Please quote as: Calma, A.; Leimeister, J. M.; Lukowicz, P.; Oeste-Reiß, S.; Reitmaier, T.; Schmidt, A.; Sick, B.; Stumme, G. & Zweig, K. A. (2016): From Active Learning to Dedicated Collaborative Interactive Learning. In: 4th International Workshop on Self-Optimisation in Autonomic and Organic Computing Systems (SAOS), Berlin.

From Active Learning to Dedicated Collaborative Interactive Learning

Adrian Calma*, Jan Marco Leimeister†, Paul Lukowicz‡, Sarah Oeste-Reiß†, Tobias Reitmaier*, Albrecht Schmidt§, Bernhard Sick*, Gerd Stumme¶ and Katharina Anna Zweig‖

* Intelligent Embedded Systems, University of Kassel, Germany, Email: {adrian.calma|tobias.reitmaier|bsick}@uni-kassel.de
† Information Systems, University of Kassel, Germany, Email: {leimeister|oeste-reiss}@uni-kassel.de
‡ Embedded Intelligence, DFKI, Kaiserslautern, Germany, Email: paul.lukowicz@dfki.de
§ Human-Computer Interaction, University of Stuttgart, Germany, Email: albrecht.schmidt@vis.uni-stuttgart.de
¶ Knowledge and Data Engineering, University of Kassel, Germany, Email: stumme@cs.uni-kassel.de
‖ University of Kaiserslautern, Germany, Email: zweig@cs.uni-kl.de

Abstract—Active learning (AL) is a machine learning paradigm where an active learner has to train a model (e.g., a classifier) which is in principle trained in a supervised way. AL has to be done by means of a data set where a low fraction of samples (also termed data points or observations) are labeled. To obtain labels for the unlabeled samples, the active learner has to ask an oracle (e.g., a human expert) for labels. In most cases, the goal is to maximize some metric assessing the task performance (e.g., the classification accuracy) and to minimize the number of queries at the same time. In this article, we first briefly discuss the state-of-the-art in the field of AL. Then, we propose the concept of dedicated collaborative interactive learning (D-CIL) and describe some research challenges. With D-CIL, we will overcome many of the harsh limitations of current AL. In particular, we envision scenarios where the expert may be wrong for various reasons. There also might be several or even many experts with different expertise who collaborate, the experts may label not only samples but also supply knowledge at a higher level such as rules, and we consider that the labeling costs depend on many conditions. Moreover, human experts may even profit by improving their own knowledge when they get feedback from the active learner.

1 Introduction

Machine learning is based on sample data. Sometimes, these data are labeled and, thus, models to solve a certain problem (e.g., a classification or regression problem) can be built using targets assigned to input data of the model. In other cases, data are unlabeled (e.g., for clustering problems) or only partially labeled. Correspondingly, we distinguish the areas of supervised, unsupervised, and semi-supervised learning. In many application areas (e.g., industrial quality monitoring processes, intrusion detection in computer networks, speech recognition, or drug discovery) it is rather easy to collect unlabeled data, but quite difficult, time-consuming, or expensive to gather the corresponding targets. That is, labeling is in principal possible, but the costs may be enormous.

This article focuses on a substantial advancement of *active learning (AL)*, a machine learning paradigm which is related to semi-supervised learning.

AL starts with an initially unlabeled or very sparsely labeled set of samples and iteratively increases the labeled fraction of the training data set by "asking the right questions". These questions are answered by humans (e.g., experts in an application domain), by simulation systems, by means of real experiments, etc., often modeled by an abstract "oracle". Basically, the "idealized" goal of AL is to obtain a model (e.g., a classifier or a regression model) with (almost) the performance of a model trained with a fully labeled data set at (almost) the cost of an unlabeled data set. In the following, the framework consisting of a knowledge model with machine learning techniques, pools of unlabeled and (when available) labeled data, and a unit that selects unlabeled samples for queries and controls the training of the model will be referred to as *active learner*.

Often, the following assumptions are made in AL:

- The labeling process starts with an initially labeled set of samples and assumes well-defined learning tasks (e.g., the number of classes is given in advance).
- The oracle labels single samples or sets of samples (called queries depending on the AL type, see Section 2) presented by an active learner.
- The oracle is omniscient and omnipresent, i.e., it always delivers the correct answers and it is always available.
- The labeling costs for all samples are identical.

These assumptions impose severe limitations for many applications. For this reason, the following key challenges regarding an extension of AL can be identified:

Challenge 1: An expert may be (more or less) wrong for various reasons, e.g., depending on her/his experience in the application domain (we still assume we have no malicious or deceptive experts that cheat or attack the active learner).

Challenge 2: There might be several or even many experts with different expertise (e.g., different degree or kind of experience) who may collaborate to provide the active learner with labels.

Challenge 3: The experts may label not only samples but also other kinds of queries to provide knowledge at a higher level (e.g., by assigning a conclusion to a presented premise of a rule).

Challenge 4: The labeling costs depend on many conditions, e.g., whether samples or rules are labeled, on the location of samples in the input space of a model (i.e., making labeling more or less difficult), the degree of expertise of a human, etc.

Challenge 5: The experts want to benefit from the active learner by receiving feedback in order to improve their own knowledge.

Challenge 6: The learning task may require a "lifelong" learning of the system (e.g., if the process or environment from which the measured data originate is time-variant).

Moreover, there may be several tasks that have to be fulfilled at the same time (e.g., movies that are assessed regarding several criteria) and different kinds of information sources (e.g., human experts and simulation systems).

The above challenges 1 to 6 will be discussed in more detail in this article.

