Teaching an Active Learner with Contrastive Examples

Anonymous Author(s)
Affiliation
Address
email

Abstract

We study the problem of active learning with the added twist that the learner is assisted by a helpful teacher. We consider the following natural interaction protocol: At each round, the learner proposes a query asking for the label of an instance x^q, the teacher provides the requested label $\{x^q, y^q\}$ along with explanatory information to guide the learning process. In this paper, we view this information in the form of an additional *contrastive example* ($\{x^c, y^c\}$) where x^c is picked from a set constrained by x^q (e.g., dissimilar instances with the same label, or similar instances with opposite labels). Our focus is to design a teaching algorithm that can provide an informative sequence of contrastive examples to the learner to speed up the learning process. We show that this leads to a challenging sequence optimization problem where the algorithm's choices at a given round depend on the history of interactions. We investigate an efficient teaching algorithm that adaptively picks these contrastive examples. We derive strong performance guarantees for our algorithm based on two problem-dependent parameters and further show that for specific types of active learners (e.g., a generalized binary search learner), the proposed teaching algorithm exhibits strong approximation guarantees. Finally, we empirically verify our bounds and demonstrate the effectiveness of our teaching framework via two numerical case studies.

1 Introduction

Active learning characterizes a natural learning paradigm where a learner *actively engages* in the learning process, and has demonstrated great success in both the educational domain (Prince, 2004) and the machine learning literature (Settles, 2012). In the context of machine learning, active learning has traditionally been studied as the *active instance labeling* problem, where the goal of an active learner is to learn a concept by selectively querying the labels of a sequence of data points (Kosaraju et al., 1999; Dasgupta, 2005; Nowak, 2008; Balcan et al., 2009; Gonen et al., 2013; Hanneke and Yang, 2014; Yan and Zhang, 2017). Recently, active learning has been investigated under richer interaction protocols, including learning from feature feedback (Poulis and Dasgupta, 2017; Dasgupta and Sabato, 2020) or from demonstrations (Chernova and Thomaz, 2014; Silver et al., 2012). Typically, these works assume that the rich feedback obtained from the teacher is adversarial (Dasgupta and Sabato, 2020) or stochastic (Poulis and Dasgupta, 2017). While these results have shown promising improvements in sample complexity by leveraging more expressive feedback, the assumptions of an adversarial or stochastic teacher could fall short in characterizing many real-world interactive learning systems (e.g. automated tutoring or multi-agent learning systems) where the teacher cooperates with the learner to achieve the learning objective (Zilles et al., 2011; Zhu et al., 2018).

Motivated by real-world cooperative learning and teaching scenarios (Chen et al., 2018a; Teso and Kersting, 2019), we consider an optimistic yet natural interaction protocol between an active learner

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

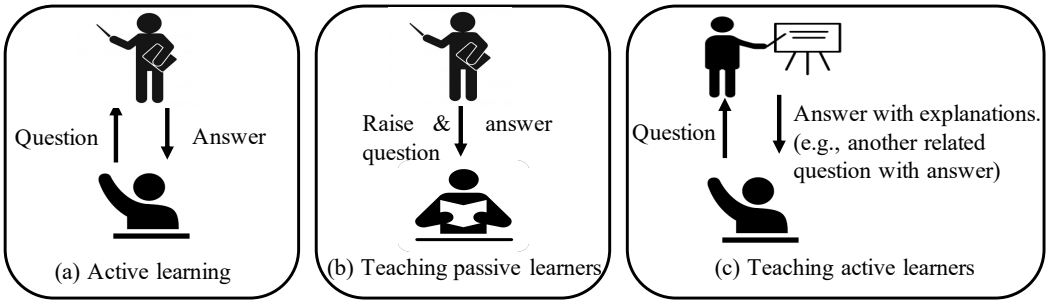

Figure 1: The difference between (a) active learning, (b) classic machine teaching and (c) our problem.

and a *helpful* teacher (see Figure 1 (c)): At each round, the learner proposes a query asking for the label of an instance x^q; the teacher then provides the requested label $\{x^q, y^q\}$ along with explanatory information to guide the learning process. In this paper, we consider this information in the form of an additional *contrastive example* ($\{x^c, y^c\}$) where x^c is picked from a set constrained by x^q (e.g., dissimilar instances with the same label, or similar instances with opposite labels). In contrast to existing work in active learning with contrastive examples (Abe et al., 2006; Liang et al., 2020; Ma et al., 2021) or with comparison queries (Jamieson and Nowak, 2011; Kane et al., 2017) which focus on developing *active learning algorithms*, we assume that the learner's querying and learning strategy is *fixed* and *known* to the teacher, and focus on designing a *teaching algorithm* that can provide an informative sequence of contrastive examples to the learner to speed up the learning process.

Note that in the absence of the learner's query, this problem reduces to the classical *machine teaching* problem (Goldman and Kearns, 1995). When teaching a (passive) version space learner (Goldman and Kearns, 1995), classical machine teaching can be viewed as a set cover problem, and the greedy algorithm is known to compute a cover that is within a logarithmic factor of the minimum cover (Chvatal, 1979). However, as we show in section 3.2, the fact that an active learner actively trying to reduce the label cost could actually *hurt* the performance of a greedy teacher.

Given an active version space learner specified by an acquisition function, we show that the design of an optimal teaching algorithm induces a challenging sequence optimization problem, where the teacher's available choices of constrastive examples at a given round depend on the history of interactions. We study the greedy teacher that adaptively picks contrastive examples in response to the learner's queries. Specifically, we introduce several natural problem-dependent parameters to characterize (1) the constraints imposed by the active learner and (2) the structure of the coverage function resulted from the hypothesis space. Based on these characterizations, we explore the theoretical properties for the greedy teacher, and demonstrate the power and limitations of the interactive teaching framework. Our main technical contributions are summarized as follows:

I. We consider a novel interactive protocol for teaching an active learner, and derive a general performance guarantee for the greedy teacher, regardless of the choice of acquisition function by the active learner. Our result generalizes existing near-optimality guarantees for the greedy teacher under the classical teaching (Goldman and Kearns, 1995) to the more challenging sequence optimization problem (section 3).

