18^{th} *Computer Vision Winter Workshop*
Walter G. Kropatsch, Fuensanta Torres, Geetha Ramachandran (eds.)
Hernstein, Austria, February 4-6, 2013

Knowledge gap detection for interactive learning of categorical knowledge

Matjaž Majnik, Matej Kristan, and Danijel Skočaj
University of Ljubljana, Faculty of Computer and Information Science
Tržaška 25, 1000 Ljubljana, Slovenia
{matjaz.majnik,matej.kristan,danijel.skocaj}@fri.uni-lj.si

Abstract. *In interactive machine learning the process of labeling training instances and introducing them to the learner may be expensive in terms of human effort and time. In this paper we present different strategies for detecting gaps in the learner's knowledge and communicating these gaps to the teacher. These strategies are considered from the viewpoint of extrospective and introspective behavior of the learner – this new perspective is also the main contribution of our paper. The experimental results indicate that the analyzed strategies are successful in reducing the number of training instances required to reach the needed recognition rate. Such a facilitation may be an important step towards the broader use of interactive autonomous systems.*

1. Introduction

Cognitive systems are often characterised by their ability to learn, interact with the environment, and act autonomously. They are able to respond to requests of human users and other cognitive agents, and are also capable of taking the initiative and engaging in a dialogue with a human. Very importantly, they are able to learn from such interactions; they are able to acquire novel knowledge and update previously learned conceptual models in an incremental manner. They can passively receive the information they need in this incremental learning process. In this case they simply rely on the environment, or on a human tutor, for being provided with appropriate information for efficient learning. However, they can also take an active part in this incremental learning process and try to infer what kind of information is needed to make the learning more efficient. The latter learning approach is known as *active learning*.

Active learning requires that the system identifies learning opportunities. This in turn requires that it must be able to *detect gaps in its knowledge*, which may indicate good learning opportunities. Typically the knowledge gaps are detected in a particular modality; they are usually grounded in a particular representations. Subsequently, the knowledge gaps have to be, in some general form, *communicated* to the rest of the system and to other agents that can plan and execute actions necessary to fill these gaps. After the system obtains the required information, it can extend its current knowledge accordingly. A crucial requirement of a system that is supposed to self-extend is therefore a certain level of self-understanding that enables the detection and communication of gaps in its knowledge.

In our work we focus on this problem. We address the problem of knowledge gap detection in the context of the active learning paradigm and address specific active learning strategies from the viewpoint of extrospective and introspective robotic behavior. This new robotics-oriented view also represents the main contribution of our paper. Since our research has been concentrated around interactive continuous learning of conceptual knowledge in dialogue with a tutor, most of this paper has been written with this learning scenario in mind. However, the proposed solutions are general enough that they can be applied to other learning domains as well.

Our final goal is to develop active learning strategies which would successfully reduce the amount of learning data needed to transfer the categorical knowledge from the human teacher to the robotic learner. These strategies should be used to construct as autonomous and as domain-independent framework for dialogue-based learning as possible.

The paper is organised as follows. In Section 2 we first discuss the related work. In Section 3 we present an approach to knowledge gap detection. In Section 4 we then discuss different ways of knowledge gap communication and propose four active learning strategies based on them in Section 5. Then we present the experimental results in Section 6 and conclude the paper with some final remarks in Section 7.

2. Related work

Active learning strategies proposed in the literature mainly address the problem of estimating classifiers using minimal amount of data. They are motivated by the fact that there are many situations in which large quantities of unlabeled data are relatively easily obtained, however, the cost of labeling each instance can be high. Two extensive survey papers on active learning literature are available providing a broad overview of the field [11, 13]. Furthermore, a survey has recently been published [7], studying and comparing utility metrics and learning strategies for selecting training instances in active learning.

In this work we focus our attention towards *interactive learning of categorical knowledge in dialogue with a teacher*, similarly as the authors in [5]. For such real-life situations, it is desirable that the learner detects good candidates for querying on the fly and updates its classifiers accordingly, while requiring minimal involvement of the teacher. Also other authors focus more on this *social aspect* of active learning. In [8] the authors present a learning strategy akin to our *LDieSel*. They have similarly combined principles from two active learning scenarios, query synthesis and pool-based sampling. However, in their case the motivation is in obtaining improved performance by eliminating specific disadvantages of each of the two scenarios, while in ours we essentially enable the communication between the robotic learner and the human teacher.

