Active learning with teacher-learner mutuality

Matjaž Majnik, and Danijel Skočaj

University of Ljubljana, Faculty of Computer and Information Science
Tržaška cesta 25, 1000 Ljubljana, Slovenia
E-mail: {matjaz.majnik,danijel.skocaj}@fri.uni-lj.si

Abstract

In active learning, the basic objective is to reach a desired performance of some learning algorithm with as little training instances as possible. The reason behind is that labeling of training instances may be expensive with respect to the amount of time and intellectual effort of a human annotator. We propose a new approach for active learning, called "mutual active learning", which helps the artificial intelligent learner to pose questions to his human teacher, which are as clear and as understandable as possible. Such learning appears to be more reliable and successful than basic active learning.

1 Introduction

The basic goal of active learning is to speed up the learning process. The term "active" comes from the fact that the learner is taking part in choosing which specific feature vectors should be labeled by the teacher and used for subsequent training. The learner usually estimates the informativeness of available, unlabeled training instances. Several different strategies, also called *query strategies*, of how to evaluate the informativeness exist [8, 5]. In our previous work, e.g. [9, 4], we already have studied and evaluated various approaches that build on these underlying query strategies and are used for the final selection of the most appropriate training instance to be labeled in each step of active learning.

We assume the socially interactive scenario where a human teacher interactively communicates with an intelligent system. In this scenario, active learning is motivated by saving the teacher's time and effort through decreasing the number of training instances that have to be labeled. However, not all instances are equally hard to classify, i.e. to assign appropriate label. In general, ambiguous instances near borders between classes (which tend to be the most informative), may be harder for the teacher to label, taking more of his or her effort and time.

If we make a real-world-like assumption that *human teachers are not all-knowing* and may make mistakes, we realize that the selection of instances should not only be based on informativeness but also on how well a teacher understands provided training instances. To this end we propose a new approach for active learning, called *mutual active learning*, which should ensure that informative instances chosen by active learning will also be understandable to the human teacher. The concept of this type of learning has, in the field of artificial intelligence, to the best of our knowledge, not been explicitly investigated. We use the term "mutual" to denote learning in which also the teacher has influence on the selection of training instances which the learner provides for labeling. This paper presents the core idea with a proof of concept. The initial evaluation has confirmed that such learning is more reliable and successful than basic active learning.

The paper is organized as follows. In Section 2 we first discuss the related work. In Section 3 we present the newly proposed approach for mutual active learning. Then, in Section 4 we present the experimental evaluation, and conclude the paper in Section 5.

2 Related work

In socially assistive robotics [10] the motivation is to construct robots able to help human users through social interaction, and thus contribute to the quality of their lives. Such robots should be capable of natural communication and social interaction, and may be especially beneficial to elder population and people with cognitive, developmental and social disabilities. In such a setting, understandability of robot's questions is greatly important, and socially assistive robotics may thus be one of the areas which could benefit from mutual active learning.

In [1] the authors discuss what a good question is from the perspective of machine learning as well as human learning. The authors identify different types of questions, however, they do not explicitly consider mutuality, i.e. tutor's influence on the selection of learner's queries.

One of the main assumptions of our paper, the assumption of noisy human teachers, has been addressed in [8] as one of the challenges for practical use of active learning. The problem has been further studied in [2], where the authors discuss behavior of human experts in real-world situations. In their study, the authors assume that when the expert is highly confident in training instances which he or she labels, the likelihood of these labels being correct should also be high. The analysis about when and how human experts provide incorrect answers is then presented. A new method is proposed on this basis for balancing the selection of instances which are very informative (but the expert is likely to mislabel some of them) and instances with high likelihood of being correctly labeled (but having little informative value). This balancing, though addressed from a different perspective, also represents the core idea of our paper. This problem is also known as *exploration/exploitation dilemma (trade-off)* and is further discussed in [6].

In [7] the authors study the scenario where non-experts are employed as robot's teachers. This scenario is related to our work, since the basic problem with non-experts is namely the understanding of robot's questions, what may result in inaccurate responses. The difference in comparison to our approach is, however, that we assume a teacher to have an *active* role in communicating what information is personally more understandable to him (the learner may adapt to the knowledge of an individual teacher).