We envision *dedicated collaborative interactive learning (D-CIL)* approaches where the above limitations no longer hold. That is, we will develop future AL processes that are

- **interactive** in the sense that there is an information flow not only from humans to the active learner but also vice versa and not only in the form of labels but in various, more complex ways,
- **collaborative** in the sense that various experts collaborate to support the active learner with information, and
- **dedicated** in the sense that the learning process is clearly defined (such as, e.g., in an industrial quality monitoring process), the group of human experts is rather small and the collaborate over a longer period of time.

As an example for a D-CIL application, consider an industrial quality monitoring problem we addressed some years ago [1]: At a last stage of a silicon wafer fabrication process the wafers have to be checked for possible defects by means of visual inspection. Anomalies such as abrasions, cracks, scratches, or dust particles must be identified in images of wafers taken under different lightning conditions in order to sort out unusable wafers. Conspicuous regions on a wafer can rather easily be detected using appropriate image processing techniques. The classification of these regions, however, is rather difficult. Human experts often fail, they disagree, or their assessment criteria vary over time, depending on parameters such as fatigue, motivation, experience, etc. that may not be known in detail. How can such a classification process be automated using, e.g., features computed from the images such as the length-width ratio describing a conspicuous region? It is cheap to obtain a large amount of images of conspicuous regions, but time-consuming and error-prone to get the corresponding labels. The solution could be a D-CIL approach as sketched above.

The field of AL has awoken the interest of many companies, such as Microsoft, IBM, Siemens, AT&T, Mitsubishi, or Yahoo. Publications of those companies show that AL can be successfully utilized to solve problems such as in text classification [2], detecting and filtering abusive user-generated content on the Web [3], sentiment analysis of texts [4], speech recognition [5], [6], image classification [7], drug design [8], [9], detection of plant diseases [10], malware detection [11], or recommender systems [12], [13].

Altogether, we can be sure that there will also be an increasing interest in AL and, as many limitations of AL are abolished, in D-CIL, too. We even believe that many problems arising in the field of *Big Data* may be solved relying on D-CIL approaches. D-CIL techniques may also advance more technical fields such as the field of *self-organizing and adaptive systems* by increasing their degree of autonomy in learning tasks. D-CIL may even be seen as a first step towards CIL in open-ended environments, an approach we call *opportunistic collaborative interactive learning (O-CIL)*. There, many technical devices (e.g., open, heterogeneous, dynamic systems such as mobile devices, e.g., smartphones) will interact in the sense sketched above by actively collecting information from other devices, from humans, or from the Internet, for instance. Although we focus on D-CIL in this article, we will briefly outline O-CIL in Section 5.

In the remainder of this article, we first present some foundations of AL in Section 2 and define D-CIL in Section 3. In Section 4 we investigate the above challenges in more detail and briefly discuss possible solutions. Finally, Section 5 concludes the article by taking a look at possible application fields and at O-CIL.

2 Overview of Active Learning Foundations

The motivation of AL is that obtaining plenty of unlabeled data is often quite cheap, while acquiring labels is a task with high costs (monetary or temporal). AL is based on the hypothesis that a process of (iteratively) asking an *oracle* for labels and refining the current model can be realized in a way such that

- the performance of the resulting model is comparable to the performance of a model trained on a fully labeled data set and
- the overall labeling costs to obtain the final model are much lower (typically simply measured by the number of labels).

Actually, to address the previous requirements it is possible to build an *active learner* that is based on a complementary pair of *model* (e.g., a classifier) and *selection strategy*. With a selection strategy, the active learner decides whether a sample is *informative* and asks the oracle for labels. Here, informative means that the active learner expects a (high) performance gain if this sample is labeled (similarly, a set of samples can also be called informative).

Basically, various kinds of models can be used for AL, but the selection strategy should always be defined depending on the model type (e.g., whether support vector machines, neural networks, probabilistic classifiers, or decision trees are chosen to solve a classification problem). AL can be used for classification problems (e.g., [14], [15], [16], [17]), to modify the results of clustering (e.g., [18]), to solve regression problems (e.g., [19], [20], [21], [22]), or for feature selection (e.g., [8], [9]).

In the field of active learning (AL), *membership query learning* (MQL) [23], *stream-based active learning* (SAL) [24], and *pool-based active learning* (PAL) [25] are the most important paradigms (see Figure 1a).

In an MQL scenario, the active learner may query labels for any sample in the input space, including samples generated by the active learner itself. Lang and Baum [26], for example, describe an MQL scenario with human oracles to classify written digits. The queries generated by the active learner turned out to be some mixtures of digits, therefore being too difficult for a human to provide reliable answers.

An alternative to MQL is SAL, which assumes that obtaining unlabeled samples generates low or no costs. Therefore, a sample is drawn from the data source and the active learner decides whether or not to request label information. In SAL the source data is scanned sequentially and a decision is made for each sample individually. Typically, SAL selects only one sample in each learning cycle.

For many practical problems a large set of unlabeled samples may be gathered inexpensively and this set is available at the very beginning of the AL process. This motivates the PAL scenario. The learning cycle of PAL is depicted in Figure 1b. Typically, PAL starts with a large pool of unlabeled and a small set of labeled samples. On the basis of the labeled samples the knowledge model (e.g., a classifier) is trained. Then, based on a selection strategy, which considers the "knowledge" of the active learner, a query set of unlabeled samples is determined and presented to the oracle (e.g., a human domain expert), who provides the label information. The set of labeled samples is updated with the newly labeled samples and the learner updates its knowledge. The learning cycle is repeated until a given stopping condition is met.