In [1] four learning modes are described, which distinguish in how frequently the learner communicates with the teacher to obtain information, i.e. how often the interaction takes place. Four passive learning strategies based on features of biological systems are implemented in [2], where the strategies differ in the way they update the informativeness of individual objects. In [3] active learning is discussed from the viewpoint of transparency of the learner's internal states, and how these states may be used to inform the teacher about the level of the learner's knowledge. The approach is implemented on a physical robot with socially expressive head and neck, and simple non-verbal gestures are employed to provide the teacher with the transparent insight into the underlying model uncertainties.

In [15] the difference between training instances selected by a human teacher and systematically collected training instances is investigated. Besides, a method for the learner to convey the information about its knowledge gap to the teacher is presented. [1, 15, 3] include experiments with one or more human teachers, whereas in [2] robot-to-robot interaction is employed with the intention of providing a controlled environment for systematically exploring of how the learner is influenced by different teacher behaviours.

The authors in [10] discuss an active learning system from the viewpoint of combining interactive social learning (with a human teacher) and autonomous, non-interactive, intrinsically motivated learning. On the other hand, a learning strategy in [4], combines interactive autonomous selection of training instances (in areas of the problem space where the learner can classify with sufficient certainty) and non-interactive teacher-driven selection (in not-well-explored areas). The underlying learning method is, similarly as our odKDE, based on Gaussian mixture models (GMMs). A system with the goal of autonomous exploration of new knowledge is presented in [12] and discussed from the viewpoint of intrinsic motivation systems for autonomous development of robotic learners. The authors are inspired by human development where intrinsic motivation plays an important role and which may be characterized as progressive, incremental, active and autonomous.

The learning strategies that we address in our work are related to several approaches presented in the above-mentioned papers. We have built on our work in [14] and further developed certain learning strategies. Additionally, the meta-learning framework has been generalized to work on high-dimensional data (i.e. comprising tens of attributes) and tested with two learning methods.

3. Knowledge gap detection

In the active learning cycle there are two very important tasks that the system has to complete in order to get novel information that would help it to improve its knowledge in an efficient way: *detection* (Section 3) and *communication* (Section 4) of knowledge gaps. It should be noted that in the active learning community the process of "knowledge gap detection" is recognized under the term "selection of informative training instances".

3.1. Extrospection and introspection

The crucial step in an active learning cycle is the detection of ignorance. The system should first self-understand – it should understand what it does and what it does not know. By using its internal modal representations it should detect what information is missing. Generally, the knowledge gap detection can be tackled in two different ways; it can either be related to a particular situation or not. We therefore may distinguish between two types of knowledge gap detection: extrospective and introspective.

In the case of **extrospective** knowledge gap detection, this process is related to a particular situation, i.e., to a particular object in a scene or to some other training instances the learner can perceive. The learner tries to detect the lack of knowledge by observing and trying to recognize a number of existing objects that it might or it might not know. This is a typical pool-based active learning approach – the learner is given a number of unlabeled objects, and it has to select one (or several) of them for labeling. Ideally, it would select the instance that would help to improve its knowledge most (by providing instance's real label).

Another way of detecting knowledge gaps is through **introspection**. In this case the detection of knowledge gaps is completely self-driven and is not related to any particular situation or task. It is not triggered by any external problem; it is triggered by an inner motivational mechanism with the goal of detecting ignorance and proposing actions that would provide the information needed to extend the current knowledge. No sensorial inputs are used in this case; the detection of knowledge gaps is based solely on the current knowledge. Since there are no real instances the learner could estimate its knowledge on, the robot could try to hallucinate sensorial inputs (basically, sample over distributions of feature values it uses) and attempt to interpret these hallucinated situations. Failing to do that would indicate a knowledge gap. This type of knowledge gap detection is thus also based on the output of the classifier, which is built on the top of the models; the only difference is that the input is hallucinated and not perceived.

