

Bayesian active learning with abstention feedbacks

Cuong V. Nguyen ^{a,1}, Lam Si Tung Ho ^b, Huan Xu ^{c,1}, Vu Dinh ^d, Binh T. Nguyen ^{e,f,g,*}



^a KFSCIS, Florida International University, USA

^b Department of Mathematics and Statistics, Dalhousie University, Canada

^c Alibaba Inc., USA

^d Department of Mathematical Sciences, University of Delaware, USA

^e AISIA Research Lab, Ho Chi Minh City, Viet Nam

^f Department of Computer Science, University of Science, Ho Chi Minh City, Viet Nam

^g Vietnam National University, Ho Chi Minh City, Viet Nam

ARTICLE INFO

Article history:

Received 21 January 2021

Revised 31 July 2021

Accepted 14 November 2021

Available online 22 November 2021

Communicated by Zidong Wang

Keywords:

Active learning

Bayesian methods

Abstention feedbacks

ABSTRACT

We study pool-based active learning with abstention feedbacks where a labeler can abstain from labeling a queried example with some unknown abstention rate. This is an important problem with many useful applications. We take a Bayesian approach to the problem and develop two new greedy algorithms that learn both the classification problem and the unknown abstention rate at the same time. These are achieved by simply incorporating the estimated average abstention rate into the greedy criteria. We prove that both algorithms have near-optimality guarantees: they respectively achieve a $(1 - \frac{1}{e})$ constant factor approximation of the optimal expected or worst-case value of a useful utility function. Our experiments show the algorithms perform well in various practical scenarios.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

We consider the problem of active learning with abstention feedbacks where a labeler can abstain from labeling queried examples with some unknown abstention rate. This problem is one of the several attempts to deal with imperfect labelers in active learning who may give incorrect or noisy labels to queried examples [1,2] or in our case, give abstention feedbacks to queries [2–4].

Learning with abstention feedbacks is important in many real-life scenarios. Below we discuss some examples where this problem is useful. In these examples, although the reasons for the abstention vary, from the learner's view, they are the same: the learner will receive no labels for some queries and the true labels for others.

Crowdsourcing: In crowdsourcing, we have many labelers, each of whom only has expertise in some certain area and therefore can only provide labels for a subset of the input domain. These labelers were also called labelers with a knowledge blind spot [3]. In this case, active learning is a good approach to quickly narrow down the expertise domain of a labeler and focus on querying examples in this region to learn a good model. By adapting active

learning algorithms to each labeler, we can also gather representative subsets of labeled data from the labelers and combine them into a final training set.

Learning with Corrupted Labels: In this problem, the abstention feedbacks do not come from the labeler but occur due to corruptions in the labels received by the learner. The corruptions could be caused by bad communication channels that distort the labels or could even be caused by attackers attempting to corrupt the labels [5]. The setting in our paper can deal with the case when the corrupted labels are completely lost, i.e., they cannot be recovered and are not converted to incorrect ones.

In this paper, we consider the pool-based active learning with a fixed budget setting, where a finite pool of unlabeled examples is given in advance, and we need to sequentially select N examples from the pool to query their labels. Our setting assumes that abstention feedbacks count towards the budget N , so we need to be careful when selecting the queried examples. Our work takes a Bayesian approach to the problem and learns both the classification model and the unknown abstention rate at the same time. We call this approach the *Bayesian Active Learning with Abstention Feedbacks* (BALAF) framework. Our framework can be used to instantiate different algorithms for active learning with abstention feedbacks problems.

Unlike previous work, which took a frequentist approach to the problem and did not provide any theoretical insights, our work

* Corresponding author at: Department of Computer Science, University of Science, Vietnam National University, Ho Chi Minh City, Viet Nam.

E-mail address: ngtbinh@hcmus.edu.vn (B.T. Nguyen).

¹ Work done before joining current organizations.

takes the Bayesian approach and also contributes to the understandings of this problem both algorithmically and theoretically. Algorithmically, we develop two novel greedy algorithms for active learning with abstention feedbacks based on our BALAF framework. Each algorithm uses a different greedy criterion to select queried examples that can give information for both the classification model and the abstention rate. Theoretically, we prove that our proposed algorithms have theoretical guarantees for version space reduction, a useful property of the selected examples in active learning. Version space reduction is a popular concept for analyzing properties of active learning algorithms [6–8] and is defined as the volume of all functions in the hypothesis space that disagree with the selected labeled examples. To the best of our knowledge, ours are the first theoretical results for active learning with abstention feedbacks in the Bayesian pool-based setting.

The first greedy algorithm that we propose in this paper aims to maximize the expected version space reduction [6] of the joint deterministic space deduced from the spaces of possible classification models and abstention rates. In essence, the proposed algorithm is a generalization of the maximum Gibbs error algorithm [7] where we incorporate the terms controlling the estimated abstention rate into the greedy criterion. When there is no abstention, our algorithm recovers the maximum Gibbs error algorithm. By using previous theoretical results for adaptive submodularity [6], we are able to prove that our algorithm has an *average-case near-optimality* guarantee: the average utility value of its selected examples is always within a $(1 - \frac{1}{e})$ constant factor of the optimal average utility value.

In contrast to the first algorithm, the second algorithm that we propose aims to maximize the worst-case version space reduction utility. This algorithm is a generalization of the least confidence active learning algorithm [9] with the main difference that we also incorporate the estimated abstention rate into the greedy criterion. When there is no abstention, our algorithm can recover the least confidence algorithm. From previous theoretical results for pointwise submodularity [8], we can prove that the proposed algorithm has a *worst-case near-optimality* guarantee: the worst-case utility value of its selected examples is always within a $(1 - \frac{1}{e})$ constant factor of the optimal worst-case utility value.

We conduct experiments to evaluate our proposed algorithms on various binary classification tasks under three different realistic abstention scenarios. The experiments show that our algorithms are useful compared to the passive learning and normal active learning baselines with various abstention rates under these scenarios.

2. Related work

2.1. Theory

The theoretical guarantees considered in this paper have been studied for normal Bayesian pool-based active learning where the labeler always gives labels to queried examples [6,10,7,8,11–13]. The theory for the average case was originally developed by [6] with adaptive submodular utilities, while that of the worst case was developed by [8] for pointwise submodular utilities. In both cases, $(1 - \frac{1}{e})$ -factor approximation guarantees were proven for the corresponding greedy algorithms.

The problem of active learning with abstention feedbacks was previously investigated in [3,4,2]. [3] considered a setting similar to ours where the labeler may have knowledge blind spots and would be incapable of labeling examples in such blind spots. On the other hand, [4] studied a situation where the learner may interrupt the labeler rather than waiting for his response, thus allowing the possibility of receiving “I don’t know” labels. In both papers,

greedy algorithms were used to create a balance between maximizing the received information and minimizing the abstention probability; however, no theoretical guarantees about their performance were obtained. Our work in this paper, although similar to theirs in spirit, provides theoretical guarantees for the proposed algorithms. In [14], the authors consider a different framework for active learning with abstention where the learner has an option to withhold making a decision for a fixed cost. Our framework, on the other hand, allows the labeler to abstain from labeling while the learner does not have this option.

