IAL

INTERACTIVE ADAPTIVE LEARNING (IAL 2023)-Tutorial

Collocated with ECML PKDD 2023

Mirko Bunse, Georg Krempl, Alaa Tharwat Othman, Amal Saadallah

September 26, 2023

Time	Program	Presenter / Author
09:00–11:00	**Session 1: Tutorials & Poster Session**	
09:00–09:30	Tutorial Part I: Foundations of Active Learning	A. Tharwat
09:30–10:30	Tutorial Part II: Beyond Pool-Based Scenarios	G. Krempl
11:30–11:00	Poster Session	
	Coffee Break (11:00–11:30)	
11:30–13:00	**Session 2: Tutorials**	
11:30–12:00	Tutorial Part III: Beyond Active Labelling	M. Bunse
12:00–12:30	Tutorial Part IV: Towards Explainable Active Learning using Meta-Learning	A. Saadallah
12:30–13:00	Tutorial Part V: Practical Challenges and New Research Directions	A. Tharwat
	Lunch Break (13:00–14:00)	

Lunch Break (13:00–14:00)

14:00–16:00	**Session 3: Keynote & Workshop Contributions**	
14:00–14:40	Keynote: From Insights to Impact: A Metrics-Driven Active Learning Journey	A. Abraham
14:40–15:00	Active Learning for Survival Analysis with Incrementally Disclosed Label Information	K. Dedja, F.K. Nakano & C. Vens
15:00–15:15	Towards Enhancing Deep Active Learning with Weak Supervision and Constrained Clustering	M. Aßenmacher, L. Rauch, J. Goschenhofer, A. Stephan, B. Bischl, B. Roth & B. Sick
15:15–15:30	Who knows best? A Case Study on Intelligent Crowdworker Selection via Deep Learning	M. Herde, D. Huseljic, B. Sick, U. Bretschneider & S. Oeste-Reiß
15:30–15:45	Role of Hyperparameters in Deep Active Learning	D. Huseljic, M. Herde, P. Hahn & B. Sick
15:45–16:00	Challenges for Active Feature Acquisition and Imputation on Data Streams	C. Beyer, M. Büttner & M. Spiliopoulou

Coffee Break (16:00–16:30)

Coffee Break (16:00–16:30)

16:30–17:40	**Session 4: Workshop Contributions & Closing**	
16:30–16:50	Active Learning with Fast Model Updates and Class-Balanced Selection for Imbalanced Datasets	Z. Huang, Y. He, M. Herde, D. Huseljic & B. Sick
16:50–17:10	Interpretable Meta-Active Learning for Regression Ensemble Learning	O. Saadallah & Z. Rouissi
17:10–17:30	Look and You Will Find It: Fairness-Aware Data Collection through Active Learning	H. Weerts, R. Theunissen & M. Willemsen
17:30–17:40	Closing	

IAL

Foundations of Active Learning

Alaa Tharwat Othman

IAL

Content

Foundations of Active Learning

- The Motivation for Active Learning (AL)
- Basic Workflow of Active Learning
- The main components of Active Learning
- Different types of active learning
- What is the benefits of AL?
- Simple AL example

The Motivation for Active Learning

- Recently, there is huge amount of free unlabeled data (i.e., raw data) that could be collected (e.g., from IoT devices like sensors), but labeling data is
  - time-consuming
  - expensive
  - difficult to collect
- This labeling problem could be solved by reducing the size of the training data and keeping only the high-quality training data (how?)
- The active learning (AL) technique offers searches within the unlabeled data for the most informative and representative points for labelling/annotating them