3 Mutual active learning

The motivation and objectives in active learning and mutual active learning may be summarized as follows.

1. Active learning (AL):
   (a) general motivation: Transfer of knowledge from the teacher to the learner with minimum possible number of training instances.
   (b) human teacher: "Learner, please learn as fast as possible, my time and effort are precious."
   (c) artificial intelligent learner: "Teacher, please provide instances which will help me to learn as fast as possible."
2. Mutual active learning (MAL):
   (a) general motivation: As in AL + with as little intellectual effort for the teacher as possible.
   (b) human teacher: As in AL + "Learner, please pose questions which are frustrating as little as possible (so that I will be able to answer each question as reliably and as quickly as possible)."
   (c) artificial intelligent learner: As in AL.

In Section 3.1 we outline the simplest approach of incorporating "mutuality" into active learning. Our hypotheses are stated in Section 3.2.

3.1 Incorporating the teacher's feedback

The simplest manner of incorporation is the following. Instead of only one, the learner offers several unlabeled training instances to its teacher in every step of learning. These instances should fulfil active learning requirements, i.e. they should all be approximately equally informative. The teacher then selects and labels the one which is the most understandable from his perspective. This instance is then used for training in a given step.

In the implementation we use directed uncertainty sampling (based on the calculation of posterior probability over all possible classes) in the form of a Monte Carlo–like method, described in [4]. This sampling of a feature space is used for generating training instances which are then offered to the teacher for labeling. The process of generation is fully introspective (see [4]).

3.2 Our hypotheses

As on each step of interactive learning we only operate on instances which are similarly informative (and are therefore equally useful for the learner), there should generally be no risk that the proposed method could have a negative influence on the final classification accuracy. Our hypotheses are thus the following.

1. Classification accuracy should increase due to the fact that training instances the teacher obtains are likely to be more understandable to him. The teacher can thus classify them more reliably.
2. Training time should decrease as it will take less time for the teacher to label training samples that he or she better understands (or is more familiar with, or are less ambiguous). Whether the training time actually decreases depends on the domain, more specifically, on how important it is for the human annotator to avoid possible misclassifications. We assume that with increasing importance the training time should decrease in comparison to basic active learning.
3. Mental workload on a human teacher should decrease due to more understandable training instances.

4 Experimental evaluation

The hypotheses have been tested on the problem of recognizing color samples taken from the (three-dimensional) HSL space (with H, S, and L being attributes). The task has been to label the provided samples as being of one of the following eight colors (classes): red, yellow, blue, orange, green, pink, black, and white. This domain possesses the following two properties which make it suitable for our purpose:

- The notion of colors is, to some extent, inherently subjective. There are no precise borders between different colors.
- Labels of some typical colors are missing (e.g. brown, purple, and gray). The teachers have to choose which of the available labels best describes a given sample, or have to opt for "None of the available".

In the evaluation we have used four learning approaches, two of them (*random* and *multi-random*) being auxiliary.

1. *mutual* – Mutual active learning. In each step multiple feature vectors (FVs) are obtained as described in Section 3.1. These FVs are then offered for labeling.
2. *pure* – Pure active learning. In each step only one FV is obtained (with the same procedure as in *mutual*). This FV is then offered for labeling.
3. *random* – Random learning. In each step a FV is generated randomly. There is one FV offered for labeling.
4. *multi-random* – Multi-random learning. In each step multiple FVs are generated randomly. All of these multiple FVs are then offered for labeling.

Eight teachers have participated in experiments. Five of them were in the age group 20–30 years, and the other three in the age group 30–40 years. Three teachers were female, five were men. Each teacher has in total assigned labels to 840 FVs:

- initial training set ($L1$) with 80 FVs (balanced class distribution with 10 FVs for each of the 8 classes),
- interactive training sets ($L2$) with 50 FVs for each of 3 repetitions and 4 learning approaches with altogether $3 \times 4 \times 50$ FVs (generally imbalanced class distribution),
- and test set (T) with 160 FVs (balanced class distribution with 20 FVs for each of the 8 classes).

For more reliable comparison of results, FVs of the initial training set and the test set were common for all teachers (i.e. the teachers were given the same FVs but the labels to these FVs have been assigned individually by each teacher). On the other hand there has been produced a unique interactive training set for each *(teacher, repetition, approach)* combination based on the four above-mentioned learning approaches, i.e. altogether $8 \times 3 \times 4 = 96$ interactive training sets.

