Mixed-Initiative Active Learning

Maya Cakmak and Andrea L. Thomaz {MAYA,ATHOMAZ}@CC.GATECH.EDU
School of Interactive Computing, Georgia Institute of Technology, 801 Atlantic Dr., Atlanta, GA 30318, USA

Abstract

We propose a learning paradigm in which the responsibility of choosing samples is shared between the teacher and the learner. We present an experiment that demonstrates the potential of this approach with human teachers and discuss interesting related research problems.

1. Introduction.

A range of Machine Learning (ML) applications involve learning from data provided by a human teacher (*e.g.* document/image classification, recommender systems, programming by demonstration). Improving the sample efficiency of learning algorithms is particularly important in these applications, as providing data can be a cumbersome task for the human. Active Learning (AL) has proven useful towards this objective. Although theoretical results supporting the benefits of AL have been limited (Castro & Nowak, 2006; Balcan et al., 2010), many papers published over that last decade have demonstrated the practical strength of AL with human oracles (Settles, 2010). A number of directions have been explored with the purpose of improving techniques within the AL paradigm; however, one question that has not been asked frequently is "Can we do better than AL?"

The indication that we can do so comes from the field of Algorithmic Teaching (Balbach & Zeugmann, 2009; Goldman & Kearns, 1995) which studies the teachability of concepts. The teaching problem, unlike the learning problem, involves producing a sequence of examples from a *known* target concept such that the concept can be learned by a learner. Finding an optimal teaching sequence for an arbitrary concept is NP-hard (by reduction to the minimum cover set problem (Goldman & Kearns, 1995)) however efficient algorithms have been proposed for particular concept classes. An important insight is that an active learner can never learn faster than when a passive learner is taught optimally (Goldman & Kearns, 1995; Angluin, 2004). In other words, teachability indicates an upper bound on how fast a concept can be learned. Thus we believe that the key for going beyond AL will be in "good teaching." To this end, we hope to exploit the flexibility and intelligence of human teachers.

Presented at the *ICML 2011 Workshop on Combining Learning Strategies to Reduce Label Cost*, Bellevue, WA, USA, 2011. Copyright 2011 by the author(s)/owner(s).

We propose a learning paradigm in which the responsibility of choosing examples is shared between the teacher and the learner – Mixed-Initiative AL (MIAL). MIAL stands between passive supervised learning, where all learning examples are chosen by the teacher and AL where all examples are chosen by the learner and labeled by the teacher. In this setup, the learner receives an example even when it does not make a query. Therefore it needs to decide *when* to make a query, besides *what* query to make.

In this paper, we first give a few illustrative examples of good teaching outperforming AL. We then overview our previous work on human teaching which motivates a mixed-initiative approach and present a follow-up study that demonstrates the strength of MIAL.

2. Good Teaching

In this section we discuss the potential of "good teaching" in different problem settings.

2.1. Optimal Teaching

Quantifying the teachability of concepts is a central problem in Algorithmic Teaching. A number of teachability metrics have been proposed (Natarajan, 1989; Anthony et al., 1995; Balbach, 2008), the most popular being the *Teaching Dimension* (Goldman & Kearns, 1995). This is the smallest number of examples needed to uniquely identify *any* concept in a concept class. The shortest sequence of examples that uniquely identifies a concept is referred to as an *optimal teaching sequence*. A major concern is finding polynomial-time algorithms that produce an optimal teaching sequence for any concept in a concept class. Such algorithms exist for certain concept classes such as conjunctions, monotone decision lists and monotone K-term DNFs.

To illustrate the strength of optimal teaching we consider the concept class of conjunctions for which the teaching dimension is identified as

$$TD(C_n) = min(r + 2, n + 1) \tag{1}$$

where r is the number of relevant features (*i.e.* number of variables in the conjunction) and n is the total number of features (Goldman & Kearns, 1995). The sample complexity of this concept class in the PAC-Learning model is characterized by the inequality

$$\frac{1}{\epsilon}(n \ln 3 + \ln(1/\delta)) \leq m \tag{2}$$

which says that the number of examples m to teach a conjunction on n variables to a consistent learner such that its error is bounded by ϵ, with probability δ, is at least the number specified by the left hand side of the inequality (Mitchell, 1997). For instance, the number of required examples is 96 for a desired 0.95 probability that a hypothesis with at most 0.10 error will be learned for conjunctions with $n = 6$. On the other hand, the optimal teaching algorithm guarantees to teach any conjunction (in a sample space with 6 features) with at most 7 examples. From a human teacher's perspective, this can be a significant difference.