Basic Workflow of Active Learning

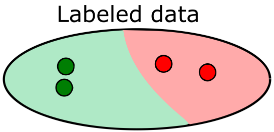

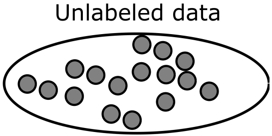

IAL

Basic Workflow of Active Learning

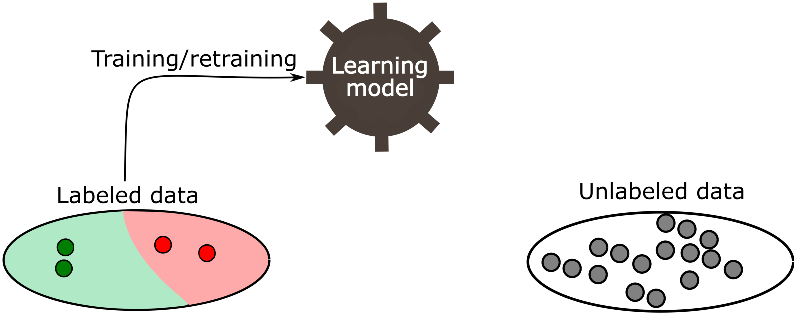

IAL

Basic Workflow of Active Learning

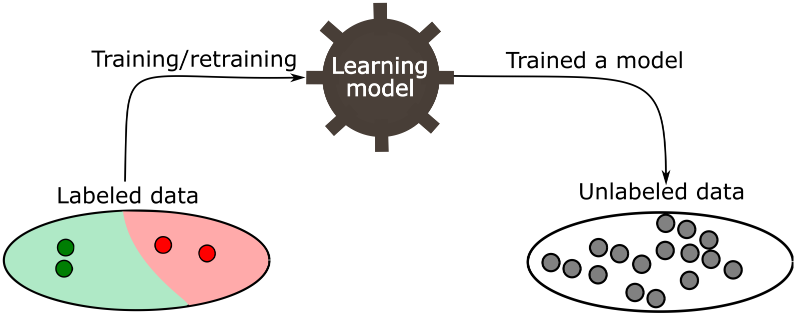

IAL

Basic Workflow of Active Learning

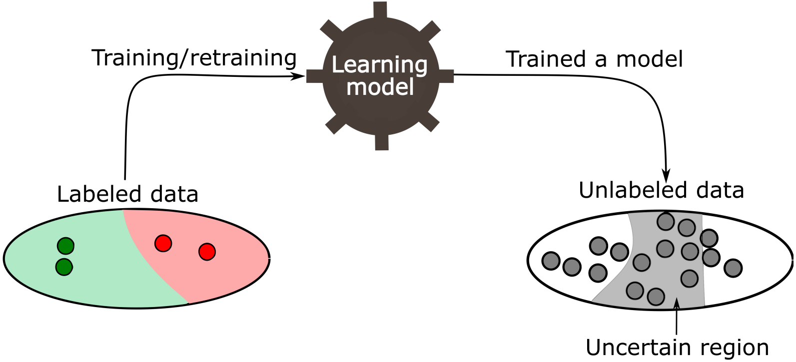

Basic Workflow of Active Learning

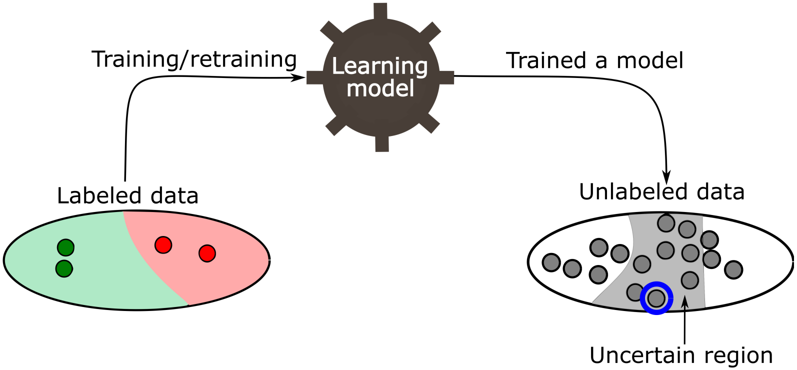

Basic Workflow of Active Learning

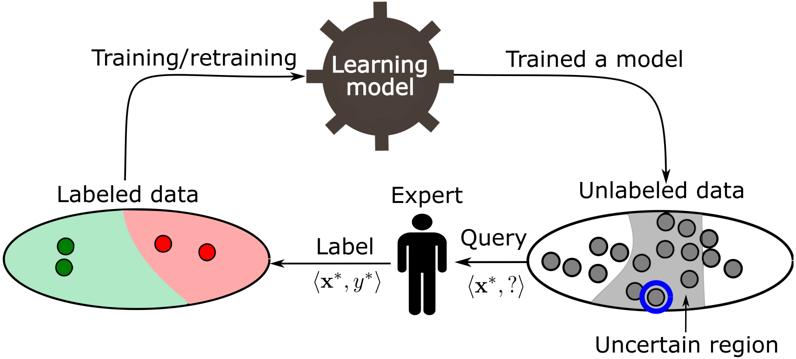

IAL

Basic Workflow of Active Learning

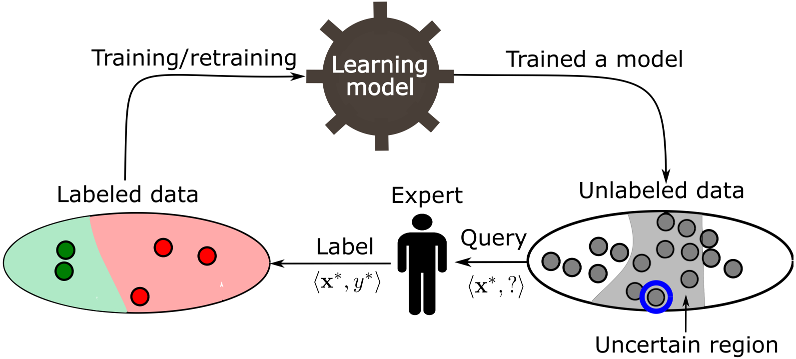

IAL

The main components of Active Learning

- **Data**: (i) unlabeled data (D_U), which represents the pool from which a new point is selected and (ii) labeled data (D_L) is used to train a model (h)
- **Learning algorithm** (h): The learning model (h)is trained on D_L. this component is mostly used to evaluate the current annotation process and find the most uncertain instances/regions
- **Query strategy** (or **acquisition function**): This uses a specific utility function for evaluating the instances in D_U for selecting and querying the most informative and representative point(s) in D_U
- **Annotator/labeler/oracle/Expert**: Who annotates/labels the queried unlabeled points

IAL

Different types of active learning

- **Active labelling**: A model actively selects and requests labels for specific data points from a human annotator in order to improve its performance
  - To build a spam email classifier to automatically identify spam emails without having labelled dataset of emails (as spam or not), instead of labeling the entire dataset manually, active labeling to make the process more efficient

Different types of active learning

