ECAI 2008
M. Ghallab et al. (Eds.)
IOS Press, 2008

doi:10.3233/978-1-58603-891-5-773

Active Concept Learning For Ontology Evolution [1]

Murat Şensoy and **Pınar Yolum** [2]

Abstract. This paper proposes an approach that enables agents to teach each other concepts from their ontologies using examples. Unlike other concept learning approaches, our approach enables the learner to elicit the most informative examples interactively from the teacher. Hence, the learner participates to the learning process actively. We empirically compare the proposed approach with the previous concept learning approaches. Our experiments show that using the proposed approach, agents can learn new concepts successfully and with fewer examples.

1 Introduction

In current approaches to concept learning, the learner is passive. That is, the training examples are solely chosen by the teacher. However, this assumes that the teacher has an accurate view of what the learner knows, which concepts are confusing for it, and so on. We propose to involve the learner in the learning process by enabling it to interact with the teacher to elicit the most useful examples for its understanding of the concept to be learned.

In our approach, each agent represents its domain knowledge using an ontology and manages this ontology using a network of *experts*. Each expert is a stand-alone learner composed of one or more classifiers. Main task of an expert is to learn how to discriminate between the sub-concepts of a specific concept. An agent learns a new concept from another agent using our approach as follows:

1. The learner agent asks for the positive examples of the concept from the teacher agent.
2. After receiving the positive examples, the learner determines the new concept's parent in its ontology using those positive examples. Then, the expert related to the parent concept is entitled to learn the new concept.
3. This expert determines some negative examples of the new concept using the positive examples and a *semi-supervised learning approach*. Hence, it first learns the new concept roughly without receiving any negative examples from the teacher.
4. The expert iteratively enhances its knowledge on the new concept by eliciting the most useful negative examples from the teacher.
5. After learning the new concept sufficiently, it is placed into the learner's ontology and the ontology is modified accordingly.

2 Representing Knowledge

In the current instance-based concept learning approaches, one classifier is trained to learn each concept independently [3]. Although the concepts are related through parent-child relationships, their classifiers are regarded as independent of one another. Such approaches require each classifier to learn how to discriminate instances of one concept from those of every other concept in the ontology. Therefore, in order to learn a single concept, the agent uses the whole domain knowledge.

In this paper, we envision that the domain knowledge related to an ontology is managed by a set of *experts*, each of which is knowledgeable in a certain concept. By knowledgeable in a concept, we mean that the expert can correctly report which of the concept's subclasses an instance belongs to. Hence, each expert is trained with examples of the concept and nothing else. For example, an expert on motorcycles can tell us correctly that *Burgman 400* is a scooter.

3 Actively Learning A Concept

While teacher teaches a new concept to the learner, it first selects a set of positive examples of the concept. This is relatively easier than selecting negative examples, which are chosen among instances of any other concept. Then, the teacher gives the selected positive examples to the learner. In our approach, negative examples are not directly given by the teacher, because the teacher cannot estimate which examples are more useful or informative for the learner. The given positive examples are classified using the experts of the learner and the most specific concept in the learners ontology is determined so that this concept subsumes all of the positive examples. Assume that, the teacher wants to teach *Motorcycles* concept to the learner, so it first provides examples of motorcycles to the learner. The learner realized that all of the provided examples are instances of $Car\&Motorsports$ concept in its ontology. Hence, learning task is delegated to the expert of $Car\&Motorsports$ concept. The expert examines the other instances of $Car\&Motorsports$ to differentiate given motorcycles from the others as much as possible.

Motorcycle instances should have some features in common that make them separate from the other instances of $Car\&Motorsports$ concept. In order to determine which features are more important for the $Motorcycles$ concept, the differences of the feature distributions between the positive examples and the unlabeled examples can be used [2, 5]. We can estimate how significant an instance I is as a motorcycle example, using the significance of its features. After computing the significance value for each known instance of $Car\&Motorsports$, the obvious negative examples of $Motorcycles$ are chosen among the instances that have the least significance values. Using these negative examples and the positive examples provided by the teacher, the expert tries to learn the new concept roughly. Note that until now, the teacher has not provided any negative examples.

Using the positive examples of $Motorcycles$ and the obvious negative examples, the expert trains a classifier. This classifier can

[1] .This research has been partially supported by Boğaziçi University Research Fund under grant BAP07A102 and The Scientific and Technological Research Council of Turkey by a CAREER Award under grant 105E073.
[2] Department of Computer Engineering, Boğaziçi University, Bebek, 34342, Istanbul,Turkey, email: {murat.sensoy,pinar.yolum}@boun.edu.tr

roughly discriminate instances of $Motorcycles$ from other instances of $Car\&Motorsports$. However, the boundary between these two classes is not learned precisely yet, because only the obvious negative examples are used for training. Moreover, some of these negative examples can be wrongly chosen. This may seriously affect the performance of the trained classifier. Therefore, the expert iteratively elicits more useful negative examples from the teacher and learns this boundary more precisely and correctly. Specifically, at each iteration, the expert samples instances of $Car\&Motorsports$ and then using the classifier, it labels these sampled instances as instance of $Motorcycles$ or not. Then, the teacher instructs the expert about the correct labels of these examples. The feedback from the teacher is used to refine and improve the knowledge of the expert about the new concept $Motorcycles$. This, iterative active learning phase continues until the teacher makes sure that the learner correctly learns the concept. Then, the new concept is placed into the learner's ontology as a new subconcept of $Car\&Motorsports$. Lastly, we test whether $Motorcycles$ concept subsumes some subconcepts of $Car\&Motorsports$ or not. If this is the case, concept-subconcept relationships are rearranged.