In the remainder of this article we focus on PAL for classification problems. This is only done to simplify the discussion of the challenges. Basically, MQL and SAL suffer from the same limitations and they will benefit from D-CIL. Also, many of the solution ideas for classification problems may be transferred to other kinds of problems such as regression.

A selection strategy for PAL has to fulfill several tasks, two of which shall be given as an example: At an early stage of the AL process, samples have to be chosen in all regions of the input space covered by data (*exploration phase*). At a late stage of the AL process, a fine-tuning of the decision boundary of the classifier has to be realized by choosing samples close to the (current) decision boundary (*exploitation phase*). Thus, "asking he right question" (i.e., choosing samples for a query) is a multi-faceted problem and various selection strategies have been proposed and investigated. We want to emphasize that a successful selection strategy has to consider structure in the (un-)labeled data.

Typically, very limiting assumptions (cf. Section 1) are made concerning the oracle and the labeling costs (omniscient, omnipresent oracle that labels samples on a fixed cost basis). Moreover, some other aspects of real-world problems are often more or less neglected by current research:

- In real-world applications, AL has often to start "from scratch", i.e., with no labels at all. This requires sophisticated selection strategies with different behaviors at different phases of the AL process.
- Parameters of the active learner (including parameter of training algorithms for the classifier and the selection strategy) cannot be found by trial-and-error. AL only allows for "one shot".

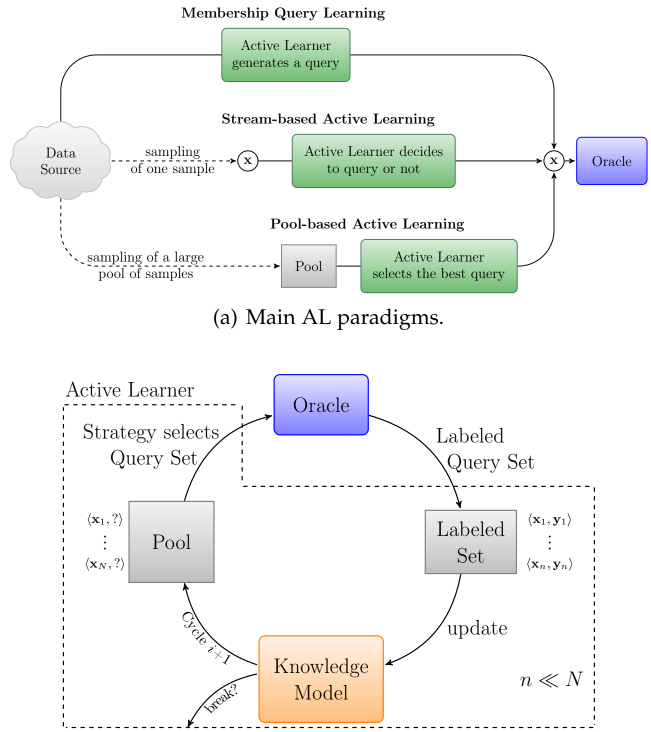

(a) Main AL paradigms.

(b) Learning cycle of standard PAL.

Fig. 1. Overview of main AL scenarios with focus on PAL.

There are a number of articles that assess the state-of-the-art in AL:

- A general introduction to AL, including a discussion of AL scenarios and an overview of query strategies is provided in [27].
- A detailed overview of relevant PAL techniques is part of [14]. In addition to single-view/single-learner methods, alternative approaches are outlined: multi-view/single-learner, single-view/multi-learner, and multi-view/multi-learner.
- For certain problem areas it makes sense to use AL in combination with semi-supervised learning (SSL). AL techniques that integrate SSL techniques are presented in [15].
- Work that uses AL in combination with support vector machines (SVM) for solving classification problems is summarized in [28], [29].

3 Characterization of Dedicated Collaborative Interactive Learning

In this section we describe our vision of future AL that we call *dedicated collaborative interactive learning* (D-CIL).

To overcome the unrealistic limitations made by conventional AL techniques (cf. Sections 1 and 2) we have to integrate multiple "uncertain oracles" (e.g., human domain experts) into the AL process (see Figure 2). That is, these oracles possibly make errors due to various causes, e.g., they are differently experienced in coping with the learning task, or their work quality depends on daily condition, motivation, etc. Therefore, D-CIL explicitly models information uncertainty, i.e. uncertainty regarding samples, labels, or parametrization of models. This uncertainty is then taken into account when either (1) the expertise of the human domain experts has to be identified or (2) their knowledge is required to provide labels.

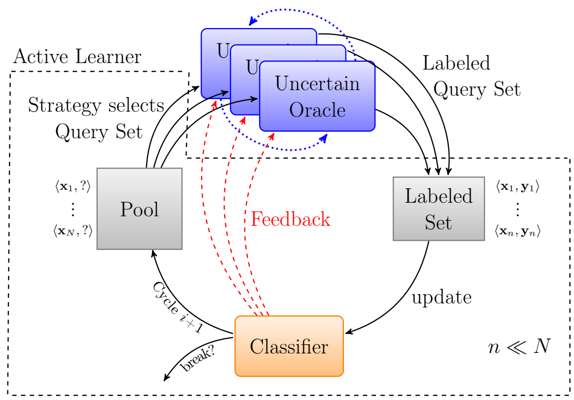

Fig. 2. Learning cycle of D-CIL with multiple collaborating uncertain oracles (human experts, simulation systems, etc.).