II. We then show that, perhaps surprisingly, a greedy teacher may fail badly when interacting with an active learner compared to the optimal teacher, due to the constrained choices of contrastive examples imposed by the learner's query (section 3). In fact, there exist pessimistic scenarios where an active learner alone is better off without the help of a greedy teacher (section 4).

III. When interacting with a generalized binary search (GBS) learner (Dasgupta, 2005; Nowak, 2008), we provide a strong performance guarantee for the greedy teacher, and show that it cannot perform arbitrarily badly (section 4). Our bound establishes a rigorous connection between the problem-dependent parameters introduced in section 3 to the classical *k-neighborly condition* and *coherence parameter* which are widely used in the active learning literature (Nowak, 2008, 2011; Mussmann and Liang, 2018).

IV. We empirically verify our bounds via two numerical case studies. These examples illustrate the intuitive sequential teaching process and demonstrate the effectiveness of our teaching framework.

1.1 Related Works

Teaching a sequential learner Machine teaching has shown significant development in the past two decades in both theoretical and practical domains. In a quest to achieve rich teach-learner interactions, various different models for machine teaching have been proposed under the sequential teaching setting (e.g., local preference-based model for version space learners (Chen et al., 2018b), models for gradient learners (Liu et al., 2017, 2018; Kamalaruban et al., 2019), models inspired by control theory (Zhu, 2018; Lessard et al., 2019), models for sequential tasks (Cakmak and Lopes, 2012; Haug et al., 2018; Tschiatschek et al., 2019), and models for human-centered applications that require adaptivity (Singla et al., 2013; Hunziker et al., 2019)). Compared to our setting, in these works, the teacher is responsible for designing all of the learning data. Recently, Peltola et al. (2019) investigate the task of teaching sequential learners such as multi-armed bandits or active learners where the learners also make query, but they formulate the problem as a variant of reward shaping, and allow the teacher to provide teaching examples that are inconsistent with the true data distribution. While our work falls into the category of teaching sequential learners, it could be viewed as a form of explanation-based machine teaching, and hence inherits different teaching constraints and objectives from the above settings.

Explanation-based teaching and learning Explanation-based teaching captures a rich class of teaching protocols where labels are coupled with additional information (such as highlighting regions or features on an image). Having an additional mode in the feedback allows the learner to perceive more intuitive explanations of the target hypothesis, and hence dramatically improves the learner's ability to learn a new concept. Explanations has been shown to be effective in helping a human student to improve classification performance by steering the student's attention (Grant and Spivey, 2003; Roads et al., 2016; Chen et al., 2018a; Mac Aodha et al., 2018); likewise, it also plays an important role in helping machine learners to efficiently learn some concept classes that would otherwise be intractable (Poulis and Dasgupta, 2017; Dasgupta and Sabato, 2020). Our work is motivated by the success of explanation-based learning, and consider contrastive examples as a simple yet effective form of explanations. In comparison with existing work on explanation-based teaching, our work has an additional twist featured by the interaction with an active learner.

Submodular set and sequence optimization Our theoretical framework in section 3 is inspired by recent results on sequence submodular function maximization (Zhang et al., 2015; Tschiatschek et al., 2017; Hunziker et al., 2019). In particular, Zhang et al. (2015) introduced the notion of string submodular functions, which, analogous to the classical notion of submodular set functions (Krause and Golovin, 2014), enjoy similar performance guarantees for maximization of deterministic sequence functions. Our setting has three key differences in that (1) we focus on a constrained planning problem where the set of available actions is dynamic (i.e. depending on the learner's query), (2) our objective function no longer exhibits diminishing returns properties and (3) we consider a min-cost coverage problem as opposed to the budgeted maximization problem.

2 Active Learning with a Teacher: Problem Formulation

In this section, we formally introduce the new problem of active learning with a helpful teacher along with the corresponding interaction protocol. Then, we formulate the problem of teaching an active learner as a constrained sequential optimization problem. Throughout the entire paper, we consider binary classifiers, of which the output is either $+1$ or -1 (i.e., $h(x) \in \{+1, -1\}$), for simplicity.

2.1 Learner-Teacher Interaction Protocol

We study the problem of teaching an active version space learner with a teacher. In contrast to the standard active learning protocol, our framework not only allows the learner to ask questions but also enable the teacher to provide additional explanatory information to guide the learning process. Specifically, we model the explanatory information using a contrastive example picked from a set contrained by the learner's query (e.g., dissimilar instances with the same label, or similar instances with opposite labels). The full interaction protocol is presented in Protocol 1. At each round t, the learner asks a query $x_t^{\mathrm{q}} = q(x_{1:t-1}^{\mathrm{c}}, x_{1:t-1}^{\mathrm{q}}, H_{t-1}^{q})$ with query function q conditioning on the received contrastive examples ($x_{1:t-1}^{\mathrm{c}}$) and version space H_{t-1}^{q}. Then the teacher return the ground-truth label

Protocol 1 An interaction protocol between the teacher and the active learner

An active learner with a query function q; Initial version space $H_0^q = \mathcal{H}$; Constraint function ξ for the contrastive example; Target hypothesis $h^\star$.
for $t = 1, 2, 3, \ldots$ **do**
  learner asks a query $x_t^{\mathrm{q}} = q(x_{1:t-1}^{\mathrm{c}}, x_{1:t-1}^{\mathrm{q}}, H_{t-1}^q)$
  teacher provides the label y_t^q for the query x_t^{q}
  teacher provides a contrastive example $(x_t^{\mathrm{c}}, y_t^c)$ from a constrained set $\xi(x_t^{\mathrm{q}})$.
  if $H_t^q = \{h^\star\}$ **then** teaching process terminates.