A theoretical work on active learning with abstention feedbacks is [2], where both noisy labels and abstention feedbacks are considered. However, they only examined a simple one-dimensional classification problem and made a low-noise assumption on both the labeling noise and the abstention rate. Furthermore, the labeler’s abstention in their model is non-persistent and that allows the learner to repeatedly query an example until a label is received. Under this framework, the paper derived an algorithm with a near-optimal asymptotic convergence rate for estimating model parameters. In contrast, our work in this paper investigates the persistent label scenario, which is much less understood and more difficult to resolve [11,15], and we focus on near-optimal query strategies with a finite budget. Thus, our results are not directly comparable to those in [2]. Theory for learning with abstention has also been considered in other settings, such as boosting [16] and online learning [17].

2.2. Methods

Besides abstention feedbacks, there were other works on active learning with unreliable labelers. For examples, many authors considered labelers that give incorrect or corrupted labels from various types of noise models [18–20,13,15]. [21] considered a setting where the labeler can return both labels and confidences, while in [22,1], multiple labelers with different fidelity are available and the learner is given the option of obtaining labels from either weak or strong labelers. Our work also relates to other works on active learning and adaptive sampling in crowdsourcing such as [23–27]. Similar to our work, [28,29] combine an additional objective function with a measurement of uncertainty into the active learning process and also utilize greedy algorithms. However, they focus on different learning settings where there is no abstention from the labelers.

3. Pool-based active learning with abstention feedbacks

In pool-based active learning, we are given a finite set (called a pool) \mathcal{X} of unlabeled examples and a budget N , and we need to sequentially query the labels of N examples from \mathcal{X} to learn a good classifier. Normal active learning assumes the human labeler would always give labels for queried examples. By contrast, in this paper, we consider active learning with abstention feedbacks where the labeler is allowed to abstain from labeling a queried example. In other words, the labeler may return “no label” to a queried example. Our work considers the case where abstention feedbacks count towards the budget N , so we need to select queried examples to obtain as many useful labels as possible.

To define the problem, let $\mathcal{Y} = \{1, 2, \dots, \ell\}$ be the set of all possible labels. Assume there is an unknown true labeling $f_{\text{true}} : \mathcal{X} \rightarrow \mathcal{Y}$ of the whole pool \mathcal{X} that is used by the labeler to label queried examples, and the labeler will return $f_{\text{true}}(x)$ for a queried example x if he decides to label it. Also assume there is an unknown true abstention pattern $k_{\text{true}} : \mathcal{X} \rightarrow \{0, 1\}$ used by the labeler to decide whether or not to label a queried example. That is, $k_{\text{true}}(x) = 1$ if

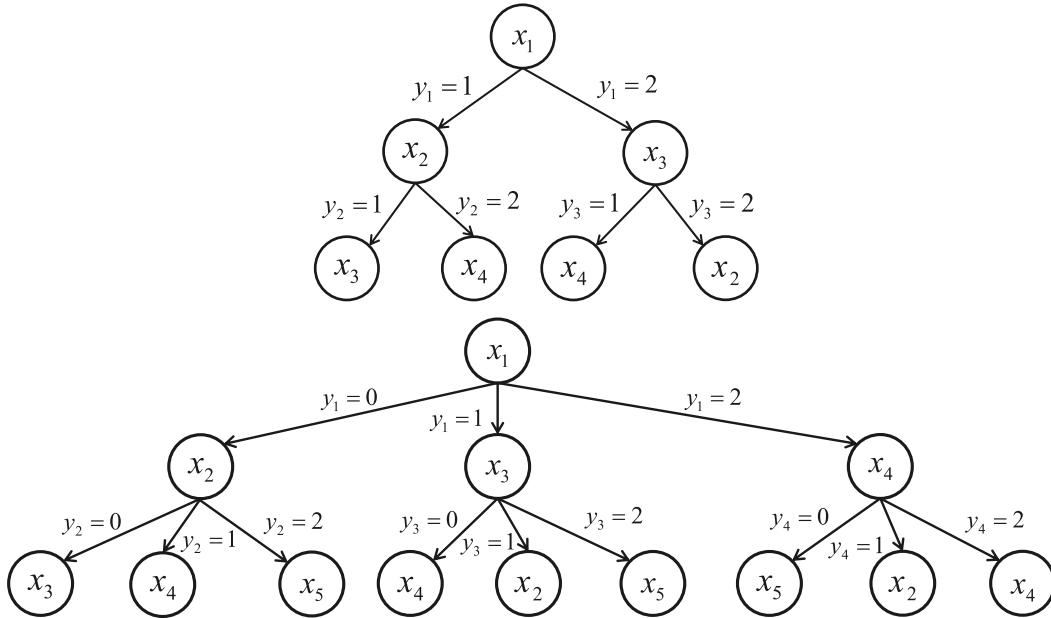


Fig. 1. Policy trees for pool-based active learning without (upper) and with (bottom) abstention feedbacks. $y_i = 0$ means the labeler abstains from labeling x_i while $y_i = 1$ or 2 means the labeler gives label 1 or 2 for x_i respectively.

the labeler abstains from labeling x and $k_{\text{true}}(x) = 0$ if he decides to label it.

In this setting, an active learning algorithm is a policy for choosing queried examples from \mathcal{X} , and these chosen examples depend on the labels as well as abstention feedbacks of previously selected examples. By definition, a policy is a mapping from a set of examples and the labeler's corresponding responses to the next unlabeled example to be queried, and it can be represented by a policy tree. During the active learning process, a learner (a policy or algorithm) sequentially selects unlabeled training examples one at a time from \mathcal{X} and asks the labeler for their labels. The labeler would use k_{true} to decide whether or not to give the labels, which in turn are determined by f_{true} . Illustrations of policy trees for normal active learning and active learning with abstention feedbacks are given in Fig. 1. Pool-based active learning with abstention feedbacks aims to design algorithms (policies) for selecting queried examples that can give us as much information about f_{true} (and in some cases, k_{true}) as possible.

4. Bayesian active learning with abstention feedbacks (BALAF)

We shall take the Bayesian approach to pool-based active learning with abstention feedbacks, which we call the *Bayesian active learning with abstention feedbacks* (BALAF) framework. In our framework, we consider a (possibly infinite) set \mathcal{H} of probabilistic hypotheses, where each hypothesis $h \in \mathcal{H}$ is a random function from \mathcal{X} to \mathcal{Y} . Formally, for any $x \in \mathcal{X}$, $h(x)$ is a categorical distribution with probability mass function $\mathbb{P}[h(x) = y]$ for all $y \in \mathcal{Y}$.

Following the Bayesian approach, we assume a prior distribution $p_0[h]$ on \mathcal{H} . If we observe a label y of an example x , we can use the following Bayes' rule to obtain a posterior distribution:

$$p_0[h|y, x] = \frac{p_0[h]\mathbb{P}[h(x) = y]}{\int p_0[h]\mathbb{P}[h(x) = y]dh} \propto p_0[h]\mathbb{P}[h(x) = y].$$

To deal with the abstention feedbacks, we also take the Bayesian approach and consider a set of possible abstention hypotheses \mathcal{R}

from \mathcal{X} to $[0, 1]$, where each function $r : \mathcal{X} \rightarrow [0, 1]$ gives us the *abstention rate* of the examples in \mathcal{X} . More specifically, $r(x)$ is the probability that the labeler abstains from labeling x , according to the abstention hypothesis r . We also assume a prior $p_0[r]$ on \mathcal{R} . Note that we have slightly abused the notation p_0 for both priors on \mathcal{H} and \mathcal{R} . In this case, p_0 can be thought of as a joint distribution on $\mathcal{H} \times \mathcal{R}$ where the two elements are independent, i.e., $p_0[h \wedge r] = p_0[h]p_0[r]$ for $h \in \mathcal{H}$ and $r \in \mathcal{R}$.