In real-world applications D-CIL has to start "from scratch", e.g., without any label information. We assume that, during the AL process more and more (uncertain) labels become available, but no "ground truth". Therefore, the ***collaboration*** of various human experts will be essential for the success of D-CIL. Experts not only collaborate with the active learner. Other kinds of collaboration could be the mutual support of two or more experts, according to the idea of pair programming, in order to achieve higher accuracies in solving the learning task. These collaboration processes are indicated by the dotted, blue arrows in Figure 2.

D-CIL integrates multiple experts into the AL process and models their expertise explicitly. Consequently, D-CIL needs more sophisticated selection strategies than those used in conventional AL. In D-CIL, the selection strategy has not only to choose the most informative samples (from the pool of unlabeled samples) considering the current knowledge of the model (here, a classifier) in order to build a query set in each AL cycle, but also to decide which experts shall be queried depending on their expertise. That is, we have an exploration/exploitation problem again. In addition, in D-CIL we query not only samples but also knowledge at higher abstraction levels (e.g. premises of rules), such that more sophisticated cost schemes are needed, too.

Particularly, D-CIL differs from conventional AL in the fact that D-CIL gives targeted feedback (cf. the dashed, red arrows in Figure 2) which will improve the experts' level of expertise. To emphasize this difference, we use the term ***interactive*** learning. Therefore, the main goal of D-CIL is to maximize the accuracy of the actively trained classifier *and* to maximize the benefit of the human domain experts with minimal costs.

We assume that the humans who collaborate to solve a learning task are actually have the knowledge about that specific learning task (e.g., a certain industrial problem). We use the term "expert" to emphasize this fact. These experts are assumed to be motivated to collaborate over a longer time period, i.e., they are regarded as being ***dedicated*** to the learning task. The learning task itself may be time-variant, i.e., it may change its characteristics over time. An example are new classes that have to be detected or classes that are not relevant any more (resulting in novel or obsolete clusters of samples in the input space of a classifier). So, we need online learning techniques to solve such problems.

Altogether, D-CIL targets a specific class of applications where we may assume that the following assumptions hold: We have rather small, quite homogeneous groups of experts that collaborate over a longer period of time to solve a specific application problem. Though we act on these assumptions for D-CIL, we will abandon them for O-CIL which will be sketched in Section 5.

4 Challenges for Future D-CIL Research

In the field of D-CIL, we will answer many questions, most of which caused by the harsh limitations of AL sketched in Section 1. In the following, we examine the six key challenges (see Section 1).

4.1 Challenge 1: Uncertain Oracles

In a first step, we address the obvious fact that oracles are not always right. In principal, labels are subject to uncertainty. Here, the meaning of the term uncertainty is adopted from [30]. That is, "uncertain" is a generic term to address aspects such as "unlikely", "doubtful", "implausible", "unreliable", "imprecise", "inconsistent", or "vague".

In real-world applications, labels may come from various sources, often but not always humans. Therefore, a new problem arises: The labels are subject to uncertainty for different reasons. For example, the performance of human annotators depends on many factors: e.g., expertise/experience, concentration/distraction, boredom/disinterest, fatigue level, etc. Furthermore, some samples are difficult for both experts and machines to label (e.g., samples near the decision boundary of a classifier). Results of real experiments or simulations may be influenced, too: There may be stochasticity which is inherent to a certain process, sensor noise, transmission errors, etc., just to mention a few. Thus, we face many questions: How can we make use of uncertain oracles (annotators that can be erroneous)? How do we decide whether an already queried sample has to be labeled again? How do we deal with noisy experts whose quality varies over time (e.g., they gather experience with the task, they get fatigued)? How does remuneration influence the labeling quality of a noisy expert (e.g., if they are payed better, they are more accurate)? How can we decide whether the expert is erroneous or an observed process itself is nondeterministic?

As a starting point, we may assume that the "expertise of an expert" (i.e., the degree of uncertainty of an oracle) is time-invariant and global in the sense that it does not depend on certain classes, certain regions of the input space of the model to be learned (e.g., a classifier), etc. Then, we may ask experts for, e.g.,

- one class label with a degree of confidence,

- membership probabilities for each class (with or without confidence labels),
- lower bounds for membership probabilities (cf. [31]),
- a difficulty estimate for a data object that is labeled, or
- relative difficulty estimates for two data objects ("easier" or "more difficult" to label).

Then, we have to define appropriate ways to model that uncertainty (e.g., second-order distributions over parameters of class distributions in a probabilistic framework) and to consider it in selection strategies (e.g., with additional criteria) and for the training of a classifier (e.g., with gradual labels).

4.2 Challenge 2: Multiple Uncertain Oracles

In a second step, we address situations where several, individually uncertain oracles (e.g., several human experts with different degree of expertise) contribute their knowledge. Thus, the learning process will now rely on the collective intelligence of a group of oracles. We see this step as a first important step towards true collaboration between human experts to support such a learning process.

In various applications, different, uncertain oracles may contribute labels (cf. Figure 2). These experts may not only have different degrees of expertise. They also may have more or less expertise for different parts of the problem that has to be solved, e.g., for different classes that have to be recognized, for different regions of the input space, for different dimensions of the input space (attributes), etc. Also, experts collaborate with others, which stimulates a learning from others and results in a knowledge gain for the expert. Now, we face many new questions: What are appropriate mechanisms to identify the expertise of the human expert? Which are the criteria for identifying the "optimal" human expert? Which experts should collaborate with each other in a labeling process in order to constitute a high-performance group? How can exploration (identifying expertise) and exploitation phases (using the experts' knowledge) be interwoven? How can we merge uncertain information obtained from several experts? How can this process be designed in order to be independent of time and place (e.g., for experts are only available on a part time basis)?