y_t^{q} for x_t^{q} and also provides a contrastive example $(x_t^{\mathrm{c}}, y_t^{\mathrm{c}})$. Under this setting, the version space at iteration t with teaching sequence $x_{1:t}^{\mathrm{c}} = \{x_i^{\mathrm{c}}\}_{i=1}^t$, can be computed by (we omitted y for clarity)

$$H_t^q = H^q(x_{1:t}^{\mathrm{c}}) = \mathcal{H} \setminus \left(\bigcup_{i=1}^{t} S(x_i^{\mathrm{q}}) \cup S(x_i^{\mathrm{c}}) \right), \tag{1}$$

where $\mathcal{H}$ is the initial version space of the active learner and $S(x) = \{h \in \mathcal{H} | h(x) \neq h^\star(x)\}$ denotes the hypotheses covered by x (i.e., the hypotheses predict different labels with $h^\star$ on x).

2.2 Optimization Objective

Our goal is to reduce the learner's version space such that only $h^\star$ (the target hypothesis) is left with minimal cost (e.g., # of interaction rounds or # of samples). The definition of the cost can be problem-dependent. For the general case, we consider the following form of the cost function

$$J^q(x_{1:t}^{\mathrm{c}}) = \sum\nolimits_{i=1}^{t} c\left(x_i^{\mathrm{c}}\right) + c\left(x_i^{\mathrm{q}}\right), \tag{2}$$

where $x_{1:t}^{\mathrm{c}}$ is the teaching sequence given by the teacher, and $c : \mathcal{X} \to \mathbb{R}_+$ computes the cost of the example. We don't write the dependencies on learner's query sequence explicitly, as the queries can be induced by the teaching sequence and the query function q accordingly.

Given the active learner, we seek the optimal teaching sequence that achieves minimal cost by solving the following optimization problem.

$$\min_{X \subseteq \mathcal{X}^\star} J^q(X) \qquad s.t. \qquad H^q(X) = \{h^\star\} \quad \text{and} \quad x_t \in \xi(x_t^{\mathrm{q}}) \;\; \forall t \in [|X|], \tag{3}$$

where we use $\mathcal{X}^\star$ to denote the set of all possible sequences of teaching examples, and $[n]$ denotes all the integers in $[1, n]$. At every iteration, we can compute the constraint set by the function $\xi : \mathcal{X} \to 2^{\mathcal{X}}$ given the learner's query. Then, the teacher must select the contrastive example from the constrained set. This is a natural consideration, since without the constraint, the teacher might have selected some contrastive examples that are not interpretable to the learner. In practice, the constrained set could be the points in a hamming ball or a euclidean ball around the learner's query.

To be noted, since the order matters in our problem, it expands the search space exponentially. This makes solving the above sequence optimization problem is far more difficult than solving the corresponding set optimization problem (e.g., classical machine teaching problem), which is known to be NP-hard. Hence, we turn to approximation algorithms that with polynomial time complexity.

3 Greedy Teaching Algorithm and Analysis: A General Bound

In this section, we propose a simple yet effective greedy algorithm for solving the sequential optimization problem defined in the last section. Then, we further provide theoretical results to upper and lower bound the number of rounds of the interaction between the learner and the teacher or the sample complexity, for the general cases (i.e., we don't have any assumptions on the active learner).[1]

3.1 Greedy Teaching Policy

We now present a simple yet effective greedy teaching algorithm for teaching an active learner (See Algorithm 2). For clarity of presentation, we focus on unit cost case, i.e., $c(x) = 1$ for all x. For this

[1] All the proofs are deferred to the Appendix due to page limit.

Algorithm 2 Greedy teaching algorithm for active leaner

An active learner with a query function q; Initial version space $H_0^q = \mathcal{H}$; Constraint function ξ for the contrastive example; Target hypothesis $h^\star$.
for $t = 1, 2, 3, \ldots$ **do**
  learner asks a query $x_t^{\mathrm{q}} = q(x_{1:t-1}^{\mathrm{c}}, x_{1:t-1}^{\mathrm{q}}, H_{t-1}^q)$.
  teacher provides the label $y_t^q = h^\star(x_t^{\mathrm{q}})$.
  teacher provides a contrastive example $x_t^{\mathrm{c}} = \arg\max_{x \in \xi(x_t^{\mathrm{q}})} \Delta(x|x_{1:t-1}^{\mathrm{c}}, q)$ (equation 5).
  Update learner's version space by equation 1.
  if $H_t^q = \{h^\star\}$ **then** teaching process terminates.

case, we focus on minimizing the number of rounds of interactions between the teacher and active learner. The extension to a more general cost will be discussed in the Appendix.

Therefore, to obtain the teaching sequence, we will solve the following minimum-cost cover problem

$$\min_{X \in \mathcal{X}^\star} |X| \quad s.t. \quad H^q(X) = \{h^\star\} \quad \text{and} \quad x_t \in \xi(x_t^{\mathrm{q}}) \;\; \forall t \in [|X|]. \tag{4}$$

For iteration $t+1$ with teaching history $x_{1:t}^c$, we define the marginal gain as the following

$$\Delta(x^c|x_{1:t}^c, q) = f(x_{1:t}^c \oplus x^c|q) - f(x_{1:t}^c|q) = |H^q(x_{1:t}^c) \setminus H^q(x_{1:t}^c \oplus x^c)|, \tag{5}$$

where $f(x_{1:t}|q) = |\mathcal{H} \setminus H^q(x_{1:t})|$ is the number of hypotheses removed from the version by the teaching examples $x_{1:t}$ together with the corresponding learner's queries (defined by equation 1), and $\oplus$ denotes the concatenation between two sequences. To be noted, when the function f is maximized, it is equivalent to that only the target hypothesis $h^\star$ is left in the version space.