- **Active labelling**: A model actively selects and requests labels for specific data points from a human annotator in order to improve its performance

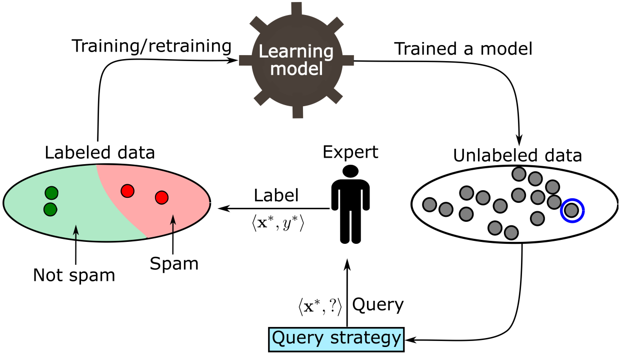

IAL

Different types of active learning

- **Active labelling**: A model actively selects and requests labels for specific data points from a human annotator in order to improve its performance
- **Active feature acquisition**: Here, the model actively selects and acquires additional features (input variables) to improve its performance
  - To build a face recognition model from images, we could extract a lot of features. Let we build a model based on only extract some features from eyes. After training the model and analyzing the feature importance scores, we find that adding (extracting) the ”nose” features has the potential to improve the model

IAL

Different types of active learning

- **Active labelling**: A model actively selects and requests labels for specific data points from a human annotator in order to improve its performance
- **Active feature acquisition**: Here, the model actively selects and acquires additional features (input variables) to improve its performance
- **Active class selection**: Instead of requesting labels for existing instances, or explicitly querying the feature space by creating instances to be labeled by an annotator, ACS create/generate instances for a particular class
  - To train a model in smart factories to classify two classes (negative and positive), the initial training data may be balanced and let we assume that the negative class is more critical; hence, it is better to actively generate and annotate more negative items to improve the model's performance in identifying the items of this class

IAL

What is the benefits of AL?

