots, (Z_q, Y_q)$  be the examples that  $\mathcal{A}_s$  receives and the labels it gets, in order. Let  $Y = (Y_1, \dots, Y_q)$ . Let  $\alpha = \mathbb{P}[Z_1 = a_{1,0} \mid S_X]$ . If  $Z_1 = a_{1,0}$ , then  $\mathbb{P}[Y_1 = 0 \mid S_X] = \frac{1}{2}$ . If  $Z_1 \neq a_{1,0}$ , then  $\mathbb{P}[Y_1 = 0 \mid S_X] \geq 3/4$ . Let  $\mathbb{H}$  be the base-2 entropy, and  $\mathbb{H}_b$  be the binary entropy. Then  $\mathbb{H}_b(Y_1 \mid S_X) = \mathbb{H}_b((\alpha + 1)/4)$ , and so

$$\begin{aligned} \mathbb{H}(H \mid Y, S_X) &= \mathbb{H}(H, Y \mid S_X) - \mathbb{H}(Y_1 \mid S_X) - \mathbb{H}(Y_1, \dots, Y_q \mid Y_1, S_X) \\ &\geq q - \mathbb{H}_b((\alpha + 1)/4) - (q - 1) \\ &= 1 - \mathbb{H}_b((\alpha + 1)/4). \end{aligned}$$

From the Taylor expansion of the binary entropy around  $1/2$ ,  $\mathbb{H}_b(p) \leq 1 - (1 - 2p)^2/2$ , therefore  $\mathbb{H}(H | Y, S_X) \geq (1 - \alpha)^2/8$ . We have  $\mathbb{P}[\bar{h} \neq H] \leq \delta$ , hence  $\mathbb{P}_{S_X}[\mathbb{P}[\bar{h} \neq H | S_X] \leq 2\delta] \geq \frac{1}{2}$ . By Fano's inequality, for any  $S_X$  such that  $\mathbb{P}[\bar{h} \neq H | S_X] \leq 2\delta$ ,

$$(1 - \alpha)^2/8 \leq \mathbb{H}(H | Y, S_X) \leq \mathbb{H}_b(2\delta) + 2\delta q \leq 2\delta(\log_2(\frac{1}{2\delta}) + 2 + q).$$

Where the last inequality follows from  $\mathbb{H}_b(p) \leq p \log_2(1/p) + 2p$ . From the definition of  $\delta$ , we have  $\delta \leq |\mathcal{X}| \exp(-m/n)$ . Setting  $T = \lceil \log_2(q) \rceil$ , and noting that  $|\mathcal{X}| \leq q2^T/2 \leq q^2$  and so  $m \geq n \log(128q^3|\mathcal{X}|)$ , we have  $\delta \leq \frac{1}{128q^3}$ .

Therefore, for  $q \geq 22$ ,  $1 - \alpha \leq \frac{1}{2q}$ .

It follows that  $\mathbb{P}_{S_X}[\mathbb{P}[Z_1 \neq a_{1,0} | S_X] \leq \frac{1}{2q}] \geq 1/2$ . Now, the same argument holds for any round  $t$  conditioned on  $I \bmod 2^t = 0$  and  $Z_1 = a_{1,0}, \dots, Z_t = a_{t,0}$ , since in this case after  $t$  labels, the algorithm has  $q - t$  queries left, and needs to select from  $\mathcal{H}'$ , which is equivalent to  $\mathcal{H}$ , with  $q - t$  instead of  $q$ . Moreover,  $\mathbb{P}[\bar{h} = H | I \bmod 2^t = 0] \leq 1 - \delta$  as well, since this holds for every  $H$  individually. We conclude that for every  $t \leq q$ , with a probability at least  $\frac{1}{2}$  over  $S_X$ ,

$$\mathbb{P}[Z_t \neq a_{t,0} | S_X, H = h_0] \leq \frac{1}{2q}.$$

It follows that with a probability at least  $\frac{1}{2}$  over  $S_X$ ,  $\mathbb{P}[Z_1 = a_{1,0}, \dots, Z_q = a_{q,0} | S_X, H = h_0] \geq 1/2$ . Hence  $\mathbb{P}[Z_1 = a_{1,0}, \dots, Z_q = a_{q,0} | H = h_0] \geq 1/4$ .

Now, suppose  $H = h_0$ . Let  $W_t$  be the number of elements that  $\mathcal{A}_s$  observes after selecting element  $t - 1$ , until observing the next element. Let  $X_j \sim \mathcal{D}_X$  be the  $j$ 'th element observed after selecting the first  $t - 1$  elements, and let  $B_j = \mathbb{I}[X_j = a_{t,0}]$ .  $B_j$  are independent Bernoulli random variables with  $\mathbb{P}[B_j = 1] = 1/n$ , and  $\mathbb{P}[B_{W_t} = 1] \geq \mathbb{P}[E] = \beta \geq \frac{1}{4}$ . By Lemma 9, if  $\frac{1}{n} \leq \frac{1}{8}$ , then  $\mathbb{E}[W_t] \geq \frac{n}{32}$ . It follows that the expected number of iterations over  $q$  selections is at least  $\frac{qn}{32} \geq \frac{q}{32} \left\lfloor \frac{m}{7 \log(2q)} \right\rfloor$ . ■