As a starting point, we may initially assume that the "expertise of an expert" is known. We may use generative, probabilistic models, for example, to describe the individual knowledge of experts and the "global" knowledge of the active learner (cf. [14], [15]). Uncertainty may again be captured with second-order approaches. New selection strategies must then not only choose samples, but also oracles. If the expertise of an expert is not known, it must be revealed either by asking for difficulty or confidence estimates or by comparing it to the knowledge of others (e.g., by asking an expert who has to be assessed some questions with already known answers). In order to explore solutions to challenge 2, we may not only rely on real experiments with humans (cf. the field of crowdsourcing, for instance). In addition, we may also be confronted with the problem of simulation: We have to simulate several uncertain oracles with the different characteristics mentioned above.

4.3 Challenge 3: Alternative Query Types

If we have to explore the knowledge of oracles as sketched above, the costs of AL increase substantially. In the other hand, we might ask oracles such as human experts for more abstract knowledge with the goal to reduce the number of queries this way.

In many applications, active learners could ask for more "valuable" knowledge. Examples are conclusions that a human expert gives for a presented rule premise, or correlations between different features or features and classes that an expert provides in order to identify important or redundant features. Questions that arise in this context are: Which questions can be asked? How can we provide (i.e., visualize, for instance) the required information to the expert? How can we combine different kinds of expert statements, e.g., about samples, rules, relations between features, etc? How can we use this information to initialize the models that are trained or to restrict the model capabilities in an appropriate way (e.g., if features are known not to be correlated?

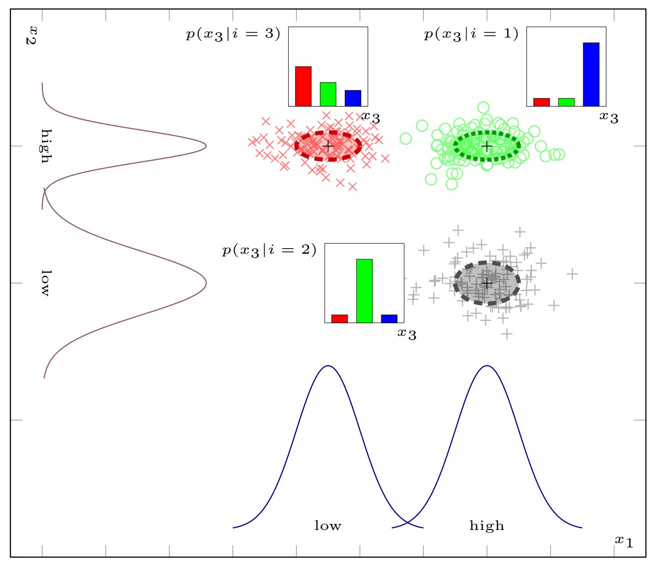

Fig. 3. Asking for conclusions of rule premises.

As a starting point, we could investigate the case of annotating rule premises with conclusions. To stay in a probabilistic framework we could obtain user-readable rule premises by marginalization of density functions from a generative process model. Figure 3 gives an example for a density model consisting of three components in a three dimensional input space. The first two dimensions x_1 and x_2 are continuous and, thus, modeled by bivariate Gaussians whose centers are described by larger crosses (+). The ellipses are level curves (surfaces of constant density) with shapes defined by the covariance matrices of the Gaussians. Here, due to the diagonality of the covariance matrices these ellipses are axes-oriented and their projection onto the axes is also shown. The third dimension x_3 is categorical with categories A (red), B (green), and C (blue). The distributions of the third dimension x_3 are illustrated by the histograms next to every component. Here, only categories with a probability strictly greater than the average are considered in rules in order to simplify the resulting rules. We assume that the components modeling sets of circles (green) and crosses (red) are already labeled, resulting in two rules for these components:

if x_1 is *low* and x_2 is *high* and x_3 is A or B
then class = red,
if x_1 is *high* and x_2 is *high* and x_3 is C
then class = green.

Now, the active learner presents the following rule premise and asks for a conclusion in form of a class assignment:

x_1 is *high* and x_2 is *low* and x_3 is B.

This information could then be used to (re-)train a classifier, e.g., in a transductive learning step.

To investigate relations between features, i.e., between input dimensions of a classifier, we may rely on statistical measures, but also adopt ideas from the field of concept exploration (cf. the field of attribute exploration in [32]).

4.4 Challenge 4: Complex Cost Schemes for Queries

In many real-world applications obtaining information may be possible at different costs, e.g., some class information is more expensive than other or the labeling costs depend on the location of the sample in the input space. This already applies to a "conventional" AL setting without the many ideas discussed above. In a D-CIL setting, considering complex cost schemes is even more important.

We must consider costs that depend on

1) *samples with their classes:* As mentioned above, labeling costs may depend on the class (e.g., some kinds of error classes in an industrial production process may be more difficult to detect than others) or on the location of the sample in the input space (e.g., samples close to the decision boundary require higher temporal effort), for instance.
2) *query types:* It is obvious that different labeling costs have to be foreseen for samples (with or without certainty estimates) and for more complex queries such as rule premises. The cost schemes have to be even more detailed in a D-CIL setting with feedback to the humans (e.g., with queries such as "Can you confirm that ...?").
3) *oracles (experts):* The costs of humans may depend on their expertise, their temporal effort, their availability (e.g., working hour may be modeled with finite costs, otherwise costs are infinite), etc.

In principle, all these costs may change over time, too. The basic questions in this context are: How can a cost schema be defined and which different types are existent? How should compensation mechanisms for the differentiated expertise of a human be designed? How can these compensation mechanisms be implemented? How must the selection strategies of an active learner be adapted?