The number of examples required by a learner to uniquely identify any concept in a concept class C with membership queries is characterized by the inequality

$$\log_2 |C| \leq \#MQ(C) \leq |X| \tag{3}$$

where X denotes the instance space (Angluin, 2004). For conjunctions over n binary features, $|X| = 2^n$ and $|C| = 3^n$, since a feature can either be irrelevant or have one of the binary values in a hypothesis. For example, for $n = 6$ the number of membership queries to learn any concept is between 10 and 64. In this setting, the lower bound on samples to be labeled by a human teacher is reasonable and close to the optimal value. However the range is large and this number can become unreasonable for a human as it gets closer to the upper bound. Thus, optimal teaching can also significantly reduce the number of examples to learn a concept, as compared to AL.

Note that a crucial assumption in this example is that there exists an efficient algorithm that achieves optimal teaching. As we will discuss in Sec. 3, our goal is to use these algorithms as teaching guidance for humans in order to make them better teachers for machine learning algorithms.

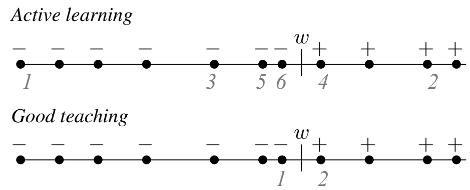

$$c_w(x) = \begin{cases} 1 & \text{if } x \geq w \\ 0 & \text{if } x < w \end{cases}$$

Figure 1. An example of empirically optimal teaching outperforming Active Learning: finding a decision boundary in $\mathbb{R}^1$ (Dasgupta, 2006).

2.2. Empirically Optimal Teaching

In some cases no optimal teaching algorithms exist (*e.g.* infinite concept classes and continuous instance spaces). For these cases we consider limiting the teaching problem to choosing examples from a finite set drawn from the instance space. This is a realistic setting often employed for studying AL algorithms.

As a simple example for this setting, consider learning a decision boundary in $\mathbb{R}^1$ illustrated in *Figure* 1. This is commonly cited as a case where AL provides logarithmic advantage over random sampling (Dasgupta, 2006). We can easily see that a good teacher can directly provide the examples closest to the decision boundary to achieve the smallest possible error on the rest of the data set. We refer to this as *empirically optimal teaching*. In general empirically optimal teaching requires enumerating all possible subsets of the data. However in some cases, such as in this example, the optimal teaching set can directly be identified by going over all examples once. This is achieved by exploiting the known structure of the sample and hypothesis spaces. Human teachers can also exploit such structure when they are teaching. Note that the visualization of data in this example makes teaching even easier for humans (*i.e.* in comparison to having an unsorted list of real numbers).

2.3. Heuristic Teaching

In practice finding such structure and visualizations is not trivial and empirically optimal teaching in polynomial time might not be possible. Nevertheless, teaching algorithms that outperform AL might exist. To discuss this scenario, we recreate the example in (Settles, 2010) that illustrates the advantage of AL over random sampling. It involves the classification of emails between two newsgroups from the *20 Newsgroups* dataset, using logistic regression. The features are counts of words in the email. *Figure* 2 shows learn-

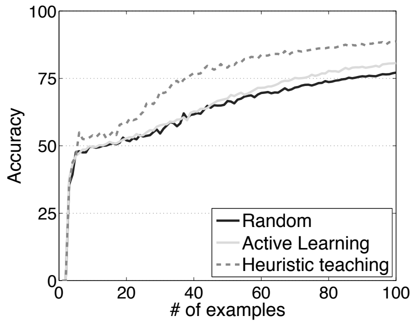

Figure 2. Progress of accuracy over the number of examples provided to the learner, averaged over 5-folds using cross-validation, in the 2-class classification of emails between two newsgroups using Logistic regression.

ing curves for random and uncertainty sampling, as well as for a learner trained by a greedy teacher that always presents the example that provides the most improvement in classifying the rest of the data set. While this algorithm might be unrealistic for humans to use, the result shows the existence of a dataset better than the one produced with AL. We hypothesize that these data sets can be captured by human teachers using teaching heuristics. Such heuristics might also allow good teaching in problems where the input set is not finite (*i.e.* the teacher both instantiates and labels an example for the learner, rather than picking an example from a finite set). We believe that good teaching heuristics can be derived from intuition based on the properties of the state space and the concept class. Another option is to inspect good teaching sets (*e.g.* the set that outperforms AL in *Figure* 2) to characterize properties of good teaching examples at a higher level.

3. Optimality of Human Teaching

The examples in the previous section demonstrate the potential of "good teaching." As mentioned earlier, our motivation in analyzing optimal teaching algorithms and algorithmic teaching heuristics is in deriving teaching strategies and heuristics that are usable by humans. We refer to this as *teaching guidance*. In this section we overview an experiment from our previous work (Cakmak & Thomaz, 2010) that explores the idea of teaching guidance for humans.