- reduces the need for large labeled datasets by selecting the most informative data points for labeling ⇒ "**Cost+time saving**"
- can be particularly beneficial when dealing with limited resources, as it allows for the targeted collection of valuable data ⇒ "**Scalability**"
- lead to faster model convergence by actively selecting informative data points, allowing the model to learn more quickly ⇒ "**Faster Model Convergence**"
- results in models with better performance, as they are trained on the most valuable and informative data points ⇒ "**Improved Model Performance**"
- reduce annotation bias by actively seeking diverse examples, leading to a more balanced and representative dataset ⇒ "**Reduced Annotation Bias**"

Simple AL example

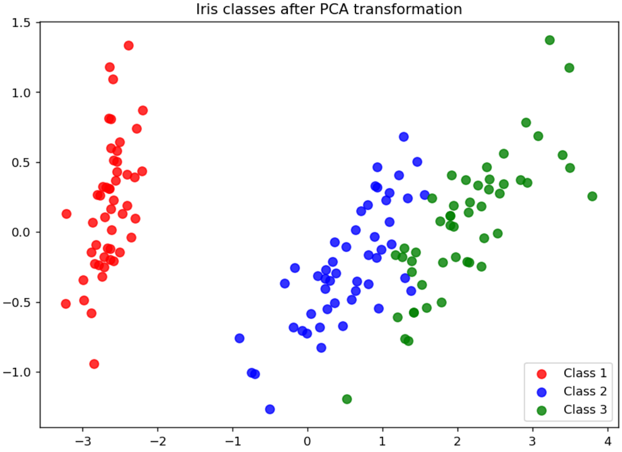

IAL

Simple AL example

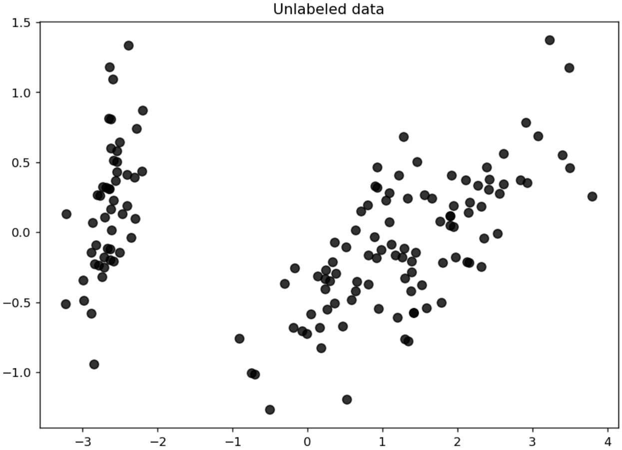

IAL

Simple AL example

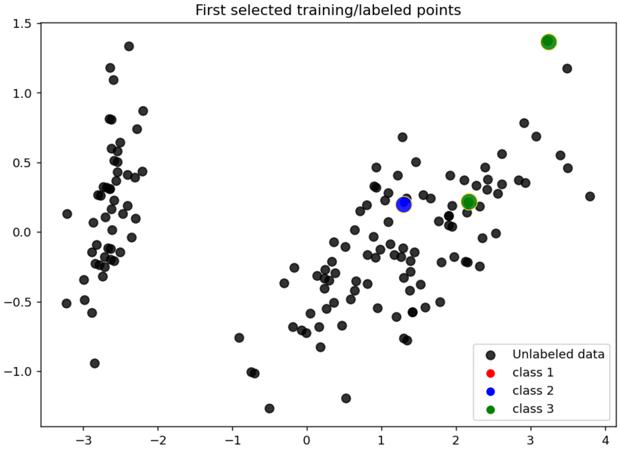

IAL

Simple AL example

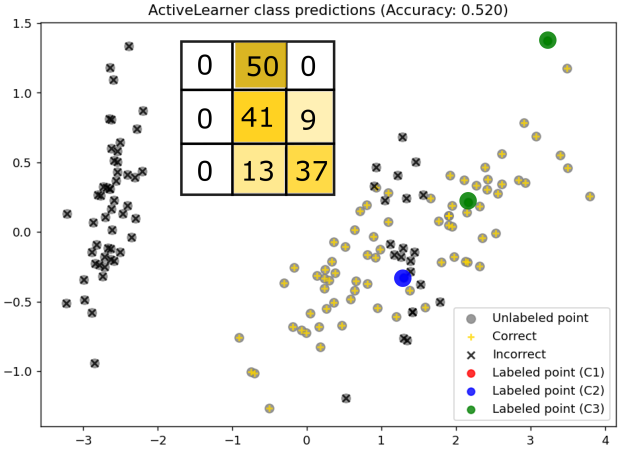

IAL

Simple AL example

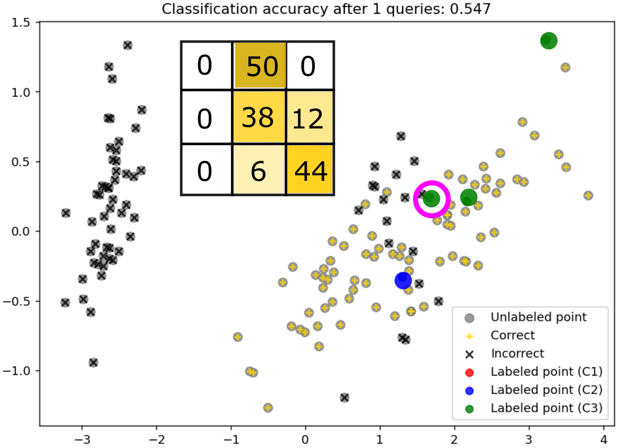

IAL

Simple AL example

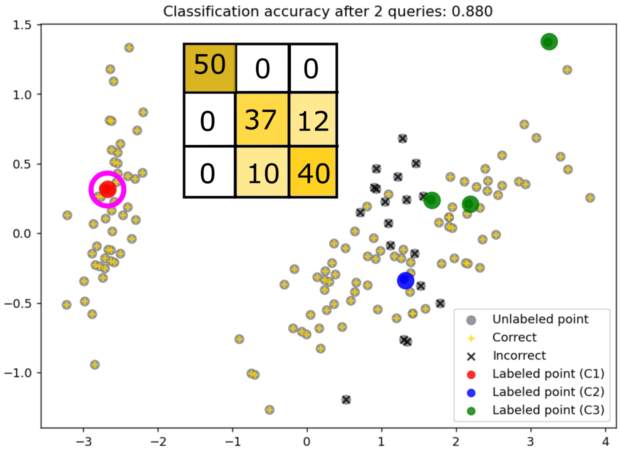

IAL

Simple AL example

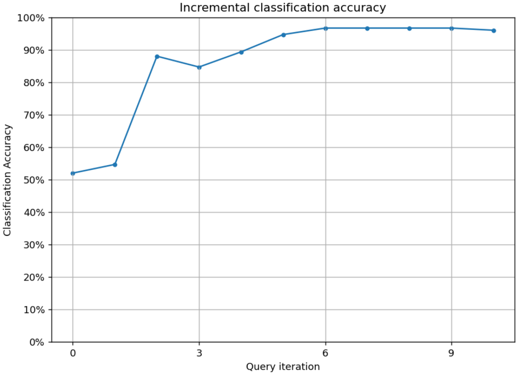

Tharwat, A., & Schenck, W. (2023). A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions. Mathematics, 11(4), 820.

The code is available here: `https://github.com/Eng-Alaa/AL_SurveyPaper/blob/main/AL_IrisData_SurveyPaper.ipynb`

IAL

Beyond pool-based scenarios

Georg Krempl