During the active learning process, if we receive a label for an example x , we can update the posterior distribution using the following Bayes' rule:

$$p_0[r|x \text{ has a label}] = \frac{p_0[r](1 - r(x))}{\int p_0[r](1 - r(x))dr} \propto p_0[r](1 - r(x)).$$

Otherwise, if the labeler abstains from labeling x , we can update the posterior distribution using:

$$p_0[r|x \text{ has no label}] = \frac{p_0[r]r(x)}{\int p_0[r]r(x)dr} \propto p_0[r]r(x).$$

We summarize the general BALAF framework in Algorithm 1, where N examples are chosen sequentially and the posteriors are updated according to the above rules. The framework returns the final posteriors $p_N[h]$ and $p_N[r]$ which can be used to make prediction on new examples or to serve as priors in future active learning processes. For example, the label distribution of a new example x can be predicted using the posterior $p_N[h]$ by:

$$p_N[y|x] = \int p_N[h]\mathbb{P}[h(x) = y]dh.$$

The posterior $p_N[r]$, on the other hand, can be used as a prior on r in future active learning processes if the same labeler is employed to give labels. This would enable the learning algorithm to use the prior knowledge about the labeler's preferences to select the most suitable queried examples while avoiding re-learning his abstention patterns from scratch. The posterior $p_N[r]$ can also be transferred and adapted to other labelers who may have similar abstention patterns.

Algorithm 1.

General BALAF framework

input: Priors $p_0[h]$ and $p_0[r]$, budget N .
output: Final posteriors $p_N[h]$ and $p_N[r]$ after N queries.
for $i = 1$ to N **do**
 Select unlabeled data point x^* to query based on some active learning criterion.
 $y^* \leftarrow \text{Query-label}(x^*)$.
 if received label y^* **then**
 Update $p_i[h] \propto p_{i-1}[h] \mathbb{P}[h(x^*) = y^*]$, and
 $p_i[r] \propto p_{i-1}[r] (1 - r(x^*))$.
 else
 Update $p_i[r] \propto p_{i-1}[r] r(x^*)$.
 end if
end for
return $p_N[h], p_N[r]$.

5. BALAF algorithms with theoretical guarantees

In this section, we propose two specific instances of the BALAF framework above that can achieve near-optimality guarantees for reducing the hypothesis spaces that contain f_{true} and k_{true} . Our first algorithm provides an average-case near-optimality guarantee, while the second algorithm provides a worst-case near-optimality guarantee. The algorithms only differ in the ways we choose the queried data point x^* in [Algorithm 1](#). In what follows, for any $S = \{x_1, x_2, \dots, x_n\} \subseteq \mathcal{X}$ and $h \in \mathcal{H}$, we define $h(S) \triangleq \{h(x_1), h(x_2), \dots, h(x_n)\}$. We shall assume $h(x_i)$ and $h(x_j)$ are independent for any fixed h and $i \neq j$. Thus, $h(S)$ is also a categorical distribution with probability mass function $\mathbb{P}[h(S) = \mathbf{y}] = \prod_{i=1}^n \mathbb{P}[h(x_i) = y_i]$ for all $\mathbf{y} = \{y_1, y_2, \dots, y_n\} \in \mathcal{Y}^{|S|}$. We call \mathbf{y} a labeling of S as it contains the labels of the examples in S . For any $S \subseteq \mathcal{X}$ and any distribution $p[h]$ on \mathcal{H} , let \mathbf{Y} be the random variable for the labeling of S with respect to the distribution p . We note that \mathbf{Y} takes values in $\mathcal{Y}^{|S|}$ with probability mass function:

$$p[\mathbf{Y} = \mathbf{y}; S] = \int p[h] \mathbb{P}[h(S) = \mathbf{y}] dh \quad (1)$$

for all $\mathbf{y} \in \mathcal{Y}^{|S|}$. This is also the marginal probability that the labeling of S is \mathbf{y} . As a special case, if S is a singleton $\{x\}$ and Y is the random variable for the label of x , we write $p[Y = y; x]$ for $y \in \mathcal{Y}$ to denote the probability mass function of Y .

5.1. The average-case BALAF algorithm

In this average-case BALAF algorithm, at each iteration i in [Algorithm 1](#), we select the queried data point x^* as follows. First, we estimate the average abstention function $\tilde{r}(x)$ based on the current posterior $p_{i-1}[r]$:

$$\tilde{r}(x) \triangleq \mathbb{E}_{r \sim p_{i-1}}[r(x)] = \int p_{i-1}[r] r(x) dr. \quad (2)$$

Then we select the example x^* to query using the following greedy criterion:

$$x^* = \arg \max_{x \in \mathcal{X}} \{1 - \tilde{r}(x)^2 - (1 - \tilde{r}(x))^2 \sum_{y \in \mathcal{Y}} p_{i-1}[Y = y; x]^2\}, \quad (3)$$

Intuitively, this criterion maximizes the expected one-step utility increment, with the utility function being defined in [Eq. \(4\)](#) below. [Eq. \(3\)](#) resembles the maximum Gibbs error criterion [\[7\]](#) which selects

$$x^* = \arg \max_x \{1 - \sum_y p_{i-1}[Y = y; x]^2\},$$

except that we incorporate the terms $\tilde{r}(x)^2$ and $(1 - \tilde{r}(x))^2$ into the criterion. In fact, our criterion [\(3\)](#) is a generalization of the maximum Gibbs error criterion above. When there is no abstention, i.e. $\tilde{r}(x) = 0$, our algorithm will recover the maximum Gibbs error algorithm.

If we fix the distribution $p_{i-1}[Y = y; x]$, the criterion value in [Eq. \(3\)](#) achieves its maximum when $\tilde{r}(x) = \sum_y p_{i-1}[Y = y; x]^2 / (1 + \sum_y p_{i-1}[Y = y; x]^2)$ and achieves its minimum when $\tilde{r}(x) = 1$. [Fig. 2](#) (top) illustrates this criterion value as a function of $\tilde{r}(x)$. Thus, given $\tilde{r}(x)$ reasonably approximates the true abstention rate, this algorithm would give more preference to the examples with abstention rate near $\sum_y p_{i-1}[Y = y; x]^2 / (1 + \sum_y p_{i-1}[Y = y; x]^2)$ and less preference to the examples with abstention rate near 1.

Near-optimality Guarantee. We now show the average-case near-optimality guarantee for this algorithm. In the context of this paper, near-optimality means the algorithm can achieve a constant factor approximation to the optimal algorithm with respect to some objective function.