For the greedy teaching algorithm, as described in Algorithm 2, at each round, it selects the contrastive example from the constrained set that maximizes the marginal gain defined in equation 5. Such process will be repeated until the function f is maximized or the budget is used up.

3.2 Theoretical Guarantees

We now turn to understand the theoretical performance of greedy algorithm. In general, for active learning, there is no guarantee for the sample complexity of arbitrary learners, and sometimes, the worst case complexity can be up to $\mathcal{O}(|\mathcal{H}|)$ (Dasgupta, 2005).

To understand the role of greedy teacher in active learning, we rely on the following two properties of the objective function. These two properties are usually used in the characterizing the submodularity or string submodularity of the underlying function (Zhang et al., 2015; Hunziker et al., 2019).

Definition 1 (**Pointwise submodularity ratio**). *For any sequence function f, the pointwise submodularity ratio with respect to any sequences σ, σ' and query function q is defined as*

$$\rho(\sigma, \sigma', q) = \min_{x \in \mathcal{X}} \frac{\Delta(x|\sigma, q)}{\Delta(x|\sigma \oplus \sigma', q)}. \tag{6}$$

Definition 2 (**Pointwise backward curvature**). *For any sequence function f, the pointwise backward ratio with respect to any sequences σ, σ' and query function q is defined as*

$$\gamma(\sigma, \sigma', q) = 1 - \frac{f(\sigma' \oplus \sigma|q) - f(\sigma|q)}{f(\sigma'|q) - f(\emptyset|q)}. \tag{7}$$

The submodularity ratio measures the degree of diminishing returns of the underlying sequence function by computing the ratio between the marginal gain. When the sequence function f is string/sequence submodular (Zhang et al., 2015; Alaei et al., 2010), the minimum pointwise submodularity ratio over all possible sequences is 1. Whereas the pointwsie backward curvature measures the degree of diminishing returns of the difference between marginal gains. For postfix monotone sequence functions, the pointwise backward curvature is smaller than 1 (Zhang et al., 2015).

In addition, at each round, the teacher can only select examples from a constrained set, which limits its power. To accommodate such limitations, we can view the teacher as an α-approximate greedy

teacher in the analysis. The parameter α characterizes the ratio between the marginal gain of a unconstrained greedy teacher and that of the constrained greedy teacher. We can compute it by

$$\alpha = \max_t \frac{\max_{x \in \mathcal{X}} \Delta(x|x^c_{1:t-1}, q)}{\max_{x \in \xi(x^q_t)} \Delta(x|x^c_{1:t-1}, q)}. \tag{8}$$

The following theorem summarizes our first theoretical result, which provides a problem-dependent upper bound for the sample complexity of greedy teaching algorithms.

Theorem 1. *The sample complexity of α-approximate greedy algorithm with any active learner is at most*

$$\left(\frac{\alpha \cdot \mathcal{O}\left(\log|\mathcal{H}| \cdot \log\left(|\mathcal{H}|/\gamma^g\right)\right)}{\rho^g \gamma^g \log(\gamma^g/(\gamma^g - 1))} + \frac{\alpha \cdot \mathcal{O}\left(\log\left(|\mathcal{H}|/\gamma^g\right)\right)}{\rho^g \gamma^g} \right) \cdot \mathtt{OPT}^{\mathtt{T+AL}}, \tag{9}$$

where γ^g and ρ^g are computed with respect to the greedy teaching sequence $x^c_{1:T}$ (let T be the length of the greedy teaching sequence), the active learner and the optimal teaching sequence σ by,

$$\rho^g = \min_{i,j \geq 0} \rho(x^c_{1:i}, \sigma_{1:j}, q), \quad \gamma^g = \max_{i \geq 1} \gamma(\sigma, x^c_{1:i}, q), \tag{10}$$

and $\mathtt{OPT}^{\mathtt{T+AL}}$ is the optimal sample complexity for teaching the corresponding active learner without the constraints on the contrastive examples (it also holds when with constraints).

Remark 1. *For any active learner, $1 \leq \gamma^g \leq |\mathcal{H}| - 1$ and $\rho^g \leq 1$. When $\gamma \to 1$, the l.h.s. term in the parentheses of equation 9 will vanish. Whereas when $\gamma^g \to |\mathcal{H}|$, the r.h.s. term will vanish.*

Theorem 1 bounds the worst case sample complexity of greedy teaching algorithms. To be noted, when $\gamma^g \to 1$, the r.h.s. term in equation 9 will dominate, whereas when $\gamma^g \to |\mathcal{H}|$, the l.h.s. term will become dominating. In the next, we show that the linear dependency on α cannot be avoided, showing that Theorem 1 is non-vacuous.

Theorem 2. *There exists constraint function ξ, version space $\mathcal{H}$ and active learner with query function q, such that the sample complexity of the greedy teacher is at least $\Omega(\sqrt{|\mathcal{H}|}) \cdot \mathtt{OPT}^{\mathtt{T+AL}}$.*

Theorem 2 shows that the linear dependency on α cannot be improved when without any assumptions on the constraint function ξ, hypotheses space and the active learner. To show this, we constructed an special case to demonstrate Theorem 2 in the Appendix. In the example, the optimal teacher only uses a constant number of examples to teach the active learner, while the greedy teacher needs to provide all the examples, whose number is of the order $\Omega(\sqrt{|\mathcal{H}|})$.