Beyond pool-based scenarios

IAL

Beyond pool-based scenarios

Aims

- **Broadening view** on active learning
- **Overview** on different variants of the active learning task
- **Pointers** to surveys / key papers for each variant
- **Challenges/caveats** and exemplary approaches

IAL

Active Learning: Broadening the Scope

IAL

Active Learning: Broadening the Scope

active learning

IAL

Active Learning: Broadening the Scope

pool-based

active learning

IAL

Active Learning: Broadening the Scope

processing scenarios

pool-based ✖ active learning

stream-based

query synthesis

IAL

Active Learning: Broadening the Scope

processing scenarios

pool-based ✖ inductive active learning

stream-based

query synthesis

IAL

Active Learning: Broadening the Scope

processing scenarios		learning objective		
pool-based	✖	inductive	✖	active learning
stream-based		transductive		
query synthesis				

IAL

Active Learning: Broadening the Scope

processing scenarios		learning objective		initiation of interaction
pool-based	✖	inductive	✖	active learning
stream-based		transductive		machine teaching
query synthesis				

IAL

Active Learning: Broadening the Scope

processing scenarios	learning objective	initiation of interaction	selected information
pool-based	inductive	active learning	of labels (labelling)
stream-based	transductive	machine teaching	of features (feature acquistion)
query synthesis			of instances (class selection)
			of classification features / instances