3.2. Measuring uncertainty

In knowledge gap detection the crucial task is to measure how good the knowledge is, i.e., how certain or uncertain is the classification of particular observation or hallucination. In the case of extrospective knowledge gap detection the learner tries to estimate the certainty of recognition of all available data. Similarly, in the case of introspective knowledge gap detection the learner tries to estimate the certainty of hallucinated instances. In both cases, it has then to select the deepest gap in its knowledge based on the estimated certainties.

Several certainty measures can be employed. As mentioned above, our knowledge gap detection mechanism requires that the underlying learning and recognition methods provide posterior probability over all categories, therefore the responses $p(M_i|\mathbf{z})$ of all k models M_i for every given observation $\mathbf{z}$. Our method for calculating the certainty by analysing the posteriors is presented below.

First, we determine two models with the highest response:

$$M_{maxap} = \arg\max_{M_i}\{p(M_i|\mathbf{z})\} \tag{1}$$

$$M_{maxap2} = \operatorname*{arg\,max}_{M_i, M_i \neq M_{maxap}} \{p(M_i|\mathbf{z})\} \tag{2}$$

Based on these responses we are able to look for two types of knowledge gaps. A low response of the best model M_{maxap} indicates that the particular region of the feature space is not well modeled. In this case the measure for certainty could be expressed simply as

$$C(\mathbf{z}) = p(M_{maxap}|\mathbf{z}) \; . \tag{3}$$

However, even if the response of the best model is not low, but is on the other hand similar to the response of the second best model, we can consider the particular region in the feature space as a knowledge gap. The reason for such a conclusion is that the models are ambiguous, and can not provide reliable classification. In this case we can express the certainty as

$$C(\mathbf{z}) = p(M_{maxap}|\mathbf{z})/p(M_{maxap2}|\mathbf{z}) \; , \tag{4}$$

which is very similar measure to the margin sampling, known from the active learning literature. In this literature, also the third certainty measure in uncertainty sampling is often used, which is based on the entropy:

$$C(\mathbf{z}) = \sum_{i=1}^{k} p(M_i|\mathbf{z}) \log(M_i|\mathbf{z}) \; . \tag{5}$$

Once the certainties of all samples are estimated, the deepest knowledge gap can be found by looking for the most uncertain sample:

$$\mathbf{z}^* = \arg\min_{\mathbf{z}} \{C(\mathbf{z})\} \ . \qquad (6)$$

In this paper we do not commit to specific underlying learning methods that are actually used to train classifiers using training instances; we rather focus on a higher layer of the proposed learning framework. We try to be as agnostic with respect to the underlying learning method as possible. Any incremental learning method and the classifier that can return posterior probability over all possible classes can, in principle, be used.

3.3. Directed uncertainty sampling

In the case of introspective knowledge gap detection the learner has to sample the feature space to produce the hallucinated instances. This sampling cannot be random, especially in the case of high-dimensional feature spaces; it should be driven by the structure of the current knowledge and by the output of previous classifications.

We have designed the following Monte Carlo–like method to deal with possible high-dimensional feature spaces. In our method the learner executes the following four steps.

1. Take M random samples from the feature-space. These are the first collected samples.

2. Calculate the depth of knowledge gap for all collected samples (as described in Section 3.2).

3. Choose M samples from the set of collected samples; tend to choose samples with deep knowledge gaps. Around each of M samples take a new Gaussian sample from the feature space, calculate the depth of knowledge gap for the sample and add it to the set of collected samples.

4. Repeat the previous step N times or until convergence.

This simple algorithm does not guarantee to find the optimal solution, i.e., the global certainty minimum. But in this task finding the optimal solution is not really necessary; what we actually want is a good enough solution, which is always provided by the algorithm. In fact, this algorithm very often finds knowledge gaps that are very close to the optimal ones.

4. Knowledge gap communication

The second problem we address in this paper is the one of knowledge gap communication. How can the learner communicate the knowledge gap? How can it notify others (e.g., the teacher) about what kind of information is needed? Also in this case we can distinguish between the extrospective and the introspective case.

In **extrospective** knowledge gap communication, the learner refers to existing training instances. It therefore selects one of the available training instances from the pool of instances, or it labels a generated instance based on the label of one of the existing instances.