4 Teaching Greedy Active Learners

In this section, we consider the problem of teaching an greedy active learner, concretely, the Generalized Binary Search (GBS) (Nowak, 2008), which selects the query in a greedy manner (i.e., the one that maximizes the utility). At each iteration $t+1$ with the version space H_t, the GBS learner selects the query according to the expected number of hypotheses can be covered by the query (i.e., the utility of the query). Specifically, for binary classification problem, for each query x^q_{t+1}, the utility can be computed as

$$u(x^q_{t+1}|H_t) = \frac{2}{|H_t|} \cdot \left(\sum_{h \in H_t} \mathbb{1}[h(x) = -1] \right) \cdot \left(\sum_{h \in H_t} \mathbb{1}[h(x) = 1] \right). \tag{11}$$

The above utility function measures how discriminative the selected query is. GBS selects the query that maximizes the utility, and hence it lower bounds the minimal number of hypotheses to be removed after the query. However, as pointed out in Dasgupta (2005), the worst case sample complexity for GBS learner is $\mathcal{O}(|\mathcal{H}|)$. However, for practical problems, there are usually some structures in the hypotheses space and the sample space we can exploit to get a tighter bound. Specifically, we consider the problems with the k-neighborly structure (Nowak, 2008; Mussmann and Liang, 2018).

Definition 3 (k**-neighborly**). *We define the following graph (V, E). The vertices $V = \mathcal{X}$, and the edges $E = \{(x, x')|d_{\mathcal{H}}(x, x') \leq k, \ \forall \ x, x' \in \mathcal{X}\}$, where $d_{\mathcal{H}}(x, x') = |\{h|h(x) \neq h(x'), \ \forall \ h \in \mathcal{H}\}|$. The query and hypotheses space $(\mathcal{X}, \mathcal{H})$ is k-neighborly if the induced graph is connected.*

Intuitively, the k-neighborly structure ensures that for any two points $x, x' \in \mathcal{X}$, there is a path formed by similar pairs connecting them. With such structures, Nowak (2011) further introduce the *coherence parameter*, which plays an important role in bounding the worst case complexity of GBS.

Definition 4 (**Coherence parameter**). *The coherence parameter for $(\mathcal{X}, \mathcal{H})$ is defined as*

$$c^\star(\mathcal{X}, \mathcal{H}) := \min_P \max_{h \in \mathcal{H}} \left| \int_{\mathcal{X}} h(x) dP(x) \right|, \tag{12}$$

where we minimize over all possible distribution on $\mathcal{X}$.

The coherence parameter $c^\star$ measures how evenly the query selected by GBS bisects the hypotheses. When with the above assumptions, Nowak (2008) shows that the sample complexity of GBS can be upper bounded by $\mathcal{O}(\log(|\mathcal{H}|)/\log(1/\alpha))$, where $\alpha = \max\{(1+c^\star)/2, (k+1)/(k+2)\}$. This implies that, the smaller the $c^\star$ and k, the more efficient the GBS is.

In the next, we seek to understand what's the sample complexity when with a greedy teacher? Can we also upper bound its sample complexity, even though the interaction between the learner and the teacher introduces extra complexity? The result is summarized in the following theorem.

Theorem 3. *For GBS learner, if $(\mathcal{X}, \mathcal{H})$ is k-neighborly and with coherence parameter $c^\star$, then the sample complexity of greedy teaching algorithm with any constraint function ξ is at most*

$$\frac{\alpha}{\epsilon} \cdot \mathcal{O}\left(\log^2 |\mathcal{H}|\right) \cdot \mathtt{OPT}^{\mathtt{T+AL}}, \tag{13}$$

where $\epsilon = \min\{(1-c^\star)/(1+c^\star), c^\star/(k-c^\star)\}$, and $\alpha \leq \max\{k/c^\star, 2/(1-c^\star)\}$.

Theorem 3 gives a near-optimal bound for the greedy teaching algorithm for GBS learners, which guarantees that its performance cannot be arbitrarily bad. However, the bound cannot be directly compared to that of GBS alone, as we don't know $\mathtt{OPT}^{\mathtt{T+AL}}$. Unfortunately, our following remark shows that the sample complexity of GBS with a greedy teacher is not guaranteed to be smaller compared to GBS alone, and for some extreme cases, it can be even larger.

Remark 2. *The sample complexity of GBS with greedy teacher (even without the constraints) is not guaranteed to be smaller than that of GBS alone.*

5 Numerical Experiments

In this section, we empirically evaluate the greedy teaching algorithm with β-greedy active learner (a rich class of active learners including GBS) in simulated environment. We seek to understand how does the type of active learners (i.e., varying β) and the constraint on contrastive examples affect the performance of the greedy teacher. To dig deeper, we further demonstrate how does these two factors affect the value of α and ρ^g, which plays an important role in our theoretical results.

5.1 Experimental Setup

β-greedy active learners: In the experiments, we consider the class of β-greedy active learners, of which the selected query x^q_{t+1} satisfies

$$u(x^q_{t+1}|H_t) \geq \frac{1}{\beta} \cdot \max_{x \in \mathcal{X}} u(x|H_t). \tag{14}$$

We can interpolate β from 1 to $+\infty$ to cover a broad range of active learners with different properties. Specifically, when $\beta = 1$, we recover the GBS active learner. Whereas when $\beta = +\infty$, we recover the random active leaner. In general, as β increases, the strategy becomes less greedy. For β-greedy active leaner, the query function q returns a query that satisfies the inequality (14).

Datasets: We conducted experiments on two datasets: `Vespula-Weevil` and `Butterfly-Moth`, which were adopted in Singla et al. (2014) for verifying the algorithm for teaching the crowd to classify images. Specifically, `Vespula-Weevil` is a synthesized dataset with 200 images, of which are generated by sampling a 2-d features from two bi-variate Gaussian distributions. Whereas `Butterfly-Moth` consists of 160 real images with two categories Butterfly and Moth. The 2-d embedding of each image is inferred from the annotation process using a generative Bayesian model proposed in Welinder et al. (2010). More details about the datasets can be found in the Appendix.