To define an objective function that is useful for active learning with abstention feedbacks, we first induce a deterministic hypothesis space equivalent to the original probabilistic hypothesis space \mathcal{H} . In particular, consider the hypothesis space $\mathcal{F} \triangleq \{f : \mathcal{X} \rightarrow \mathcal{Y}\}$ consisting of all deterministic functions from \mathcal{X} to \mathcal{Y} . We induce a new prior q_0 on \mathcal{F} from the original prior p_0 such that $q_0[f] \triangleq p_0[\mathbf{Y} = f(\mathcal{X}); \mathcal{X}]$, the marginal probability that the labeling of the whole pool \mathcal{X} is $f(\mathcal{X})$. For any $S \subseteq \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}^{|S|}$, we can

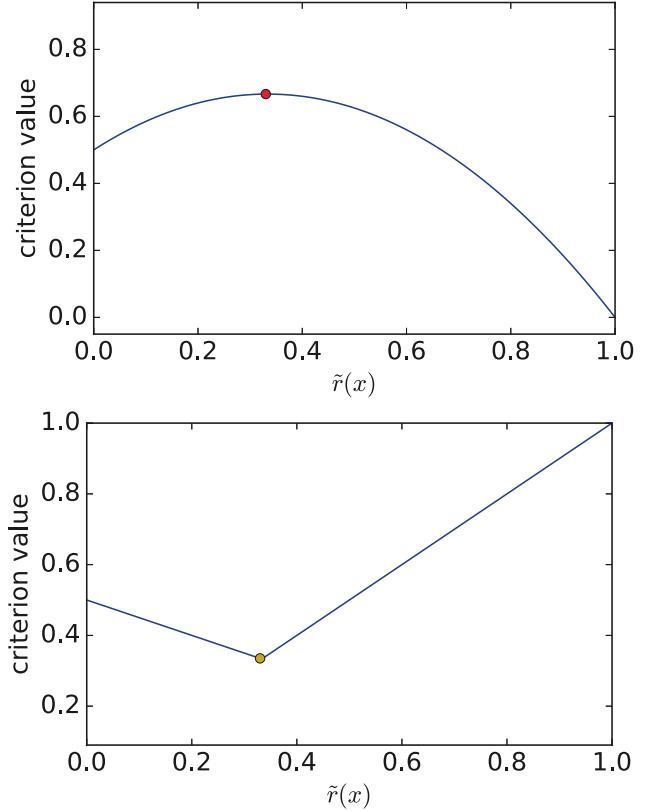


Fig. 2. Graphs showing the greedy criterion values in the average-case BALAF (upper) and worst-case BALAF (bottom) algorithms as a function of $\tilde{r}(x)$ in a binary classification problem. The graphs are plotted with the fixed distribution $p_{i-1}[Y = 1; x] = p_{i-1}[Y = 2; x] = 0.5$. The red and yellow points indicate the maximum and minimum points in the graphs respectively.

define $q_0[\mathbf{Y} = \mathbf{y}; S]$ similarly to Eq. (1) with the hypothesis space \mathcal{F} and distribution q_0 .

Also consider the space $\mathcal{K} \triangleq \{k : \mathcal{X} \rightarrow \{0, 1\}\}$ consisting of all deterministic functions from \mathcal{X} to $\{0, 1\}$. In essence, $k(x) = 1$ means the labeler abstains from labeling x while $k(x) = 0$ means the labeler gives a label for x . We will call each $k \in \mathcal{K}$ an *abstention pattern*. The prior $p_0[r]$ also induces a probability distribution $q_0[k]$ on \mathcal{K} where:

$$\begin{aligned} q_0[k] &\triangleq \int p_0[r] P_r[k] dr, \text{ and} \\ P_r[k] &\triangleq \prod_{x \in \mathcal{X}} (1 - r(x))^{1-k(x)} r(x)^{k(x)} \end{aligned}$$

is the probability (with respect to the rate r) that the labeler gives or abstains from giving labels to the whole pool \mathcal{X} according to the abstention pattern k . For any $S \subseteq \mathcal{X}$ and $\mathbf{z} \in \{0, 1\}^{|S|}$, we can define $q_0[\mathbf{Z} = \mathbf{z}; S]$ similarly to Eq. (1) with the hypothesis space \mathcal{K} and distribution q_0 , where \mathbf{Z} is the random variable for the abstention pattern of S . Note that the induced prior q_0 can also be thought of as a joint prior on $\mathcal{F} \times \mathcal{K}$ where the two elements are independent, i.e., $q_0[f \wedge k] = q_0[f] q_0[k]$.

For $S \subseteq \mathcal{X}, f \in \mathcal{F}$, and $k \in \mathcal{K}$, we consider the utility function:

$$g(S, (f, k)) \triangleq 1 - q_0[\mathbf{Y} = f(S) \wedge \mathbf{Z} = k(S); S], \quad (4)$$

where $q_0[\mathbf{Y} = f(S) \wedge \mathbf{Z} = k(S); S] \triangleq q_0[\mathbf{Y} = f(S); S] \times q_0[\mathbf{Z} = k(S); S]$ is the joint marginal probability (with respect to q_0) that the labeling of S is $f(S)$ and the abstention pattern of S is $k(S)$. This is a useful utility function for active learning because it is the version space reduction utility with respect to the joint prior $q_0[f \wedge k]$ on the joint space $\mathcal{F} \times \mathcal{K}$ [6].

With this utility, our objective function is defined as:

$$G_{\text{avg}}(\pi) \triangleq \mathbb{E}_{f_{\text{true}}, k_{\text{true}} \sim q_0} [g(\mathbf{x}_{f_{\text{true}}, k_{\text{true}}}^{\pi}, (f_{\text{true}}, k_{\text{true}}))], \quad (5)$$

where for all f and k , $\mathbf{x}_{f, k}^{\pi}$ is the set of examples selected by the policy π given that the true labeling is f and the true abstention pattern is k . This objective function is the average of the above utility with respect to the joint prior $q_0[f \wedge k]$. Note that in this objective, f_{true} and k_{true} are drawn from the prior since we operate in the Bayesian setting. The following theorem proves the average-case near-optimality guarantee for the average-case BALAF algorithm. The proof of this theorem is given in the Appendix.

Theorem 5.1. For any budget $N \geq 1$, let π be the policy selecting N examples using the average-case BALAF algorithm and let π_{avg}^* be the optimal policy with respect to G_{avg} that selects N examples. We have:

$$G_{\text{avg}}(\pi) > (1 - \frac{1}{e}) G_{\text{avg}}(\pi_{\text{avg}}^*).$$

5.2. The worst-case BALAF algorithm

The worst-case BALAF algorithm is essentially similar to the previous average-case BALAF algorithm, except that we replace the greedy criterion in Eq. (3) by the following greedy criterion:

$$x^* = \arg \min_{x \in \mathcal{X}} \{ \max \{ \tilde{r}(x), (1 - \tilde{r}(x)) \max_{y \in \mathcal{Y}} p_{i-1}[Y = y; x] \} \}. \quad (6)$$

Intuitively, this criterion maximizes the worst-case one-step utility increment, with the version space reduction utility in Eq. (4). The criterion (6) resembles the least confidence criterion [9], which selects

$$x^* = \arg \min_x \{ \max_y p_{i-1}[Y = y; x] \},$$

except that we also incorporate the terms $\tilde{r}(x)$ and $1 - \tilde{r}(x)$ into the criterion. Our criterion (6) is a generalization of this least confidence

criterion and will recover it when there is no abstention, i.e. when $\tilde{r}(x) = 0$.

If we fix the distribution $p_{i-1}[Y = y; x]$, the criterion value in Eq. (6) achieves its maximum when $\tilde{r}(x) = 1$ and achieves its minimum when $\tilde{r}(x) = \max_y p_{i-1}[Y = y; x] / (1 + \max_y p_{i-1}[Y = y; x])$.