Our experiment involves the teaching problem discussed in Sec. 2.1. An optimal teaching sequence that satisfies the lower bound in Equation 1 is produced by the following algorithm (Goldman & Kearns, 1995):

(1) First show one positive example.
(2) Then show another positive example in which all

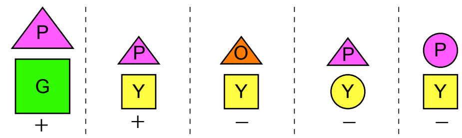

Figure 3. An optimal teaching sequence for the conjunction concept *HOUSE*.

$n - r$ irrelevant features are changed to the opposite value. This demonstrates that the features changed between the first two examples are irrelevant.
(3) Next, show negative examples that differ from a positive example by only one feature. This is repeated for each of the r relevant features. Since only one feature is different from what is known to be a positive example, these negative examples prove that the changed feature is a relevant one.

Our experiment looked at teaching conjunctions to a virtual agent. The sample space with $n = 6$ consists of objects composed of two pieces (top and bottom) where each piece is specified by three features (shape, color and size). An example concept with $r = 3$ is *HOUSE* which is defined as an object that has a pink triangle top and a square bottom piece, while other features do not matter. An optimal teaching sequence for this concept, which consists of 5 examples, is shown on *Figure* 3.

Our experiment compared *natural* human teaching with *guided* human teaching. In the natural teaching condition the subjects are not given any explicit instruction on how to teach while in the guided teaching condition they are asked to follow the optimal teaching algorithm. The algorithm is described using lay terms. Our main findings from this experiment were:

- Natural human teaching is much better than random, but not spontaneously optimal.
- Teaching guidance significantly improves human teaching, however does not make them optimal.

We observed that the improvement in the guided teaching condition was mainly due to step (2) in the described algorithm. Step (1) was already carried out during natural teaching (*i.e.* most human teachers naturally start with a positive example). Step (3), on the other hand, was not as intuitive for our subjects. Most subjects could not keep track of which features they had already varied, and either stopped early without uniquely identifying the target or defaulted to testing the learner until they found an example for which the learner was uncertain about. While this outcome points out the importance of the usability of teaching guidance, it also motivates a mixed-initiative approach.

4. Mixed-initiative Active Learning

In the experiment described above, the part of the teaching algorithm that humans were not good at, can actually be replaced with AL without loosing optimality. In other words, if the learner was to make a query after the first two examples in the sequence given in *Figure* 3, it would choose the same (or equivalent) samples as in the rest of the sequence. The part of the algorithm that humans were good at following happens to be the part that gives the advantage to the teacher in teaching efficiently. This step involves changing irrelevant features all at once. Having the knowledge of what these features are, the teacher can easily perform this step. An active learner, on the other hand, cannot risk to change more than one feature at a time, since if the resulting example is negative, it cannot infer which feature(s) were relevant.

This points towards an ideal division of labor between the teacher and the learner: the teacher should follow the steps (1) and (2) of the algorithm and then the learner should make queries. In this way, the workload of the teacher is reduced to providing two carefully chosen examples and then responding to queries, while the optimal number of examples to be labelled is maintained.

We explored this idea in a follow-up experiment where the learner has the capability to make queries. The queries are triggered by the teacher, *i.e.* the learner only makes a query when the teachers presses an "Any questions?" button. The experiment had two groups: the *active* learning group was told to use queries as they want, and the *mixed-initiative* group was given the first two steps of the teaching algorithm and was told to trigger and answer queries after that.

The results from this experiment demonstrate that the mixed-initiative approach can achieve close to optimal teaching. We observed that all subjects in the *mixed-initiative* condition fully identified the target concept, because they kept triggering queries until the learner converged to the correct hypothesis and did not have any more queries. In addition the length of the teaching sequences provided in this condition were optimal or close to optimal. As in the previous experiment, optimal teaching did not spontaneously emerge in the *active* learning condition with teacher-triggered queries.

We compare the different teaching conditions from both experiments in *Figure* 4 in terms of average informativeness of examples across subjects. We define the informativeness of an example as the ratio of the maximum accuracy gain provided by any example in the current state of the learner to the gain provided by this example. Thus, the informativeness of all examples in the optimal teaching sequence is 1 and the informativeness of a redundant example is 0. We see that the average informativeness in the mixed-initiative condition is close to 1 and is significantly higher than the other conditions.