4.1 Implementation of the experiment

The experiment has been conducted on the above-described interactive learning system running on a personal computer. For the underlying learning algorithm we have used the odKDE [3]. The learning system provided interactive lessons and the human teachers were answering the questions posed by the system.

The following rules were presented to the teachers:

- Approaches *random* and *pure* (type A) : Label the color sample as being of one of the 8 available colors. If you cannot decide for any of these colors, label the sample as “None of these”.
- Approaches *multi-random* and *mutual* (type B) : Choose one of the 15 offered samples for which you are able to predict its color. Again, there are 8 colors available. Label the chosen sample as being of one of these colors. If you cannot assign any of these colors to any of the 15 samples, choose “None of these”. Try to teach the system to distinguish between concepts of all 8 colors (i.e., if possible, try to choose samples of diverse colors).

Firstly, each teacher labeled 240 instances; 80 have been dedicated to the initial training set ($L1$), while 160 have been used for the test set (T). Afterwards, each teacher went over each of the four learning approaches (discussed in the beginning of Section 4) in three subsequent repetitions. Each of these approaches determines, how a FV (or several FVs) is generated in every step. A color sample corresponding to the generated FV is then shown to a teacher. In total, a teacher labeled 12 unique interactive training sets, each of them corresponding to one of 12 $(repetition, approach)$ pairs. No additional information concerning the supposed quality of the underlying approaches were provided to the teachers. At the end each teacher has answered four questions about his or her interaction with the system.

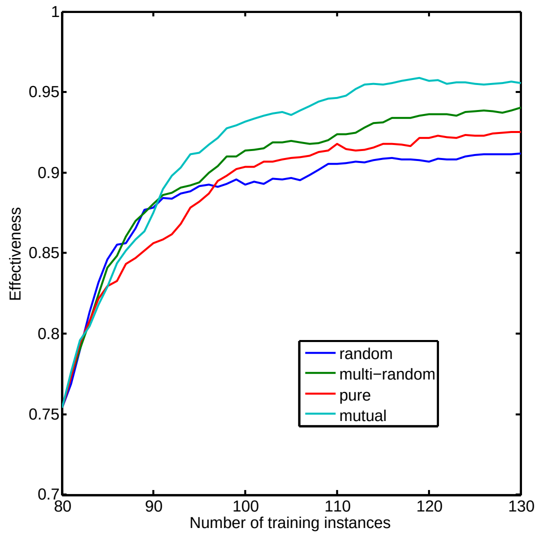

Figure 1: Average effectiveness rates

4.2 Objective and subjective measures for evaluation

The objectively measured quantities in the final analysis (for the constant number of training samples) are: training time, and effectiveness. Training time has been measured from the moment the teacher was acquainted with the first question until the learner got the last answer. The effectiveness has been determined by calculating the classification accuracy. Subjective evaluation of mental workload on human teachers has been made on the basis of a questionnaire at the end of the learning process.

4.3 Results and discussion

On Fig. 1 there are average effectiveness rates obtained from individual results of all 8 teachers and all 3 repetitions made by each teacher. As expected, *pure* active approach is performing better than *random* approach as the learning is tailored to the learner. Due to human effort *multi-random* approach also outperforms the *random* one. And finally, the *mutual* active approach has the best performance, better than *pure* (where only the learner directs the process) and *multi-random* (where only the teacher directs the process) approaches. These results confirm our first hypothesis.

Tab. 1 shows mean time for labeling one training instance over all teachers and all repetitions. We may notice that the mean gross time for labeling individual training instances is, as opposed to our second hypothesis, for about 42% higher in case of *mutual* (4.4 s) than in case of *pure* (3.1 s) approach. Gross time in the case of the *pure* approach comprises of a labeling decision time and a click time. On the other hand, at *mutual* approach the gross time includes selection decision time, a click time, labeling decision time, and finally another click time. We believe, that in domains where the consequences of mislabeling a sample may be serious, such as, e.g., safety-critical and medical applications, the gross time should be larger for the *pure* than for the *mutual* approach. In our opinion, the difficulty of labeling a typical actively obtained sample would outweigh the overhead of selection decision time and the second click time. However, on the basis of experimental results, the second hypothesis cannot be confirmed.