4 Evaluation

In order to evaluate our approach, we conduct several experiments in online shopping domain. For this purpose, we derive domain knowledge from Epinions[3]. In our experiments, there is one teacher agent and one learner agent. In the implementation of the agents and the experts, we use JAVA and the C4.5 decision tree classifier of WEKA data mining project [4]. In our experiments, an instance refers to a product item such as *IBM ThinkPad T60*, which is an instance of *PCLaptops* concept. Each product item has a web page in Epinions website and this page contains specification of the product item in English. We derive a core vocabulary from these specifications automatically and each word in this vocabulary is used as a feature [2].

Figure 1 shows the performance of our approach at each iteration in terms of the probability of misclassification. In Figure 1, after the first iteration, the expert learns the new concepts roughly (with %12 error). This error rate is not acceptable for the teacher, so the expert continues with the next iteration. The second iteration results in a considerable progress in the learning performance (error drops to %4). The classification error drops to zero at the fifth iteration, which means that the teacher and the learner have exactly the same understanding for this concept.

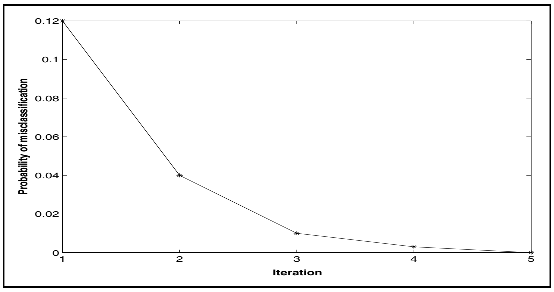

Figure 1. Probability of misclassification at different iterations.

We compare our approach with a teacher-driven concept learning approach. This approach represents the current concept learning approaches in the literature. Contrary to the proposed approach, in those approaches, the learner is inactive during the selection of the negative examples [3, 1]. The teacher selects the negative examples using its own ontology and viewpoint. Then, the learner is given positive examples and negative examples of the concept to be taught.

In order to measure how successful our approach is in learning new concept for different number of negative examples, we set up experiments where the teacher is allowed to give or label only a predefined number of negative examples. Then, these examples are given to the learner (as feedback in our approach). After training the learner with these examples, probability of misclassification is computed. Figure 2 compares the results for the teacher-driven approach and the proposed approach.

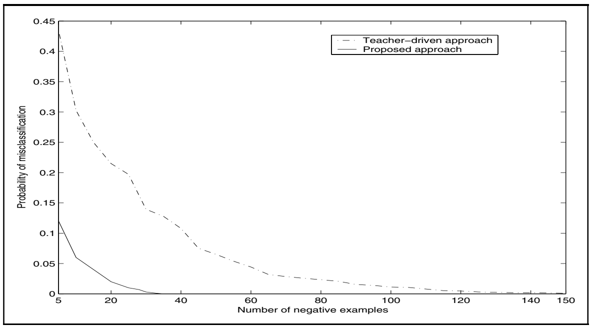

Figure 2. Probability of misclassification with different number of negative examples.

As seen in Figure 2, the teacher-driven approach requires more negative examples than the proposed approach in order to achieve an acceptable performance. With only five negative examples, the learner that uses the proposed approach fails only on the 12% of its classifications. However, in the same case, the learner using the teacher-driven approach misclassifies an instance with a probability of slightly higher than 0.4. Similarly, with only 35 negative examples, on the average, the proposed approach can learn a concept perfectly, while the teacher-driven approach requires approximately 150 negative examples for the same quality of learning.

5 Discussion

This paper develops a framework for instance-based concept learning, where a learner can estimate some negative examples of the concept to be learned and obtain feedback about these negative examples from the teacher to learn the concept accurately. Our experiments show that our approach significantly outperform a teacher-driven approach that represents other instance-based concept learning approaches in the literature by enabling learners to learn a concept with few examples.

REFERENCES

[1] A. Doan, J. Madhaven, R. Dhamankar, P. Domingos, and A. Helevy. Learning to match ontologies on the semantic web. *VLDB Journal*, pages 303–319, 2003.

[2] G. P. C. Fung, J. X. Yu, H. Lu, and P. S. Yu. Text classification without negative examples revisit. *IEEE TKDE*, 18(1):6–20, 2006.

[3] S. Sen and P. Kar. Sharing a concept. In *Working Notes of the AAAI-02 Spring Symposium*, pages 55–60, 2002.

[4] I. H. Witten and E. Frank. *Data Mining: Practical machine learning tools and techniques.* Morgan Kaufmann, San Francisco, 2005.

[5] H. Yu, J. Han, and K. C.-C. Chang. PEBL: Web page classification without negative examples. *IEEE TKDE*, 16(1):70–81, 2004.

[3] http://www.epinions.com