IAL

Processing Scenarios

processing scenarios		learning objective		initiation of interaction		selected information
pool-based	✖	inductive	✖	active learning	✖	of labels (labelling)
stream-based		transductive		machine teaching		of features (feature acquistion)
query synthesis						of instances (class selection)
						of classification features / instances

IAL

Processing Scenarios

processing scenarios		learning objective		initiation of interaction		selected information
pool-based	✖	inductive	✖	active learning	✖	of labels (labelling)
stream-based		transductive		machine teaching		of features (feature acquistion)
query synthesis						of instances (class selection)
						of classification features / instances

IAL

Processing Scenarios: Passive Learning

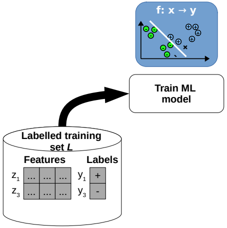

Passive Learning

- **Training set** $\mathcal{L}$ of labelled data available
- **no control** over labelling (no additional labels)

IAL

Processing Scenarios: Pool

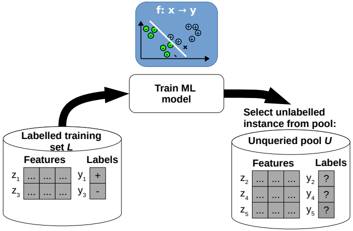

Pool-Based Scenario

- **Pool** $\mathcal{U}$ of unlabelled data
- **Static**, **repeated access**

Processing Scenarios: Pool

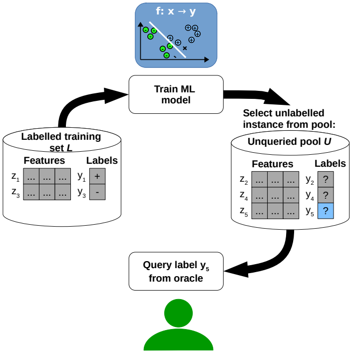

Pool-Based Scenario

- **Pool** $\mathcal{U}$ of unlabelled data
- **Static**, **repeated access**
- **Control** over labelling process

IAL

Processing Scenarios: Pool

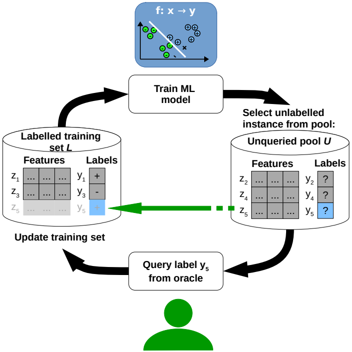

Pool-Based Scenario

- **Pool** $\mathcal{U}$ of unlabelled data
- **Static**, **repeated access**
- **Control** over labelling process
- **Oracle** provides labels
- Labelled instances pool $\rightarrow$ training set

Processing Scenarios: Query Synthesis

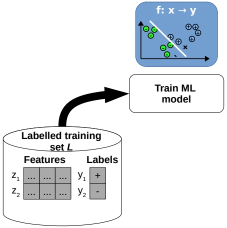

Query Synthesis Scenario

- No **pool**

Processing Scenarios: Query Synthesis

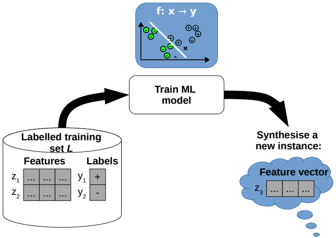

Query Synthesis Scenario

- No **pool**
- **Ad hoc generation** of queried instances

Processing Scenarios: Query Synthesis

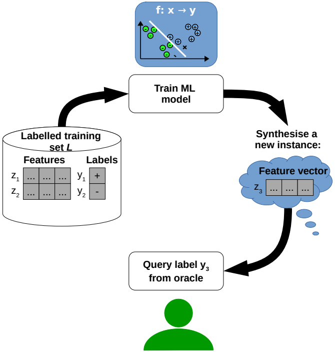

Query Synthesis Scenario

- No **pool**
- **Ad hoc generation** of queried instances
- **Membership query**: Query class membership of generated instance

Processing Scenarios: Query Synthesis

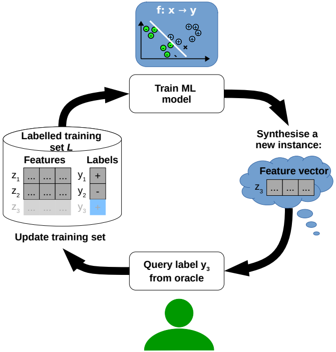

Query Synthesis Scenario

- No **pool**
- **Ad hoc generation** of queried instances
- **Membership query**: Query class membership of generated instance
- See Angluin, “Queries revisited”, 2004 (introduction)