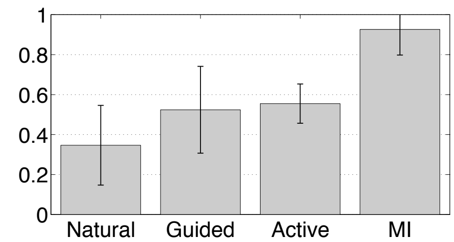

Figure 4. Average informativeness of examples provided by human teachers in four conditions in our experiments: *Natural* teaching, *Guided* teaching, *Active* learning with teacher-triggered queries, *Mixed-Initiative* teaching. The experiments involved 10 subjects in each condition.

5. Discussion

The mixed-initiative approach can be applied to other concept classes for which optimal teaching algorithms exist. This requires (i) identifying parts of an optimal teaching sequence that are obtainable with active learning, and (ii) describing the rest of the algorithm in a human-friendly manner and verifying its usability. Whether other concept classes will exhibit the nice property that led to the success of the mixed-initiative approach in our experiment is an interesting question that we would like to address in future work.

As discussed in Sec. 2, more complex scenarios require intuitive teaching heuristics, as opposed to exact teaching strategies. A potential MIAL setting with heuristics, consists of the human teacher providing a seeding set and the learner making queries in the rest. In most practical applications of AL the learner is seeded with a random set of examples from each class, which gives the learner a good place to start. However a large variance can be observed in the performance of the learner depending on the seeding set. Thus we can aim to device heuristics that let humans pick a good seeding set, such that the performance of the active learner is close to the upper bound of this variance.

We emphasize the potential of this idea with the example from Sec. 2.1 which theoretically identified the number of membership queries to learn any conjunction with $n = 6$ to be varied between 10 and 64. This number is reduced to a guaranteed 6 queries if the learner starts with a positive example. Given a positive example, an active learner can just query examples that differ from this example by only one feature to find out whether the feature is relevant or not. The learner converges to the target concept after repeating this for all n features. Thus having a good place to start in AL can make a big difference. Arguably, the practical success of many AL algorithms, is rooted in the common practice of seeding the learner with a random set of examples from each class.

In our experiment we skipped the problem of deciding *when* to make query (from the learner's perspective) in the mixed-initiative setting, by having the teacher trigger the queries. Leaving this decision to the learner raises other interesting problems, such as whether optimality can be guaranteed, the necessary and sufficient conditions to guarantee optimality, or the necessary and sufficient conditions to guarantee performance better than or equivalent to pure AL.

We note that even in cases where MIAL does not improve upon pure AL, the MI setting can provide a more balanced, possibly preferable user experience for the teacher which can be crucial in certain domains (Horvitz, 1999). In a different experiment (Cakmak et al., 2010), we demonstrated that MIAL without any guidance provides the same performance gain as pure AL, however is preferred from an interaction point of view. The constant stream of queries in pure AL is often found to be annoying and negatively affects the teacher's mental model of what has been learned. In addition, teacher-triggered queries are preferred over learner-initiated as they give full control of the interaction to the teacher.

References

Angluin, D. Queries revisited. *Theoretical Computer Science*, 313:175–194, 2004.

Anthony, M., Brightwell, G., and Shawe-Taylor, J. On specifying boolean functions by labelled examples. *Discrete Applied Mathematics*, 61(1):1–25, 1995.

Balbach, F.J. Measuring teachability using variants of the teaching dimension. *Theoretical Computer Science*, 397(1–3):94–113, 2008.

Balbach, F.J. and Zeugmann, T. Recent developments in algorithmic teaching. In *3rd Intl. Conference on Language and Automata Theory and Applications*, pp. 1–18, 2009.

Balcan, M.F., Hanneke, S., and Vaughan, J. The true sample complexity of active learning. *Machine Learning*, 80:111–139, 2010.

Cakmak, M. and Thomaz, A.L. Optimality of human teachers for robot learners. In *Proceedings of the IEEE International Conference on Development and Learning (ICDL)*, 2010.

Cakmak, M., Chao, C., and Thomaz, A.L. Designing interactions for robot active learners. *IEEE Transactions on Autonomous Mental Development*, 2(2): 108–118, 2010.

Castro, R. M. and Nowak, R. D. Upper and lower error bounds for active learning. In *The 44th Annual Conference on Communication, Control and Computing*, 2006.

Dasgupta, S. Coarse sample complexity bounds for active learning. In *In Proceedings of the NIPS Workshop on Cost-Sensitive Learning*, 2006.

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

Horvitz, E. Principles of mixed-initiative user interfaces. In *SIGCHI Conference on Human Factors in Computing Systems*, pp. 159–166, 1999.

Mitchell, Tom M. *Machine Learning*. McGraw-Hill, 1997.

Natarajan, B.K. On learning boolean functions. In *19th Annual ACM Symp. on Theory of Computing*, pp. 296–304, 1989.

Settles, Burr. Active learning literature survey. Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2010.