Near-optimality Guarantee. We now show the worst-case near-optimality guarantee for this algorithm. For this guarantee, we still make use of the version space reduction utility function $g(S, (f, k))$ defined in Eq. (4). Using this utility, we define the following worst-case objective function:

$$G_{\text{worst}}(\pi) \triangleq \min_{(f, k) \in \mathcal{F} \times \mathcal{K}} [g(\mathbf{x}_{f, k}^{\pi}, (f, k))]. \quad (7)$$

This objective function is the worst possible utility achieved by the policy π . The following theorem proves the worst-case near-optimality guarantee for this algorithm. The proof of this theorem is given in the Appendix.

Theorem 5.2. For any budget $N \geq 1$, let π be the policy selecting N examples using the worst-case BALAF algorithm and let π_{worst}^* be the optimal policy with respect to G_{worst} that selects N examples. We have:

$$G_{\text{worst}}(\pi) > (1 - \frac{1}{e}) G_{\text{worst}}(\pi_{\text{worst}}^*).$$

5.3. Remarks

A worst-case guarantee can also be derived for the average-case BALAF algorithm in Section 5.1 using Theorem 1 of [6] for adaptive monotone and submodular functions. However, this guarantee is for the min-cost coverage objective, as opposed to the maximum coverage objectives considered in our paper.

Although our version space reduction objectives on the joint abstention and labeling function space do not directly relate to the classification accuracy, previous works [7,8] have empirically shown that reducing the version space of the labeling functions lead to better accuracy for active learning in the non-abstention setting. In this work, we generalize these objective functions to the abstention setting, trying to balance between reducing the version space of the labeling functions and the version space of the abstention functions simultaneously.

6. Experiments

In this section, we experimentally evaluate the proposed BALAF algorithms. In particular, we compare four algorithms:

- PL: the passive learning baseline with randomly selected examples,
- ALg: the active learning baseline using the maximum Gibbs error criterion [7],
- ALa: the average-case BALAF algorithm, and
- ALw: the worst-case BALAF algorithm.

For binary classification, ALg is equivalent to other well-known active learning algorithms such as the least confidence [9] and maximum entropy [30] algorithms. The PL and ALg baselines do not learn the abstention probability of the examples, i.e., they ignore whether an example would be labeled or not when making a decision. In contrast, the proposed algorithms ALa and ALw take into account the estimated abstention probability $\tilde{r}(x)$ when making decisions.

To show the potential of our algorithms further, we also consider two variants of ALa and ALw that are assumed to know a good estimate of the training examples' abstention rates $r^*(x)$. In partic-

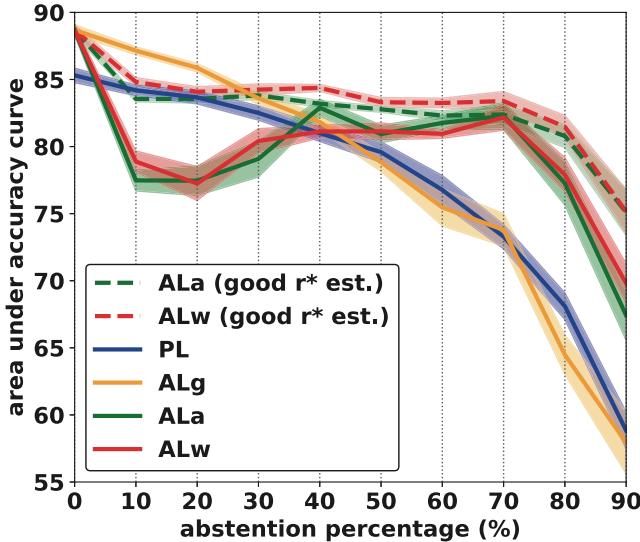


Fig. 3. AUC scores (with standard errors) for a labeler abstaining on examples unrelated to target task.

ular, for these versions of ALa and ALw (shown as dashed lines in Figs. 3 and 5), we train a logistic regression model using the actual abstention pattern on the whole training set to predict the abstention probability for each example. We keep this classifier fixed throughout the experiments and use it to estimate $\tilde{r}(x)$ in these versions of ALa and ALw.

In the algorithms, we use Bayesian logistic regression models for both \mathcal{H} and \mathcal{R} . That is, each hypothesis $h \in \mathcal{H}$ and each abstention hypothesis $r \in \mathcal{R}$ is a logistic regression model. We put an independent Gaussian prior $\mathcal{N}(0, \sigma^2)$ on each parameter of the logistic regression models (for both \mathcal{H} and \mathcal{R}). In this case, the posteriors are proportional to the regularized likelihood of the observed data with ℓ_2 penalty. Since we experiment with data sets containing very high dimensional data (more than 61,000 dimensions), running MCMC or even variational inference is very slow. Thus, for efficiency, we use the maximum a posteriori (MAP) hypotheses to estimate the probabilities in our algorithms. Finding the MAP hypotheses is equivalent to maximizing the regularized log-likelihood of the observed data.

Following previous works in active learning [31,7,8], we evaluate the algorithms using the area under the accuracy curve (AUC) scores. For each task in our experiments, we compute the scores on a separate test set during the first 300 queries and then normalize these scores so that their values are between 0 and 100. The final scores and their errors are obtained by averaging 10 runs of the algorithms using different random seeds.

We shall consider three scenarios: (1) the labeler abstains from labeling examples unrelated to the target classification task, (2) the labeler abstains from labeling easy examples, and (3) the labeler abstains from labeling hard examples.

6.1. Abstention on data unrelated to target task

We consider the binary text classification task between two recreational topics: *rec.motorcycles* and *rec.sport.baseball* from the 20 Newsgroups data [32]. In the pool of unlabeled data, we allow examples from other classes (e.g., in the *computer* category) that are not related to the two target classes. The labeler always abstains from labeling these redundant examples while always giving labels for examples from the target classes. Thus, the abstention is on examples unrelated to the target task, and this satisfies the independence assumption between h and r (or between h and k) in Sections 4 and 5. In the experiment, we fix the pool size to

be 1322 and vary the abstention percentage (%) of the labeler by changing the ratio of the redundant examples.

Fig. 3 shows the results for various abstention percentages. From the figure, our algorithms ALa and ALw are consistently better than the baselines for abstention percentages above 40%. When a good estimate of r^* is available, our algorithms perform better than all the other algorithms for abstention percentages above 30%. This shows the advantage of modeling the labeler's abstention pattern in this setting, especially for medium to high abstention percentages.

We also consider the setting where the abstention percentage is fixed at 40% and the pool size varies between 400 and 1800. Fig. 4 shows the AUC curves obtained in this setting. From the figure, our algorithms ALa and ALw are significantly better than the baselines for all pool sizes. For relatively larger pool sizes (1200 and above), ALa and ALw are also comparable to the AUC curves obtained when good estimations of r^* are available.

6.2. Abstention on easy examples

In this scenario, we test with the labeler who abstains from labeling easy data, which are far from the true decision boundary. This setting may seem counter-intuitive, but it is in fact not unrealistic. For example in the learning with corrupted labels setting discussed in Section 1, easy examples may be considered less important than hard examples and thus were less protected than hard ones. In this case, an attacker may attempt to corrupt the labels of those easy examples to bring down the performance of the learned classifier. Furthermore, under a heavy attack, we may expect a high abstention percentage. As another example, in medical diagnosis, lung cancer screening is only recommended for the high-risk group (heavy smoking, 55–74 years old, etc.) [33], so labels (cancer or no cancer) for the low-risk group (easy data) are often unavailable.