Processing Scenarios: Query Synthesis

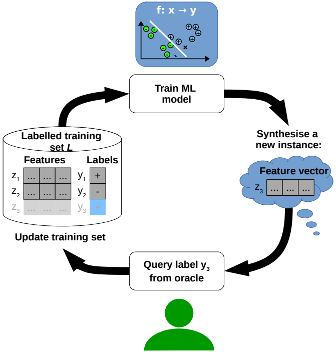

Query Synthesis Scenario

- No **pool**
- **Ad hoc generation** of queried instances
- **Membership query**: Query class membership of generated instance
- See Angluin, "Queries revisited", 2004 (introduction)
- **Challenge: creating meaningfull instances**

Processing Scenarios: Query Synthesis

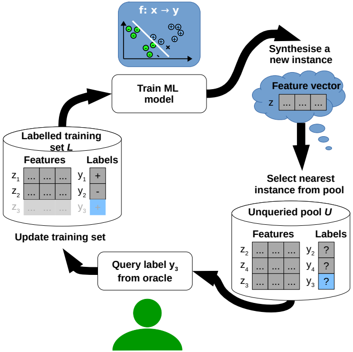

Hybrid Query Synthesis/Pool Scenario

- Aim: creating meaningfull instances

IAL

Processing Scenarios: Query Synthesis

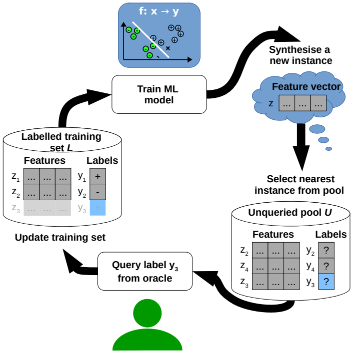

Hybrid Query Synthesis/Pool Scenario

- **Aim: creating meaningfull instances**
- **Combination with pool-based AL**: Wang et al., “Active learning via query synthesis and nearest neighbour search”, 2015
  - — given a (too) large pool of unlabelled data
  - — synthesize instance close to decision boundary
  - — select the nearest neighbouring real instance
  - — faster than pool-based AL, meaningful queries

Processing Scenarios: Stream

Stream-Based Selective Sampling Scenario

- **Sequential arrival, no repeated access**
- **Online** active learning as synonym

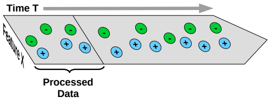

IAL

Processing Scenarios: Stream

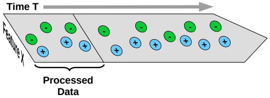

Stream-Based Selective Sampling Scenario

- **Sequential arrival, no repeated access**
- **Online** active learning as synonym
- **No/few initial labels**
- **Possibly infinite** number of instances
- **Efficient processing** and limited storage

IAL

Processing Scenarios: Stream

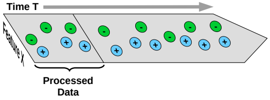

Stream-Based Selective Sampling Scenario

- **Sequential arrival, no repeated access**
- **Online** active learning as synonym

- **No/few initial labels**
- **Possibly infinite** number of instances
- **Efficient processing** and limited storage

- **Non-stationary** distributions (concept drift)
- **Adaptation** (forgetting) needed

IAL

Processing Scenarios: Stream

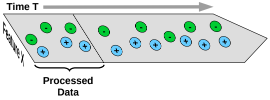

Stream-Based Selective Sampling Scenario

- **Sequential arrival, no repeated access**
- **Online** active learning as synonym
- **No/few initial labels**
- **Possibly infinite** number of instances
- **Efficient processing** and limited storage
- **Non-stationary** distributions (concept drift)
- **Adaptation** (forgetting) needed
- "Big Data" is often streaming data