In the case of extrospective knowledge gap detection such extrospective communication of the detected gap is straightforward, since the gap is located in one of the instances that actually exist.

In the case of introspection, however, it is more difficult to communicate this gap, since the detection is not based on a particular situation the robot could refer to. If the learner is able of inverse mapping from the feature values to the action parameters (as opposed to the mapping from action parameters to feature values, which is a part of the regular feature extraction process), then we say that the knowledge gap communication is also **introspective**. These kinds of learning scenarios can take full advantage of introspective knowledge gap detection, which may lead to a very efficient selection of training instances. If this is not possible, the learner can communicate the knowledge gap by referring to instances which he has access to (from the pool of training instances). In this case, therefore, the knowledge gap communication is extrospective.

5. Active learning strategies

Based on the way the knowledge gap is detected and on the means of how the request for information is communicated to the teacher, we distinguish four different active learning strategies (as also listed in Table 1):

- *LDeeSel (extrospective-extrospective instance selection).* The learner has access to a pool of non-labeled training instances. The learner measures potential knowledge gaps for feature vectors of introduced instances. Subsequently, the learner asks the teacher for the label of one of the not yet chosen instances for which the deepest knowledge gaps have been detected. In this variant the learner operates on the pool of training instances and does not sample the feature space for detecting possible more significant knowledge gaps.

- *LDieSel (introspective-extrospective instance selection).* The learner samples the feature space and tries to detect gaps in its knowledge independently of the teacher. When found, the learner looks for the most similar training instance from the pool of the instances and communicates it as a knowledge gap. This strategy should be used when the learner can not communicate the detected feature gap directly (i.e. it can not map the feature values into the input space).
- *LDieGen (introspective-extrospective instance generation).* The knowledge gap detection is performed in a similar way as in the previous case. However, the label of the feature representing the detected knowledge gap is then obtained by checking the label of the most similar training instance in the pool of instances. Note that in this case the new training instance consists of the feature vector of the knowledge gap and the label of the most similar instance. The advantage of this approach in comparison to *LDieSel* strategy is that it tends to choose *the most* informative training instances for learning and not just some approximation from the existing pool. The weakness of this strategy, however, is a risk that the nearest training instance does not necessarily belong to the same class as the new instance (i.e. the assigned label may be incorrect).
- *LDiiGen (introspective-introspective instance generation).* This is the ultimate strategy. When this one is possible, the learner introspectively communicates the knowledge gap to the teacher and asks for the label. The teacher replies with the class label of this exact, newly generated training instance. Obviously, in this case an inverse mapping from the feature space to the input space should be possible.

The principles of *LDeeSel*, *LDieSel*, and *LDiiGen* learning strategies are well known in the field of active learning. However, to the best of our knowledge, *LDieGen* provides a *novel approach for choosing informative training instances*, at least from the perspective of human-robot interaction.

Table 1. Four different active learning strategies.

strategy	detection	communication
LDeeSel	extrospective	extrospective
LDieSel	introspective	extrospective
LDieGen	introspective	extrospective
LDiiGen	introspective	introspective

When using the above presented active learning strategies, models built by the underlying learning method are not reliable at the beginning of the learning process. Instead of leaning on unreliable knowledge, a reasonable solution seems to be some level of randomness, i.e. choosing a random training instance with some probability.

6. Experimental results

In this section we present the results of the evaluation of different learning strategies, presented in the previous section. The evaluation has been performed on three data sets: spatial templates, colors, and the UCI Letter recognition data set.

For the underlying learning method we used the odKDE algorithm we have previously developed [9], and in Sec. 6.3 also the mcpIncSVM based on [6]. Both methods are able of incremental learning by updating the representations with one training instance at a time. The corresponding classifiers conveniently calculate posterior probability for all learned concepts, therefore enabling simple detection of knowledge gaps as described in Section 3.

At the beginning of the learning process we have first trained the learning method with a small batch of initial learning instances to initialise the models. Afterwards we have continued to add one instance at a time, according to different learning strategies. At every step we have evaluated the quality of obtained models by trying to classify all test instances. All experiments have been repeated several times and the results have been averaged across all runs.