We simulate the abstention pattern for this scenario by first learning a logistic regression model with regularizer $\sigma^2 = 0.5$ on the whole training data set and then measuring the distance between the model's prediction probability to 0.5 for each example. The labeler would always abstain from labeling the subset of the training data (with size depending on the abstention percentage) that have the largest such distances while he would always give labels for the other examples. Fig. 5 (first row) shows the

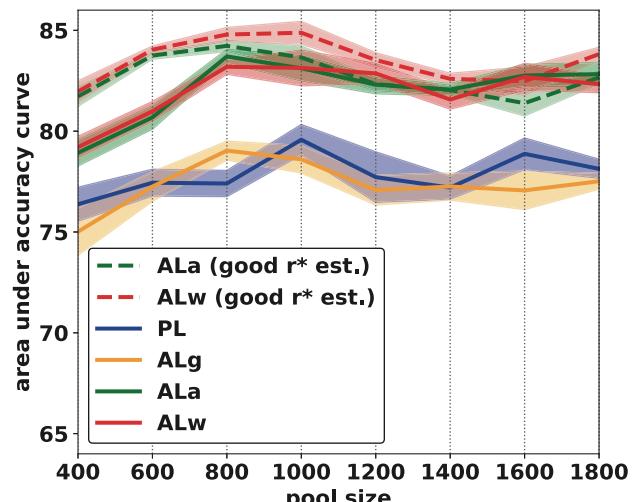


Fig. 4. AUC scores (with standard errors) vs. pool size for a labeler abstaining on examples unrelated to target task.

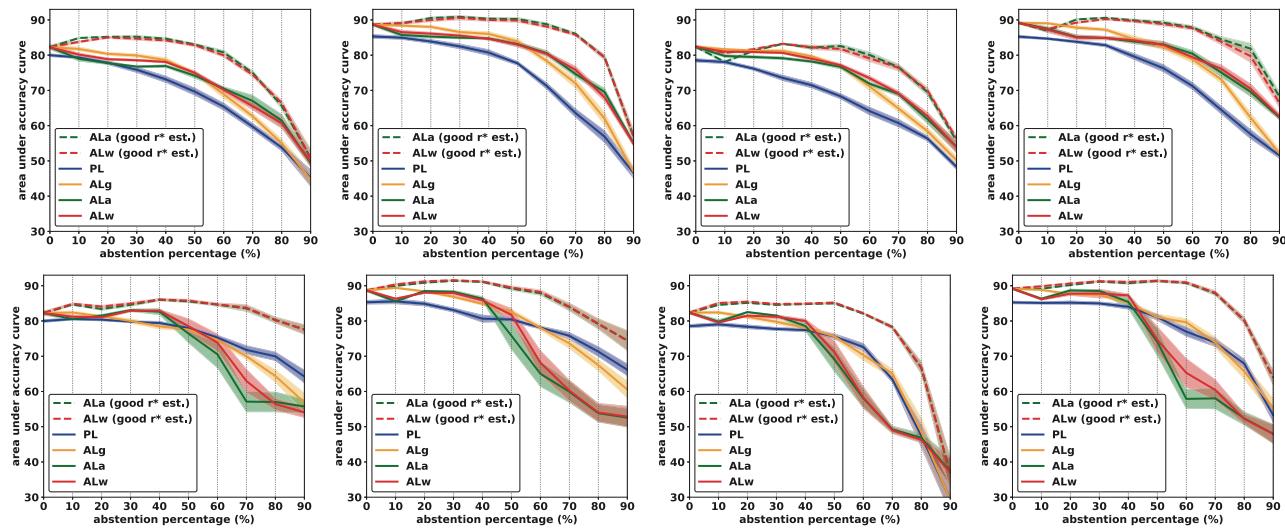


Fig. 5. AUC scores (with standard errors) for a labeler abstaining on easy examples (first row) and on hard examples (second row).

results for this setting on the following 4 binary text classification data sets from the 20 Newsgroups data (from left to right):

- comp.sys.mac.hardware/ comp.windows.x,
- rec.motorcycles/ rec.sport.baseball,
- sci.crypt/ sci.electronics,
- sci.space/ soc.religion.christian.

From the results, ALa and ALw have similar performance and work very well when the abstention percentage is above 50%. This shows that it is useful to learn and take into account the abstention probabilities when the abstention percentage is high (e.g., under heavy attacks), and our algorithms provide a good way to exploit this information. When the abstention percentage is small, the advantages of ALa and ALw diminish. This is expected because in this scenario, learning the abstention pattern is more expensive than simply ignoring it. However, when a good estimate of r^* is available, ALa and ALw perform better than all the other algorithms for most abstention percentages.

6.3. Abstention on hard examples

In this scenario, we test with the labeler who abstains from labeling hard data, which are near to the true decision boundary. This setting is common when the labeler wants to maximize the number of labels giving to the learner (e.g., in crowdsourcing where he is paid for each label provided). The abstention pattern in this experiment is generated similarly to the previous scenario, except that the labeler abstains from labeling the examples having the smallest distances above instead of those with the largest distances.

Fig. 5 (second row) shows the results for this scenario on the same 4 data sets above. These results suggest that this is a more difficult setting for active learning. From the figure, ALw performs better than ALa, but both of them are only better than the baselines when the abstention percentage is from 20–40%. For other abstention percentages, ALa, ALw, and ALg do not provide much advantage compared to PL. However, when a good estimate of r^* is available, ALa and ALw perform very well and are better than all the other algorithms.

Summary: The results above have shown that the proposed algorithms are useful for pool-based active learning with absten-

tion feedbacks when the abstention percentage is within an appropriate range that depends on the problem. The algorithms are especially useful when a good estimate of the abstention rate r^* is available. In practice, this estimate can be pre-computed from previous interactions between the learning systems and the labeler (e.g., using previous labeling preferences of the labeler), and then inputted into our algorithms as the priors $p_0[r]$. During the execution of our algorithms, this estimate will be gradually improved.

Note that in some of our plots (see Figs. 3 and 5), the performance of our algorithms decreases at abstention percentages around 10–20% and then increases for those within 30–50% before decreasing again. One possible explanation for this behavior is that our algorithms rarely observe rejected examples at a low abstention percentage, and thus it is relatively harder for the algorithms to learn the abstention rate r^* accurately, leading to a sub-optimal selection of queries.

7. Conclusion

We proposed two new greedy algorithms under the Bayesian active learning with abstention feedbacks framework. This framework is useful in many real-world scenarios, including learning from multiple labelers and under corrupted labels. We proved that the algorithms have theoretical guarantees in the average and worst cases and empirically showed that they are useful for classification, especially when a good estimate of the abstention rate is available. Our results suggest that keeping track and learning the abstention patterns of labelers are important for active learning with abstention feedbacks in practice.