On the other hand, we should be aware that the increase in effectiveness rate also causes the training time for reaching some specific effectiveness rate to be shorter. For example, to achieve 92% effectiveness rate, the teacher has to label 119 training instances in *pure* active learning mode while only 97 training instances is needed in *mutual* active approach.

Table 1: Mean time for labeling one training instance.

random	multi-random	pure	mutual
3.2 s	4.7 s	3.1 s	4.4 s

The questions used in the subjective evaluation are presented below and the results are available in Tab. 2.

Basically you have met two types of approaches for labeling color samples:

- *type A: without option to choose (always only one sample)*
- *type B: with option to choose (always 15 samples)*

Please answer with A or B to the following four questions. Explain your decision shortly.

1. *Which of the two types enabled you to more quickly decide of which color the (chosen) sample is? Why?*
2. *In your opinion, at which type the learning was more successful? Why?*
3. *Which of the two types you perceived as mentally less demanding? Why?*
4. *Which of the two types made you feel (more) uncertain / insecure? Why?*

Table 2: The subjective evaluation.

question	answer A	answer B	other answer
1	0	**6**	2
2	1	**5**	2
3	3	**4**	1
4	**7**	1	0

The results confirm our third hypothesis. We may observe that for the great majority of teachers the preferred type in questions 1, 2 and 4 is the type B. Type B is also the preferred type in question 3. At this question, however, the teachers' opinions differ the most: while in type A no decision making about which instance to choose for labeling was needed, it was easier in type B to decide what color the chosen instance is.

5 Conclusion

We have proposed a new concept for improving performance of active learning, called *mutual active learning*. We have presented its features together with its expected advantages. The empirical evaluation has supported two our hypotheses.

Since the idea of mutual active learning is domain-independent, the newly developed method may also be useful in other fields of application. In future work we will examine the performance of the mutual active learning on other domains (e.g. on the problem of recognizing handwritten digits).

Mutual active learning is approaching real world-like assumptions and is expected to help reducing time and effort a teacher has to invest in training process of an artificial intelligent system. Due to faster and more human-friendly learning these systems should become more appropriate for interactive social applications.

References

[1] M. Cakmak and A. L. Thomaz. Designing robot learners that ask good questions. In *Proceedings of the 7th ACM/IEEE International conference on human-robot interaction(HRI)*, pages 17–24, 2012.

[2] J. Du and C. X. Ling. Active learning with human-like noisy oracle. In *Proceedings of the 10th IEEE International conference on data mining (ICDM)*, pages 797–802, 2010.

[3] M. Kristan and A. Leonardis. Online Discriminative Kernel Density Estimation. In *International Conference on Pattern Recognition*, pages 581–584, 2010.

[4] M. Majnik, M. Kristan, and D. Skočaj. Knowledge gap detection for interactive learning of categorical knowledge. In *Proceedings of the 18th Computer Vision Winter Workshop*, pages 94–101, 2013.

[5] Fredrik Olsson. A literature survey of active machine learning in the context of natural language processing. Technical report SICS T2009:06, Swedish Institute of Computer Science, 2009.

[6] T. Osugi, D. Kun, and S. Scott. Balancing exploration and exploitation: A new algorithm for active machine learning. In *Proceedings of the 5th IEEE International conference on data mining (ICDM)*, pages 330–337, 2005.

[7] S. Rosenthal, M. Veloso, and A. K. Dey. Acquiring accurate human responses to robots' questions. *International journal of social robotics*, 4(2):117–129, 2012.

[8] Burr Settles. Active learning literature survey. Technical report 1648, University of Wisconsin-Madison, 2010.

[9] D. Skočaj, M. Majnik, M. Kristan, and A. Leonardis. Comparing different learning approaches in categorical knowledge acquisition. In *Proceedings of the 17th Computer Vision Winter Workshop*, pages 65–72, 2012.

[10] Adriana Tapus, Maja J. Matarić, and Brian Scassellati. The grand challenges in socially assistive robotics. *IEEE Robotics and Automation Magazine*, 14(1):35–42, 2007.