We begin this section with the summary of some teacher-driven learning strategies since they will be used as a baseline in comparison with the proposed active learning strategies.

6.1. Teacher-driven learning strategies

In the evaluation we used three strategies for teacher-driven selection of training instances, similar to those presented in [14]:

- *TDseq.* The simplest way of presenting the learning instances is to present them in a sequential order, one by one: first all of the instances from the first class, then all of the instances from the second and so on. In such a setting, the learner passively receives training instances.

- *TDrnd.* The teacher is showing not yet chosen training instances to the learner in a random order. The learner, again, passively accepts offered training instances.
- *TDfdb.* In this strategy the feedback from the learner is taken into account. The teacher shows a number of exam instances to the learner and then provides as a training instance one of the instances that was not recognised correctly. Clearly such a strategy requires more effort from the teacher and the learner, however, it should lead to better results and more efficient learning in terms of the number of knowledge updates. However, all exam instances also have to be labeled. The labels are not required by the learner but are needed by the teacher himself, therefore in this strategy the number of labeled training instances required could be high.

6.2. Learning spatial templates

Initially we have tested performance of the studied strategies on the data set of spatial templates which had been described in detail in [14]. This set contains 2 attributes and 3 classes and is the only data set of the three in our paper which allows us to compare all seven strategies. The other two (as also the vast majority of real-life domains) do not provide us with possibility of mapping from (low-dimensional) feature space to (high-dimensional) space of learning instances, and the use of fully introspective LDgen is not possible there. Fig. 1 contains averaged results of the experiment on 121 training instances with the odKDE learner over 100 runs. The training images were presented to the learning framework one by one, as selected by different learning strategies. As we may observe, all four variants of the learner-driven active strategies achieve superior recognition rate in comparison to teacher-driven strategies. *LDiiGen* strategy demonstrates the best behavior rising to 99% recognition rate after only 40 training instances. While *LDieGen* progresses quite fast in the beginning it is the only strategy which does not reach perfect classification performance at the end of learning. The reason is in incorrect labeling of some training instances, due to the risk mentioned in Sec. 5.

One thing to note is also a very good performance

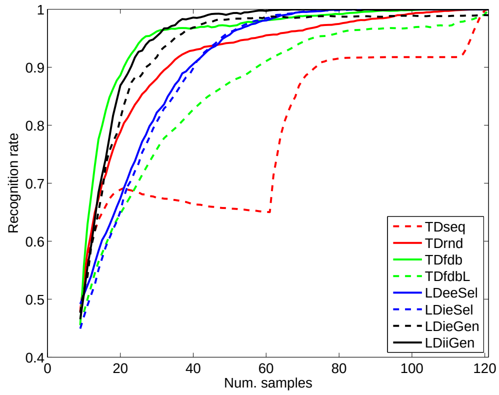

Figure 1. Recognition rate with respect to the number of training instances on the data set of spatial templates.

of the *TDfdb* strategy in terms of the number of instances used for updating the model (the solid green curve). However, the number of images that had to be labelled, in order to be presented as exam instances, was significantly higher (the dashed green *TDfdbL* curve).

6.3. Learning colors

We have made further evaluation on the color data set from [14]. The set consists of 1094 images of 129 objects. We have used 820 images (75%) as training instances and 274 images (25%) for testing the recognition performance. Eight colours were being taught, based on the H, S, and L features from HSL values of the dominant colour. In each run the set of images has been randomly split into training and test sets. We evaluated the strategies using two underlying learning algorithms, odKDE and mcpIncSVM[1]. The mcpIncSVM was set to have linear kernel.

The experimental results are presented in Fig. 2 and show the evolution of the recognition rate with respect to the number of training instances. The results have been averaged over 100 runs. In our scenario not only the final recognition rate is important but also when certain recognition rate is reached. Although the difference between strategies is not significant in view of the final recognition rate, the advantage of our strategies is that *they achieve certain recognition rate (much) earlier than random (blind) strategies*. Fig.3 depicts the number of training instances required to achieve a certain level of recognition rate for all six strategies. This level has been set to 99% of the final result of the baseline

[1]The source code is available at http://www-ti.informatik.uni-tuebingen.de/ spueler/mcpIncSVM/.