CRediT authorship contribution statement

Cuong V. Nguyen: Conceptualization, Methodology, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Lam Si Tung Ho:** Conceptualization, Investigation, Writing – original draft, Writing – review & editing. **Huan Xu:** Conceptualization, Methodology, Writing – review & editing. **Vu Dinh:** Conceptualization, Investigation, Writing – original draft, Writing – review & editing. **Binh T. Nguyen:** Formal analysis, Investigation, Writing – original draft, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

LSTH was supported by startup funds from Dalhousie University, the Canada Research Chairs program, the NSERC Discovery Grant RGPIN-2018-05447, and the NSERC Discovery Launch Supplement DGEGR-2018-00181. BTN was supported by Vietnam National University Ho Chi Minh City (VNU-HCM) under Grant No. NCM2019-18-01 throughout this paper. VD was supported by a startup fund from University of Delaware and (US) National Science Foundation grant DMS-1951474.

Appendix A. Proofs

A.1. Proof of Theorem 5.1

To prove this theorem, we first apply Theorem 5.1 in [6]. This requires us to prove that the utility function $g(S, (f, k))$ is adaptive monotone and adaptive submodular with respect to the joint prior distribution q_0 . Note that $g(S, (f, k))$ is the version space reduction function with respect to the joint prior q_0 on the joint space $\mathcal{F} \times \mathcal{K}$. From the results in Section 9 of [6], version space reduction functions are adaptive monotone and adaptive submodular with respect to the corresponding prior. Thus, the utility function $g(S, (f, k))$ is adaptive monotone and adaptive submodular with respect to the joint prior q_0 .

With the above properties of g , applying Theorem 5.2 in [6], we have:

$$G_{\text{avg}}(\pi_{\text{greedy}}) > (1 - 1/e)G_{\text{avg}}(\pi_{\text{avg}}^*),$$

where π_{greedy} is the greedy algorithm that selects the examples maximizing the expected utility gain at each step. From the proof of Theorem 5.2 of [7], this greedy algorithm is equivalent to the maximum Gibbs error algorithm that selects the examples according to the criterion:

$$x^* = \arg \max_{x \in \mathcal{X}} \{1 - p_{i-1}[Z = 1; x]^2 - \sum_{y \in \mathcal{Y}} p_{i-1}[Y = y \wedge Z = 0; x]^2\}, \quad (8)$$

where p_{i-1} is the current posterior distribution, Y is the random variable for the label of x , and Z is the random variable for the abstention pattern of x . To understand this equation, we can think of the considered problem as a classification problem with labels $(y, z = 0)$ or $(z = 1)$, where $(y, z = 0)$ indicates an example is labeled the label y and $(z = 1)$ indicates an example is not labeled.

Since y and z are independent, Eq. (8) is equivalent to:

$$\begin{aligned} x^* &= \arg \max_{x \in \mathcal{X}} \{1 - p_{i-1}[Z = 1; x]^2 - \\ &\quad \sum_{y \in \mathcal{Y}} (p_{i-1}[Z = 0; x]p_{i-1}[Y = y; x])^2\} \\ &= \arg \max_{x \in \mathcal{X}} \{1 - p_{i-1}[Z = 1; x]^2 - \\ &\quad p_{i-1}[Z = 0; x]^2 \sum_{y \in \mathcal{Y}} p_{i-1}[Y = y; x]^2\}. \end{aligned}$$

We also have:

$$\begin{aligned} p_{i-1}[Z = 1; x] &= \int p_{i-1}[r]r(x)dr \\ &= \mathbb{E}_{r \sim p_{i-1}}[r(x)] \\ &= \tilde{r}(x). \end{aligned}$$

Similarly, $p_{i-1}[Z = 0; x] = 1 - \tilde{r}(x)$.

Hence, the previous equation is equivalent to:

$$x^* = \arg \max_{x \in \mathcal{X}} \{1 - \tilde{r}(x)^2 - (1 - \tilde{r}(x))^2 \sum_{y \in \mathcal{Y}} p_{i-1}[Y = y; x]^2\},$$

which is Eq. (3). Therefore, the average-case BALAF algorithm is equivalent to π_{greedy} and Theorem 5.1 holds.

A.2. Proof of Theorem 5.2

To prove this theorem, we first apply Theorem 3 in [8]. This requires us to prove that the utility $g(S, (f, k))$ is pointwise monotone and pointwise submodular. Note that $g(S, (f, k))$ is the version space reduction function with respect to the joint prior q_0 on the joint space $\mathcal{F} \times \mathcal{K}$. From the proof of Theorem 5 in [8], version space reduction functions are both pointwise monotone and pointwise submodular. Thus, $g(S, (f, k))$ is pointwise monotone and pointwise submodular.

With the above properties of g , applying Theorem 3 in [8], we have:

$$G_{\text{worst}}(\pi_{\text{greedy}}) > (1 - 1/e)G_{\text{worst}}(\pi_{\text{worst}}^*),$$

where π_{greedy} is the greedy algorithm that selects the examples maximizing the worst-case utility gain at each step. From the proof of Theorem 5 of [8], this greedy algorithm is equivalent to the least confidence algorithm that selects the examples according to the criterion:

$$\begin{aligned} x^* &= \arg \min_{x \in \mathcal{X}} \{\max\{p_{i-1}[Z = 1; x], \\ &\quad \max_{y \in \mathcal{Y}} p_{i-1}[Y = y \wedge Z = 0; x]\}\}, \end{aligned} \quad (9)$$

where p_{i-1} is the current posterior distribution, Y is the random variable for the label of x , and Z is the random variable for the abstention pattern of x . Similar to the proof of Theorem 5.1 above, to understand this equation, we can think of the considered problem as a classification problem with labels $(y, z = 0)$ or $(z = 1)$, where $(y, z = 0)$ indicates an example is labeled the label y and $(z = 1)$ indicates an example is not labeled.

Since y and z are independent, Eq. (9) is equivalent to:

$$\begin{aligned} x^* &= \arg \min_{x \in \mathcal{X}} \{\max\{p_{i-1}[Z = 1; x], \\ &\quad \max_{y \in \mathcal{Y}} p_{i-1}[Y = y; x]p_{i-1}[Z = 0; x]\}\} \\ &= \arg \min_{x \in \mathcal{X}} \{\max\{p_{i-1}[Z = 1; x], \\ &\quad p_{i-1}[Z = 0; x]\max_{y \in \mathcal{Y}} p_{i-1}[Y = y; x]\}\}. \end{aligned}$$

From the proof of Theorem 5.1, we have $p_{i-1}[Z = 1; x] = \tilde{r}(x)$ and $p_{i-1}[Z = 0; x] = 1 - \tilde{r}(x)$.

Hence, the previous equation is equivalent to:

$$x^* = \arg \min_{x \in \mathcal{X}} \{\max\{\tilde{r}(x), (1 - \tilde{r}(x))\max_{y \in \mathcal{Y}} p_{i-1}[Y = y; x]\}\},$$

which is Eq. (6). Therefore, the worst-case BALAF algorithm is equivalent to π_{greedy} and Theorem 5.2 holds.

References

- [1] C. Zhang, K. Chaudhuri, Active learning from weak and strong labelers, in: Proceedings of the Conference on Neural Information Processing Systems, 2015.
- [2] S. Yan, K. Chaudhuri, T. Javidi, Active learning from noisy and abstention feedback, in: Proceedings of the Annual Allerton Conference on Communication, Control, and Computing, 2015.
- [3] M. Fang, X. Zhu, C. Zhang, Active learning from oracle with knowledge blind spot, in: Proceedings of the AAAI Conference on Artificial Intelligence, 2012.
- [4] M.E. Ramirez-Loaiza, A. Culotta, M. Bilgic, Anytime active learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, 2014.
- [5] M. Zhao, B. An, W. Gao, T. Zhang, Efficient label contamination attacks against black-box learning models, in: Proceedings of the International Joint Conferences on Artificial Intelligence, 2017.