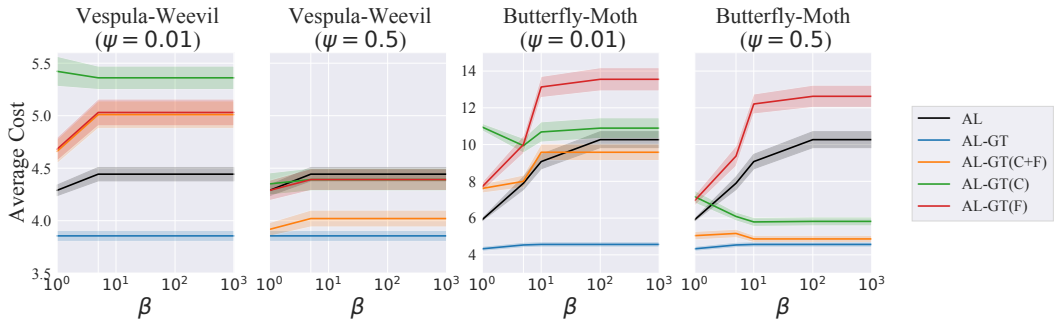

Figure 2: The average cost for different active learners and constraints. **AL** denotes the setting of active learner alone. **AL-GT** corresponds to active learner with an unconstrained greedy teacher, and **C, F, C+F** correspond to different constraints on the constrastive examples. $\psi \in [0, 1]$ is a parameter controlling the size of search space, which only affects **AL-GT(C), AL-GT(F)**, and **AL-GT(C+F)**. The shaded area is the one-standard error.

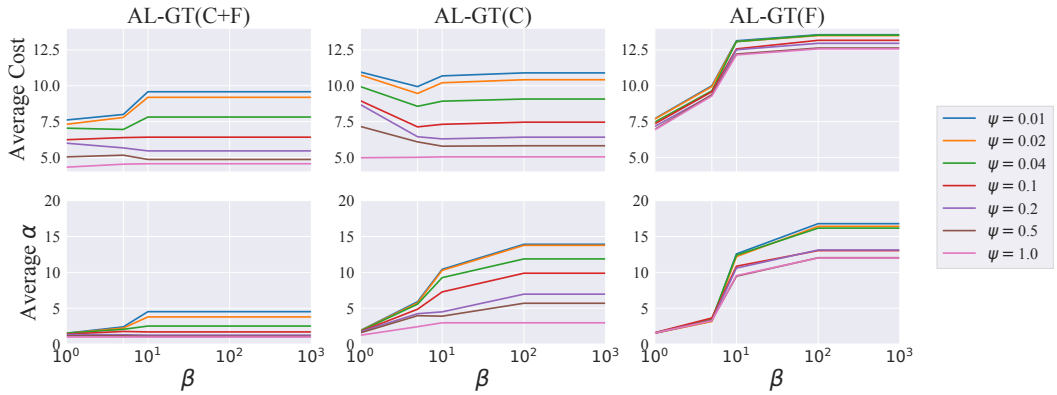

Figure 3: How do active learners (β) and constraints (ψ) affect the teaching cost (top row) and α (bottom row). The three columns are corresponding to the results of **AL-GT(C+F)**, **AL-GT(C)** and **AL-GT(F)** respectively. The results are obtained on the `Butterfly-Moth` dataset. We don't display the standard error for clarity.

Hypothesis space: Following Singla et al. (2014), for the `Vespula-Weevil` dataset, we generate the hypotheses by randomly sampling from the following eight multivariate Gaussian distributions, whose mean $\boldsymbol{\mu}_i = [\pi/4 \cdot i, 0]$, and covariance $\boldsymbol{\Sigma}_i = [2, 0; 0, 5e^{-3}]$ for i varying from 0 to 7. For the `Butterfly-Moth` dataset, the hypotheses are inferred from the annotation process with the same generative Bayesian model for inferring the 2-d embedding.

5.2 Effects of Active Learners and Constraints

We first study how does the active learner (AL) affects the teaching performance under different constraints. Specifically, we vary the β in $[1, 5, 10, 100, 1,000]$ to model different learners. In terms of the constraint on the teacher's example, we consider the following three types of the constraints: 1) examples that are close to the learner's query but with different label (denoted by **C**); 2) examples that are far away from the learner's query but with the same label(denoted by **F**); 3) the union of 1)& 2) (denoted by **C+F**). For each constraint, we also adopt a parameter $\psi \in [0, 1]$ to control the size of the search space (i.e., the closest or furthest ψ portion of the points with different or the same label as the query). We also include the unconstrained teacher, which can freely select examples.

To obtain more robust results, we iterate through all the hypotheses to serve as the target hypothesis once, and compute the averaged cost for identifying the target hypothesis in accompanying with the standard error. The results are presented in Figure 2. We can observe that the for both datasets, AL with greedy teacher (AL-GT, unconstrained) always achieves the best performance in terms the average cost. When β becomes larger, the average cost of AL-GT increase very slightly. For the others, when ψ is small, they usually perform worse than AL along. This is because when ψ is small, the candidate points in the constrained set are too restrictive and they are usually redundant to the query, making the contrastive example uninformative. Wheres when ψ is large enough, AL-GT(C+F) can significantly outperforms AL alone, showing that the teacher is helpful in this case. In summary,

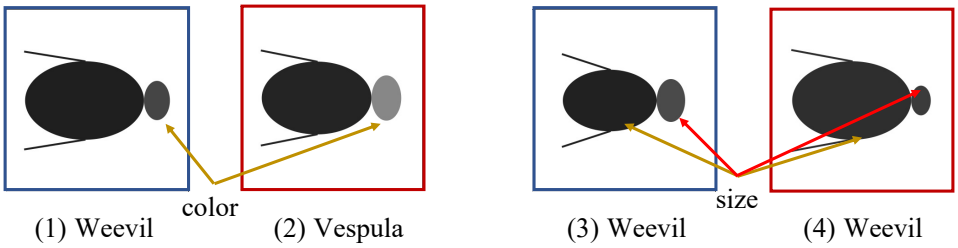

Figure 4: A visualization of a teaching sequence. (1)&(3) are learner's queries and (2)&(4) are the contrastive examples provided by the teacher (i.e., AL-GT(C+F)).

when β becomes large and the constraint is not too strict, the teacher will be more helpful (i.e., the gap between AL and AL with teacher becomes larger).