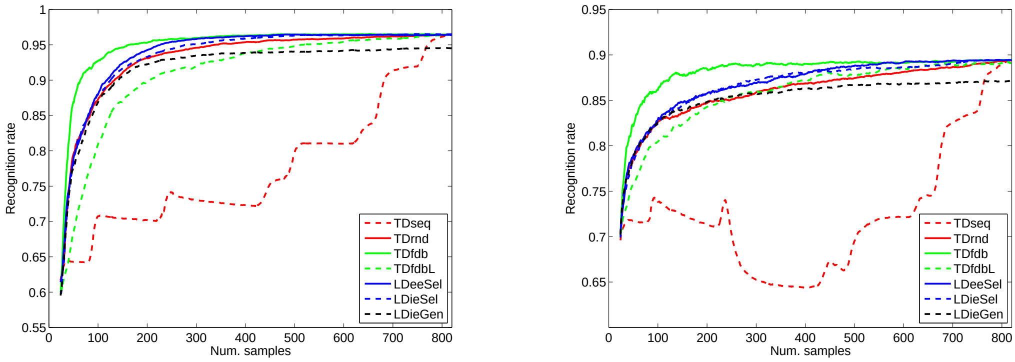

(a) odKDE (b) mcpIncSVM

Figure 2. Recognition rate in the domain of learning qualitative descriptions of object colours.

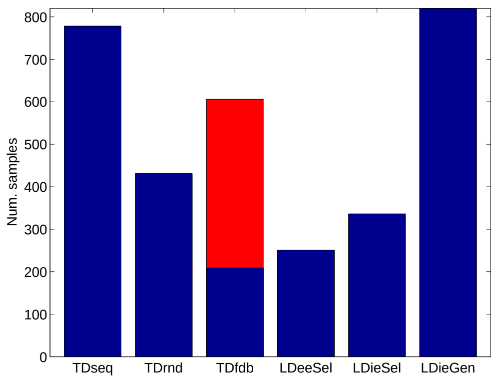

Figure 3. Number of training instances needed to achieve the same recognition rate in the case of learning colors.

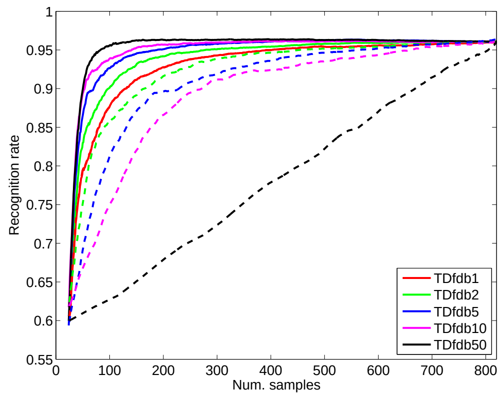

Figure 4. Recognition rate for different number of exam instances: 1, 2, 5, 10, and 50

TDrnd method. For the *TDfdb* strategy both results are shown; how many training instances have been used for updating the knowledge, as well as how many training instances have been labeled (when showing the exam instances). The figure shows that the *LDeeSel* and *LDieSel* active learning strategies clearly outperform the teacher-driven strategies.

In addition to these experiments, we have also verified the influence of the *number of exam instances (N_{exams})* on performance of the *TDfdb* strategy. This number determines how many learning instances are used for intermediate evaluations of the learner's knowledge. When N_{exams} equals 1, *TDfdb* strategy operates identical to random sampling. As the value of N_{exams} increases, the probability that at least one instance will be incorrectly classified – meaning that we will be able to take such an instance as a new training instance – also grows. The results are shown in Fig. 4. The experiments have been carried out with the odKDE learning method on the color data set with 820 training instances. Introducing more exam instances to the learner obviously speeds up the learning process. However, it has to be noted that for *TDfdb* to achieve certain recognition rate there have to be (considerably) more labeled instances available as not only training instances but also exam instances have to be labeled (denoted as the dashed curves in Fig. 4). *This experiment nicely demonstrates that by an optimal selection of the training instances one can speeds up the learning process enormously*.