As a starting point, we suggest to choose the first point from the list above and investigate solutions in a "classical" AL setting. Then D-CIL requires solutions for the second and third points, respectively. Mechanisms of crowdsourcing will provide additional insights. On the one hand, differentiated compensation mechanisms can be realized if a task with defined costs can be outsourced to the crowd [33]. On the other hand, the definition of the task requires additional research in the field of crowdsourcing.

4.5 Challenge 5: True Collaboration of Human Experts and Interaction with an Active Learner

In a next step we must pave the way for a true collaboration of human experts in AL, which will essentially be based on the capability of humans to learn and the ability of the active learner to provide appropriate feedback to the humans to enable them to learn themselves. Then, the new technique actually deserves to be called D-CIL.

In many applications, experts would be interested in getting feedback from an active learner, in improving their own knowledge, and sharing their expertise with others. As an important requirement, the active learner must be able to give feedback to the humans and asking for comments on such feedback. Some possible kinds of interactions of an active learner with humans are (cf. also [34]):

> The following rule appears to be very certain because ... !
> The following rule is in conflict with your knowledge because ... !
> Other experts are much less uncertain concerning the following rule than you are ... !
> Can you confirm the following rule ... ?
> Can you confirm that the following two features are not correlated ... ?
> Can you confirm that the following feature is very important ... ?
> Can you provide additional samples for the following input regions of the classifier ... ?

Some of the many new questions arising with this challenge are: How can we deal with time-invariant knowledge of oracles? Which information should be provided and how (e.g., with/without certainty estimates, restriction to "crisp" rules or not)? How must we adapt the active learner and the selection strategies? In particular, a compromise has to be found between modeling capabilities on the one hand and the abilities of humans to actually understand readable rules on the other. How do human experts change their behavior if they get feedback? How do human experts cooperate when they use a D-CIL system? How can an explicit "pooling" of experts in teams be realized? May we suggest solutions to experts? How can we realize a review mechanism for answers of experts? When and how can human experts be recruited? How can we measure the benefit of human experts or groups of experts?

As a starting point, we may stay within our probabilistic framework, consider the individual knowledge of humans (challenge 2) and present samples and rules (e.g., obtained by marginalization from density models to make them human-readable as sketched above, challenge 3) with fused statements (labels or conclusions) and certainty estimates. Then, the time-variance of human knowledge must be considered by extending the solutions from challenge 2. Altogether, the collaboration activities between humans and between humans and the active learner need to be designed in a structured and re-usable way [35], [36]. Again, the evaluation of any new, proposed techniques will be a challenge by itself.

4.6 Challenge 6: Online Learning for Time-Variant Learning Tasks

Above, we have sketched D-CIL which takes place in a time-variant environment in the sense that the knowledge of experts improves over time. But, the observed and modeled processes could be time-variant, too. That is, these processes may change slightly (e.g., due to increased wear of mechanical parts of an observed process), become obsolete, or new processes corresponding to known or to new, previously unknown, classes may arise during the application of the model. Then, a major challenge consists in developing online D-CIL techniques that cope with such effects. Altogether, we can say that in essence we have to to solve a learning task which changes over time, where we only have partial knowledge, and where knowledge is uncertain. That is, from the viewpoint of the humans and the active learner "lifelong" learning may be needed.

Questions that come up when we address this challenge are, for example: How can changes in the characteristics of the processes underlying the observed data be detected? How can new classes be considered online? How can we efficiently and effectively integrate the human experts in the process of detection and modeling?

As a starting point, we may adapt techniques from the fields of anomaly or novelty detection, obsoleteness detection, detection of concept drift or shift, or online clustering. However, these techniques are typically intended to work in a fully autonomous way, but we may again integrate the knowledge of human experts, for instance, to improve these techniques. Also, we may take a look at some existing multiple learner / multiple expert approaches, and adopt ideas from the field of SAL.

4.7 Further Challenges

Two additional challenges must be addressed as well:

Stopping Criterion: Currently, the stopping criterion in real-world applications is based on economic factors, e.g., the learner queries samples as long as the budget allows. The challenge consists in knowing when to stop querying for labels. One possibility may be to determine the point at which the cost of querying more labels is higher than costs for misclassification. Another possibility is to determine when the learner is at least as good as the group of annotators. For such a "self-stopping criterion", the active learner must be able to assess its own performance.

Performance Assessment: In AL, the performance of an active learner must be assessed by means of several criteria to capture effectiveness and efficiency of AL. For this purpose, we may use a ranked performance measure, a data utilization measure, the area under the learning curve, and a class distribution measure (see, e.g., [14], [15]). D-CIL requires additional measures, e.g., to assess the various learning costs or to evaluate the learning progress of human experts.

Apart from these challenges we still face the already discussed requirements such as "parameter-free" active learning techniques or self-adaptation of selection strategies to different phases of the active learning process.

5 Summary and Outlook

In this article, we have sketched our vision of D-CIL which will be elaborated in more detail in the near future. In this novel field, we would like to concentrate on developing models that take information uncertainty into consideration, identifying the annotators' level of expertise, making use of different levels of expertise and fusing possibly contradicting knowledge, labeling abstract knowledge, and improving the expertise of the experts. In the envisioned D-CIL scenarios, human domain experts should benefit from sharing their knowledge in the group. They should receive feedback which will improve their own level of expertise. We assume that in a D-CIL scenario the number of humans involved is rather low (e.g., they are specialists for certain industrial problems), they are more or less motivated, and they contribute their knowledge for a long term. In principal, many applications may benefit from D-CIL, for example, product quality control (e.g., deflectometry, classification of errors on silicon wafers or mirrors, analysis of sewing or garments in clothing industry), fault detection in technical and other systems (e.g., analysis of fault memory entries in control units of cars, analysis of different kinds of errors in cyber-physical systems, etc.), planing of product development processes (e.g., in drug design), or fraud detection and surveillance (e.g., credit card fraud, detection of tax evasion, intrusion detection, or video surveillance).