In the next, we further study how do the constraints and β affect the parameters α. We provide the results in Figure 3. In general, we can observe that as the constraints become stricter (i.e., ψ becomes smaller), α will correspondingly increase as well as the cost for teaching the active learner. This is consistent with our Theorem 1 and Theorem 3 in Section 3 and Section 4. In addition, when β becomes larger, α is also increasing, which is consistent with the trend in the plots of cost (top row in Figure 3). Besides, we also compute the empirical values of ρ^g, we find that for different β and constraints, ρ^g stays very close to 1. According our theoretical bounds, this may also explain why the trend in plots of cost is mostly consistent with that in the plots of α (bottom row in Figure 3).

5.3 Visualization of the Teaching Sequence

Lastly, we visualize a teaching sequence of AL-GT(C+F) in Figure 4. The results are obtained with $\beta = 1$ and $\psi = 0.5$ on `Vespula-Weevil` dataset. The first and third examples are the learner's queries and the second and fourth ones are the teacher's provided contrastive examples. In general, weevils have relatively short heads with similar color to the body. In contrast, vespula have big and contrasting heads. From Figure 4, we can observe that at the first round, the teacher selects an image of vespula with similar body size to the learner's query, but its head is in relatively larger and is in different color. This highlights the key features (i.e., the color and size of the head) to distinguishing weevil and vespula. In the next, the teacher provides another image of weevil which is visually very different from the query in terms the size of body and head, but are in the same category.

6 Conclusion

We introduced a novel interactive protocol for teaching an active leaner. We first derived a general performance guarantee for the greedy teacher with arbitrary active learners, and then provided a stronger performance bound for teaching GBS learners with greedy teacher under mild assumptions. In addition, we also demonstrated the limitations of greedy teachers by showing that 1) the performance of the greedy teacher can be arbitrarily worse than that of the optimal teacher due to the constrained choices of contrastive examples; and 2) the greedy teaching strategies for GBS learner is not guaranteed to be better than GBS alone. Lastly, we provided empirical results to demonstrate the effectiveness and limitations of the greedy teachers. We believe that our results have taken a significant step in bringing active learning and machine teaching closer to real-world settings in which the teacher collaborates with the learner to achieve the learning goal.

7 Social Impact

In this work, we provided an greedy teaching algorithm for active learners with theoretical guarantees. Our algorithm is derived for teaching an active leaner in a collaborative setting to achieve some learning goal, which can be potentially helpful for improving many real-world interactive learning systems, e.g., automated tutoring or multi-agent learning system.

However, just as a coin has two sides, our algorithm can be easily extended to poisoning the active learner by change the objective of the teacher. This can potentially threaten the safety of deploying machine learning models in online products.

References

Naoki Abe, Bianca Zadrozny, and John Langford. Outlier detection by active learning. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 504–509, 2006.

Saeed Alaei, Ali Makhdoumi, and Azarakhsh Malekian. Maximizing Sequence-Submodular Functions and its Application to Online Advertising. *arXiv e-prints*, art. arXiv:1009.4153, September 2010.

Maria-Florina Balcan, Alina Beygelzimer, and John Langford. Agnostic active learning. *J. Comput. Syst. Sci.*, 75(1):78–89, January 2009.

Maya Cakmak and Manuel Lopes. Algorithmic and human teaching of sequential decision tasks. In *AAAI*, 2012.

Yuxin Chen, Oisin Mac Aodha, Shihan Su, Pietro Perona, and Yisong Yue. Near-optimal machine teaching via explanatory teaching sets. In *Proc. International Conference on Artificial Intelligence and Statistics (AISTATS)*, April 2018a.

Yuxin Chen, Adish Singla, Oisin Mac Aodha, Pietro Perona, and Yisong Yue. Understanding the role of adaptivity in machine teaching: The case of version space learners. In *Advances in Neural Information Processing Systems*, pages 1476–1486, 2018b.

Sonia Chernova and Andrea L Thomaz. Robot learning from human teachers. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 8(3):1–121, 2014.

Vasek Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of operations research*, 4(3):233–235, 1979.

Sanjoy Dasgupta. Analysis of a greedy active learning strategy. In L. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems*, volume 17. MIT Press, 2005. URL https://proceedings.neurips.cc/paper/2004/file/c61fbef63df5ff317aecdc3670094472-Paper.pdf.

Sanjoy Dasgupta and Sivan Sabato. Robust learning from discriminative feature feedback. In *International Conference on Artificial Intelligence and Statistics*, pages 973–982. PMLR, 2020.

Sally A Goldman and Michael J Kearns. On the complexity of teaching. *Journal of Computer and System Sciences*, 50(1):20–31, 1995.

Alon Gonen, Sivan Sabato, and Shai Shalev-Shwartz. Efficient active learning of halfspaces: an aggressive approach. *The Journal of Machine Learning Research*, 14(1):2583–2615, 2013.

Elizabeth R Grant and Michael J Spivey. Eye movements and problem solving: Guiding attention guides thought. *Psychological Science*, 2003.

Steve Hanneke and Liu Yang. Minimax analysis of active learning. *CoRR*, abs/1410.0996, 2014.

Luis Haug, Sebastian Tschiatschek, and Adish Singla. Teaching inverse reinforcement learners via features and demonstrations. In *Advances in Neural Information Processing Systems*, pages 8464–8473, 2018.

Anette Hunziker, Yuxin Chen, Oisin Mac Aodha, Manuel Gomez Rodriguez, Andreas Krause, Pietro Perona, Yisong Yue, and Adish Singla. Teaching multiple concepts to a forgetful learner. In *Advances in Neural Information Processing Systems*, pages 4050–4060, 2019.