6.4. Experiment on the UCI Letter database

We have evaluated the strategies also on a considerably more extensive data set – UCI Letter recognition. This set is a high-dimensional data set with 16 attributes and 26 classes, where the data may not be fully visualized. There are 16000 training and 4000 test instances available in separate sets. In Fig. 5 results of testing the strategies on the UCI Letter

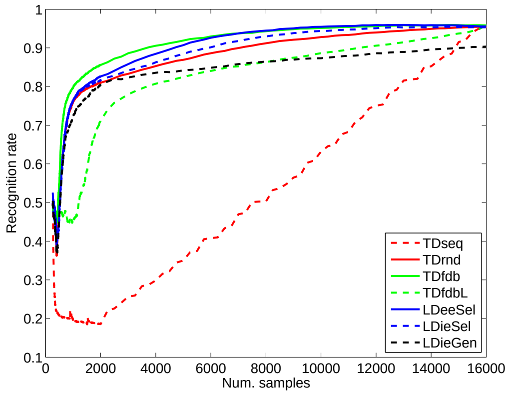

Figure 5. Recognition rate in the alphabet letters domain.

recognition data set using the odKDE learner are presented. The results have been averaged over 15 runs. It is important to notice that the results are analogue to the results on lower-dimensional color data set. These results, therefore, also confirm superior performance of the active learning strategies.

7. Conclusion

In this paper we addressed the problem of acquiring categorical knowledge from the active learning perspective. We focused on the problems of knowledge gap detection and communication. For both of these problems we proposed introspective and extrospective solutions, and based on these solutions, we presented four variants of active learning strategies.

The experimental results show that the analyzed active learning strategies effectively reduce the amount of training data that have to be introduced to the learner when reaching the required recognition rate. The results also demonstrate that the concerns about incorrect labelings in the proposed *LDieGen* strategy have came true. Further thorough test will be conducted on different domains, however, for the time being, this strategy can not be successfully employed to the problem of knowledge gap detection and communication.

References

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

[2] M. Cakmak, N. DePalma, R. I. Arriaga, and A. L. Thomaz. Exploiting social partners in robot learning. *Autonomous Robots*, 29(3-4):309–329, 2010. 2

[3] C. Chao, M. Cakmak, and A. L. Thomaz. Transparent active learning for robots. In *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 317–324, 2010. 2

[4] S. Chernova and M. Veloso. Confidence-based policy learning from demonstration using Gaussian mixture models. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, 2007. 2

[5] J. de Greeff, F. Delaunay, and B. T. Human-robot interaction in concept acquisition: a computational model. In *Proceedings of the 2009 IEEE 8th International Conference on Development and Learning*, pages 1–6, Washington, DC, USA, 2009. 2

[6] C. Diehl and G. Cauwenberghs. Svm incremental learning, adaptation and optimization. In *Proceedings of the 2003 International Joint Conference on Neural Networks*, page 26852690, 2003. 5

[7] Y. Fu, X. Zhu, and B. Li. A survey on instance selection for active learning. *Knowledge and Information Systems*, 2012. 2

[8] X. Hu, L. Wang, and B. Yuan. Querying representative points from a pool based on synthesized queries. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, pages 1–6, 2012. 2

[9] M. Kristan and A. Leonardis. Online discriminative kernel density estimation. In *International Conference on Pattern Recognition*, pages 581–584, Istanbul, Turkey, 23-26 August 2010. 5

[10] S. M. Nguyen and P.-Y. Oudeyer. Interactive learning gives the tempo to an intrinsically motivated robot learner. In *Proceedings of the IEEE-RAS International Conference on Humanoid Robots*, 2012. 2

[11] F. Olsson. A literature survey of active machine learning in the context of natural language processing. Technical report SICS Technical report T2009:06, Swedish Institute of Computer Science, 2009. 2

[12] P.-Y. Oudeyer, F. Kaplan, and V. V. Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2):265–286, 2007. 2

[13] B. Settles. Active learning literature survey. Technical report Computer sciences technical report 1648, University of Wisconsin-Madison, 2010. 2

[14] 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. 2, 5, 6

[15] A. L. Thomaz and M. Cakmak. Learning about objects with human teachers. In *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction*, pages 15–22, 2009. 2