- [6] D. Golovin, A. Krause, Adaptive submodularity: Theory and applications in active learning and stochastic optimization, *J. Artif. Intell. Res.*
- [7] N.V. Cuong, W.S. Lee, N. Ye, K.M.A. Chai, H.L. Chieu, Active learning for probabilistic hypotheses using the maximum Gibbs error criterion, in: *Proceedings of the Conference on Neural Information Processing Systems*, 2013.
- [8] N.V. Cuong, W.S. Lee, N. Ye, Near-optimal adaptive pool-based active learning with general loss, in: *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, 2014.
- [9] D.D. Lewis, W.A. Gale, A sequential algorithm for training text classifiers, in: *Proceedings of the Annual ACM SIGIR Conference on Research and Development in Information Retrieval*, 1994.
- [10] Y. Chen, A. Krause, Near-optimal batch mode active learning and adaptive submodular optimization, in: *Proceedings of the International Conference on Machine Learning*, 2013.
- [11] Y. Chen, S.H. Hassani, A. Karbasi, A. Krause, Sequential information maximization: when is greedy near-optimal?, in: *Proceedings of the Conference on Learning Theory*, 2015.
- [12] N.V. Cuong, H. Xu, Adaptive maximization of pointwise submodular functions with budget constraint, in: *Proceedings of the Conference on Neural Information Processing Systems*, 2016.
- [13] N.V. Cuong, N. Ye, W.S. Lee, Robustness of Bayesian pool-based active learning against prior misspecification, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2016.
- [14] S. Shekhar, M. Ghavamzadeh, T. Javidi, Active learning for classification with abstention, *IEEE J. Select. Areas Inform. Theory* 2 (2) (2021) 705–719.
- [15] Y. Chen, S.H. Hassani, A. Krause, Near-optimal Bayesian active learning with correlated and noisy tests, in: *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2017.
- [16] C. Cortes, G. DeSalvo, M. Mohri, Boosting with abstention, *Adv. Neural Inform. Process. Syst.* 29 (2016) 1660–1668.
- [17] C. Cortes, G. DeSalvo, C. Gentile, M. Mohri, S. Yang, Online learning with abstention, *International Conference on Machine Learning*, *PMLR* (2018) 1059–1067.
- [18] P. Donmez, J.G. Carbonell, Proactive learning: Cost-sensitive active learning with multiple imperfect oracles, in: *Proceedings of the Conference on Information and Knowledge Management*, 2008.
- [19] D. Golovin, A. Krause, D. Ray, Near-optimal Bayesian active learning with noisy observations, in: *Proceedings of the Conference on Neural Information Processing Systems*, 2010.
- [20] M. Naghshvar, T. Javidi, K. Chaudhuri, Noisy Bayesian active learning, in: *Proceedings of the Annual Allerton Conference on Communication, Control, and Computing*, 2012.
- [21] E. Ni, C. Ling, Active Learning with c-Certainty, *Advances in Knowledge Discovery and Data Mining*.
- [22] L. Malago, N. Cesa-Bianchi, J. Renders, Online active learning with strong and weak annotators, in: *NIPS Workshop on Learning from the Wisdom of Crowds*, 2014.
- [23] Y. Yan, G.M. Fung, R. Rosales, J.G. Dy, Active learning from crowds, in: *Proceedings of the International Conference on Machine Learning*, 2011.
- [24] L. Zhao, G. Sukthankar, R. Sukthankar, Incremental relabeling for active learning with noisy crowdsourced annotations, *PASSAT and SocialCom* (2011).
- [25] B. Mozafari, P. Sarkar, M. Franklin, M. Jordan, S. Madden, Scaling up crowdsourcing to very large datasets: a case for active learning, *PVLDB*.
- [26] E. Manino, L. Tran-Thanh, N.R. Jennings, Efficiency of active learning for the allocation of workers on crowdsourced classification tasks, in: *NIPS CrowdML Workshop*, 2016.
- [27] A. Singla, S. Tsiatischek, A. Krause, Noisy submodular maximization via adaptive sampling with applications to crowdsourced image collection summarization, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2016.
- [28] Y. Yang, Z. Ma, F. Nie, X. Chang, A.G. Hauptmann, Multi-class active learning by uncertainty sampling with diversity maximization, *Int. J. Comput. Vision* 113 (2) (2015) 113–127.
- [29] Y. Yan, F. Nie, W. Li, C. Gao, Y. Yang, D. Xu, Image classification by cross-media active learning with privileged information, *IEEE Trans. Multimedia* 18 (12) (2016) 2494–2502.
- [30] B. Settles, Active learning literature survey, *Computer Sciences Technical Report* 1648, University of Wisconsin-Madison, 2010.
- [31] B. Settles, M. Craven, An analysis of active learning strategies for sequence labeling tasks, in: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2008.
- [32] T. Joachims, A probabilistic analysis of the Rocchio algorithm with TFIDF for text categorization, *Tech. rep., DTIC Document* (1996).
- [33] H. Roberts, C. Walker-Dilks, K. Sivjee, Y. Ung, K. Yasufuku, A. Hey, N. Lewis, L.C. S.G.D. Group, Screening high-risk populations for lung cancer: guideline recommendations, *J. Thor. Oncol.*



Cuong V. Nguyen is an Assistant Professor of the Knight Foundation School of Computing and Information Sciences, Florida International University (FIU). Before joining FIU, he worked at Amazon Web Services (AWS) and Cambridge University. His research interests include probabilistic machine learning and artificial intelligence.



Dr. Lam Ho is currently an Assistant Professor and a Canada Research Chair (Tier 2) in Stochastic Modelling at the Department of Mathematics and Statistics, Dalhousie University. He received his Ph.D. in Statistics from the University of Wisconsin-Madison in 2014. Before joining Dal, he worked as a postdoctoral researcher at the University of California, Los Angeles. Dr. Ho is interested in statistical theory and methods for stochastic models.



Huan Xu is a senior staff algorithm engineer at the Data Technology and Product Business Unit, Alibaba. Before joining Alibaba, he was an assistant professor at the Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology (Gatech). Prior to Gatech, he was an assistant professor and then associate professor with tenure at the National University of Singapore. His research area is machine learning and operations research.



Vu Dinh is an Assistant Professor at the Department of Mathematical Sciences, University of Delaware. His research focuses on theoretical foundations and applications of machine learning algorithms in applied sciences. Before joining University of Delaware, Vu Dinh obtained his PhD in 2014 from Purdue University on computational methods for experimental design and control of biological systems. From 2015–2017, he was a post-doctoral research fellow in Matsen Group at Fred Hutchinson Cancer Research Center (Seattle, WA).



Dr. Binh Nguyen is currently the Head of Department of Computer Science, Faculty of Mathematics and Computer Science at University of Science, Vietnam National University in Ho Chi Minh City, Vietnam. In 2012, he completed his Ph.D. at Ecole Polytechnique, France. He has had 55 publications in Scientific Computing, Machine Learning, and Computer Vision.