Kevin G Jamieson and Robert Nowak. Active ranking using pairwise comparisons. In *Advances in Neural Information Processing Systems*, pages 2240–2248, 2011.

Parameswaran Kamalaruban, Rati Devidze, Volkan Cevher, and Adish Singla. Interactive teaching algorithms for inverse reinforcement learning. In *IJCAI*, pages 2692–2700, 2019.

Daniel M Kane, Shachar Lovett, Shay Moran, and Jiapeng Zhang. Active classification with comparison queries. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 355–366. IEEE, 2017.

S Rao Kosaraju, Teresa M Przytycka, and Ryan Borgstrom. On an optimal split tree problem. In *Algorithms and Data Structures*, pages 157–168. Springer, 1999.

Andreas Krause and Daniel Golovin. Submodular function maximization. In *Tractability: Practical Approaches to Hard Problems*. Cambridge University Press, February 2014.

Laurent Lessard, Xuezhou Zhang, and Xiaojin Zhu. An optimal control approach to sequential machine teaching. In *AISTATS*, pages 2495–2503, 2019.

Weixin Liang, James Zou, and Zhou Yu. Alice: Active learning with contrastive natural language explanations. *arXiv preprint arXiv:2009.10259*, 2020.

Weiyang Liu, Bo Dai, Ahmad Humayun, Charlene Tay, Chen Yu, Linda B. Smith, James M. Rehg, and Le Song. Iterative machine teaching. In *ICML*, pages 2149–2158, 2017.

Weiyang Liu, Bo Dai, Xingguo Li, Zhen Liu, James M. Rehg, and Le Song. Towards black-box iterative machine teaching. In *ICML*, pages 3147–3155, 2018.

Shuang Ma, Zhaoyang Zeng, Daniel McDuff, and Yale Song. Active contrastive learning of audio-visual video representations. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=OMizHuea_HB.

Oisin Mac Aodha, Shihan Su, Yuxin Chen, Pietro Perona, and Yisong Yue. Teaching categories to human learners with visual explanations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3820–3828, 2018.

Stephen Mussmann and Percy Liang. Generalized binary search for split-neighborly problems. In Amos Storkey and Fernando Perez-Cruz, editors, *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 1561–1569. PMLR, 09–11 Apr 2018. URL http://proceedings.mlr.press/v84/mussmann18a.html.

Robert Nowak. Generalized binary search. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pages 568–574. IEEE, 2008.

Robert D Nowak. The geometry of generalized binary search. *IEEE Transactions on Information Theory*, 57(12):7893–7906, 2011.

Tomi Peltola, Mustafa Mert Çelikok, Pedram Daee, and Samuel Kaski. Machine teaching of active sequential learners. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32, 2019.

Stefanos Poulis and Sanjoy Dasgupta. Learning with feature feedback: from theory to practice. In *Artificial Intelligence and Statistics*, pages 1104–1113. PMLR, 2017.

Michael Prince. Does active learning work? a review of the research. *Journal of engineering education*, 93(3):223–231, 2004.

Brett Roads, Michael C Mozer, and Thomas A Busey. Using highlighting to train attentional expertise. *PloS one*, 2016.

Burr Settles. Active learning. *Synthesis lectures on artificial intelligence and machine learning*, 6(1): 1–114, 2012.

David Silver, J Andrew Bagnell, and Anthony Stentz. Active learning from demonstration for robust autonomous navigation. In *2012 IEEE International Conference on Robotics and Automation*, pages 200–207. IEEE, 2012.

Adish Singla, Ilija Bogunovic, G Bartók, A Karbasi, and A Krause. On actively teaching the crowd to classify. In *NIPS Workshop on Data Driven Education*, 2013.

Adish Singla, Ilija Bogunovic, Gabor Bartok, Amin Karbasi, and Andreas Krause. Near-optimally teaching the crowd to classify. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 154–162, Bejing, China, 22–24 Jun 2014. PMLR. URL http://proceedings.mlr.press/v32/singla14.html.

Stefano Teso and Kristian Kersting. Explanatory interactive machine learning. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 239–245, 2019.

Sebastian Tschiatschek, Adish Singla, and Andreas Krause. Selecting sequences of items via submodular maximization. In *AAAI*, 2017.

Sebastian Tschiatschek, Ahana Ghosh, Luis Haug, Rati Devidze, and Adish Singla. Learner-aware teaching: Inverse reinforcement learning with preferences and constraints. In *Advances in Neural Information Processing Systems*, 2019.

Peter Welinder, Steve Branson, Pietro Perona, and Serge Belongie. The multidimensional wisdom of crowds. In J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc., 2010. URL https://proceedings.neurips.cc/paper/2010/file/0f9cafd014db7a619ddb4276af0d692c-Paper.pdf.

Songbai Yan and Chicheng Zhang. Revisiting perceptron: Efficient and label-optimal learning of halfspaces. In *Advances in Neural Information Processing Systems*, pages 1056–1066, 2017.

Zhenliang Zhang, Edwin KP Chong, Ali Pezeshki, and William Moran. String submodular functions with curvature constraints. *IEEE Transactions on Automatic Control*, 61(3):601–616, 2015.

Xiaojin Zhu. An optimal control view of adversarial machine learning. *arXiv preprint arXiv:1811.04422*, 2018.

Xiaojin Zhu, Adish Singla, Sandra Zilles, and Anna N. Rafferty. An overview of machine teaching. *CoRR*, abs/1801.05927, 2018.

Sandra Zilles, Steffen Lange, Robert Holte, and Martin Zinkevich. Models of cooperative teaching and learning. *Journal of Machine Learning Research*, 12(Feb):349–384, 2011.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] Yes, we have a paragraph to discuss the limitation and future directions.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] See Section 7.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes] The proofs are in the Appendix.
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See supplemental material for details. The code is in preparation.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See supplemental material.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See supplemental material.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] I used datasets from other papers and cited the creators.
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]