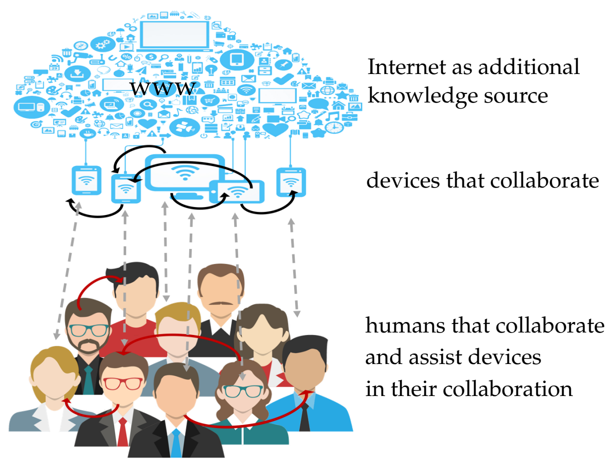

Fig. 4. Idea of Opportunistic CIL (O-CIL).

In our future world, technical systems have to evolve over time. Not all knowledge about any situation the system will face at run-time will be available at design-time. That is, the system has to detect fundamental changes in its environment and react accordingly. This requires that "never-ending" or "lifelong" learning mechanisms have to be implemented into such systems. Amongst other mechanisms (e.g., context- or self-awareness), these learning mechanisms will include appropriate active learning techniques. These future technical systems may be mobile devices, for example, that actively collect data and other kinds of information from other devices, humans (who are often non-experts in a field), or the Internet (e.g., from social networks), cf. Figure 4. These active learning processes comprise large (e.g., thousands), open (participants may leave or others may enter), and heterogeneous (e.g., different types of devices, kinds of knowledge, etc.) groups of "participants". The data that are labeled may include video, audio, text, or image data for instance. Also new kinds of human-computer interaction may come into play [37]. Each active learner built into such a future system has to make best of the available information, i.e., it has to act in an "opportunistic" way (cf. [38]). This requires an extension of D-CIL to O-CIL (*opportunistic collaborative interactive learning*), i.e., AL in open-ended environments, and also new techniques to model and analyze AL in such groups (cf., e.g. [39]).

References

[1] M. Bauer, O. Buchtala, T. Horeis, R. Kern, B. Sick, and R. Wagner, "Technical data mining with evolutionary radial basis function classifiers," *Applied Soft Computing*, vol. 9, no. 2, pp. 765–774, 2009.

[2] U. Paquet, J. V. Gael, D. Stern, G. Kasneci, R. Herbrich, and T. Graepel, "Vuvuzelas & active learning for online classification," in *Workshop on Computational Social Science and the Wisdom of Crowds*, Whistler, BC, 2010, pp. 1 – 5.

[3] W. Chu, M. Zinkevich, L. Li, A. Thomas, and B. L. Tseng, "Unbiased online active learning in data streams," in *Int. Conf. on Knowledge Discovery and Data Mining*, San Diego, CA, 2011, pp. 195–203.

[4] P. Melville and V. Sindhwani, "Active dual supervision: Reducing the cost of annotating examples and features," in *Workshop on Active Learning for Natural Language Processing*, Boulder, CO, 2009, pp. 49–57.

[5] D. Hakkani-Tür, G. Riccardi, and G. Tur, "An active approach to spoken language processing," *ACM Transactions on Speech and Language Processing*, vol. 3, no. 3, pp. 1–31, 2006.

[6] G. Tur, R. E. Schapire, and D. Hakkani-Tür, "Active learning for spoken language understanding," in *Int. Conf. on Acoustics, Speech, and Signal Processing*, Hong Kong, China, 2003, pp. 276–279.

[7] A. J. Joshi, F. Porikli, and N. P. Papanikolopoulos, "Scalable active learning for multi-class image classification," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 34, no. 11, pp. 2259–2273, 2012.

[8] R. F. Murphy, "An active role for machine learning in drug development," *Nature Chemical Biology*, vol. 7, pp. 327–330, 2011.

[9] J. D. Kangas, A. W. Naik, and R. F. Murphy, "Efficient discovery of responses of proteins to compounds using active learning," *BMC Bioinformatics*, vol. 15, no. 143, pp. 1–11, 2014.

[10] P. Schmitter, J. Behmann, J. Steinrücken, A.-K. Mahlein, E.-C. Oerke, and L. Plümer, "Aktives Lernen zur Detektion von Pflanzenkrankheiten in hyperspektralen Bildern," in *Wissenschaftlich-Technische Jahrestagung der DGPF*, Köln, Germany, 2015, pp. 398–406.

[11] N. Nissim, R. Moskovitch, L. Rokach, and Y. Elovici, "Novel active learning methods for enhanced PC malware detection in Windows OS," *Expert Systems with Applications*, vol. 41, no. 13, pp. 5843–5857, 2014.

[12] B. Lamche, U. Trottmann, and W. Wörndl, "Active learning strategies for exploratory mobile recommender systems," in *Workshop on Context-Awareness in Retrieval and Recommendation*, Amsterdam, Netherlands, 2014, pp. 10–17.

[13] H. Yu, "SVM selective sampling for ranking with application to data retrieval," in *Int. Conf. on Knowledge Discovery and Data Mining*, Chicago, IL, 2005, pp. 354–363.

[14] T. Reitmaier and B. Sick, "Let us know your decision: Pool-based active training of a generative classifier with the selection strategy 4DS," *Information Sciences*, vol. 230, pp. 106–131, 2013.

[15] T. Reitmaier, A. Calma, and B. Sick, "Transductive active learning – a new semi-supervised learning approach based on iteratively refined generative models to capture structure in data," *Information Sciences*, vol. 239, pp. 275–298, 2014.

[16] C. Constantinopoulos and A. C. Likas, "An incremental training method for the probabilistic RBF network," *IEEE Transactions on Neural Networks*, vol. 17, no. 4, pp. 966–974, 2006.

[17] Y. Zhang, H. Yang, S. Prasad, E. Pasolli, J. Jung, and M. Crawford, "Ensemble multiple kernel active learning for classification of multisource remote sensing data," *Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 2, pp. 845–858, 2015.

[18] R. Marcacini, G. Correa, and S. Rezende, "An active learning approach to frequent itemset-based text clustering," in *Int. Conf. on Pattern Recognition*, Tsukuba, Japan, 2012, pp. 3529–3532.

[19] W. Cai, Y. Zhang, and J. Zhou, "Maximizing expected model change for active learning in regression," in *IEEE Int. Conf. on Data Mining*, Dallas, TX, 2013, pp. 51–60.

[20] E. Pasolli and F. Melgani, "Gaussian process regression within an active learning scheme," in *IEEE Int. Geoscience and Remote Sensing Symposium*, Vancouver, BC, 2011, pp. 3574–3577.

[21] B. Demir and B. L., "A multiple criteria active learning method for support vector regression," *Pattern Recognition*, vol. 47, no. 7, pp. 2558–2567, 2014.

[22] F. Douak, F. Melgani, and N. Benoudjit, "Kernel ridge regression with active learning for wind speed prediction," *Applied Energy*, vol. 103, pp. 328–340, 2013.

[23] D. Angluin, "Queries and concept learning," *Machine Learning*, vol. 2, no. 4, pp. 319–342, 1988.

[24] L. Atlas, D. Cohn, R. Ladner, M. A. El-Sharkawi, and R. J. Marks, II, "Training connectionist networks with queries and selective sampling," in *Advances in Neural Information Processing Systems 2*, Denver, CO, 1990, pp. 566–573.

[25] D. Lewis and W. A. Gale, "A sequential algorithm for training text classifiers," in *ACM Conf. on Research and Development in Information Retrieval*, Dublin, Ireland, 1994, pp. 3–12.

[26] K. Lang and E. Baum, "Query learning can work poorly when a human oracle is used," in *IEEE Int. Joint Conf. on Neural Networks*, Los Alamitos, CA, 1992, pp. 335–340.

[27] B. Settles, "Active learning literature survey," University of Wisconsin, Department of Computer Science, Computer Sciences Technical Report 1648, 2009.

[28] J. Jun and I. Horace, *Active Learning with SVM*. Hershey, PA: IGI Global, 2008, vol. 3, ch. 1, pp. 1–7.

[29] J. Kremer, K. S. Pedersen, and C. Igel, "Active learning with support vector machines," *Wiley Interdisciplinary Reviews. Data Mining and Knowledge Discovery*, vol. 4, no. 4, pp. 313–326, 2014.

[30] A. Motro and P. Smets, Eds., *Uncertainty Management in Information Systems – From Needs to Solutions*. Springer US, 1997.

[31] D. Andrade and B. Sick, "Lower bound Bayesian networks – efficient inference of lower bounds on probability distributions," in *Conf. on Uncertainty in Artificial Intelligence*, Montreal, QC, 2009, pp. 10–18.

[32] G. Stumme, "Attribute exploration with background implications and exceptions," in *Annual Conf. of the Gesellschaft für Klassifikation*. Springer-Verlag, Heidelberg-Berlin, 1996, pp. 457–469.

[33] S. Zogaj, N. Leicht, B. Ivo, U. Bretschneider, and J. M. Leimeister, "Towards successful crowdsourcing projects: Evaluating the implementation of governance mechanisms," in *Int. Conf. on Information Systems*, Fort Worth, TX, 2015.

[34] T. Horeis and B. Sick, "Collaborative knowledge discovery & data mining: From knowledge to experience," in *IEEE Symposium on Computational Intelligence and Data Mining*, Honolulu, HI, 2007, pp. 421–428.

[35] J. Leimeister, *Collaboration Engineering – IT-gestützte Zusammenarbeitsprozesse systematisch entwickeln und durchfhren*. Berlin Heidelberg: Springer Gabler, 2014.

[36] S. Oeste-Reiß, M. Söllner, and J. M. Leimeister, "Development of a peer-creation-process to leverage the power of collaborative knowledge transfer," in *Hawaii Int. Conf. on System Sciences*, Kauai, HI, 2016, not yet published.

[37] A. Schimdt, "Following or leading?: the HCI community and new interaction technologies," *interactions*, no. 22, pp. 74–77, 2015.

[38] D. Roggen, G. Troster, P. Lukowicz, A. Ferscha, J. del R.Millan, and R. Chavarriaga, "Opportunistic human activity and context recognition," *Computer*, vol. 46, no. 2, pp. 36–45, 2013.

[39] M. Kaufmann and K. Zweig, "Modeling and designing real–world networks," in *Algorithmics of Large and Complex Networks*, J. Lerner, D. Wagner, and K. Zweig, Eds. Springer Berlin Heidelberg, 2009, vol. 5515, pp. 359–379.