Processing Scenarios: Stream

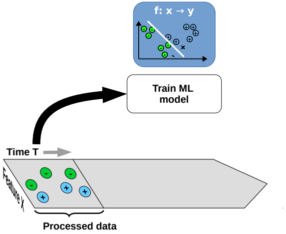

Stream-Based Selective Sampling Scenario

Processing Scenarios: Stream

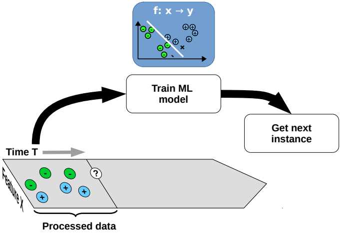

Stream-Based Selective Sampling Scenario

IAL

Processing Scenarios: Stream

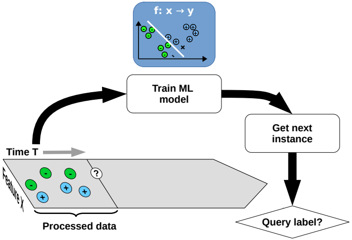

Stream-Based Selective Sampling Scenario

- **Decide upon arrival** of new instance whether to query that instance's label or not

IAL

Processing Scenarios: Stream

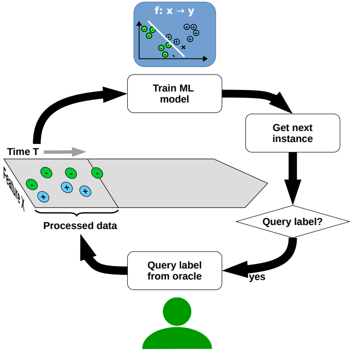

Stream-Based Selective Sampling Scenario

- **Decide upon arrival** of new instance whether to query that instance's label or not
- **Update classifier** if label was queried, otherwise skip

IAL

Processing Scenarios: Stream

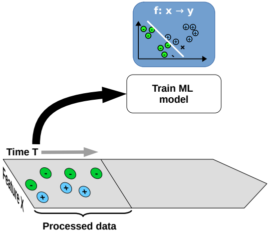

Stream-Based Selective Sampling Scenario

- **Decide upon arrival** of new instance whether to query that instance's label or not
- **Update classifier** if label was queried, otherwise skip
- **Continue** for as long as new instances arrive

IAL

Processing Scenarios: Stream

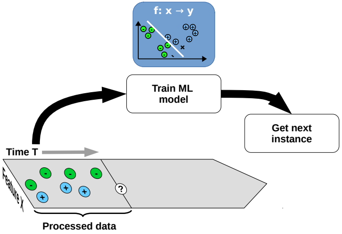

Stream-Based Selective Sampling Scenario

- **Decide upon arrival** of new instance whether to query that instance's label or not
- **Update classifier** if label was queried, otherwise skip
- **Continue** for as long as new instances arrive

Processing Scenarios: Stream

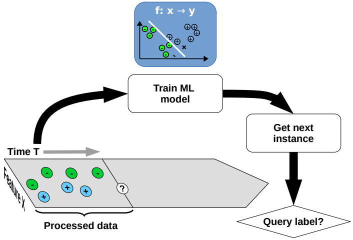

Stream-Based Selective Sampling Scenario

- **Decide upon arrival** of new instance whether to query that instance's label or not
- **Update classifier** if label was queried, otherwise skip
- **Continue** for as long as new instances arrive

Processing Scenarios: Stream

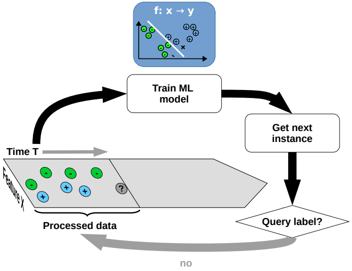

Stream-Based Selective Sampling Scenario

- **Decide upon arrival** of new instance whether to query that instance's label or not
- **Update classifier** if label was queried, otherwise skip
- **Continue** for as long as new instances arrive

Processing Scenarios: Stream

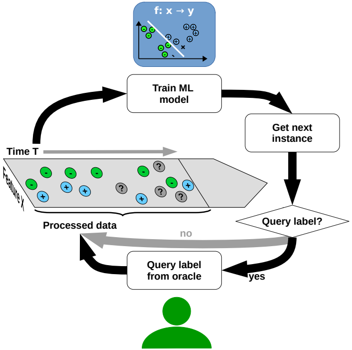

Stream-Based Selective Sampling Scenario

- **Decide upon arrival** of new instance whether to query that instance's label or not
- **Update classifier** if label was queried, otherwise skip
- **Continue** for as long as new instances arrive

IAL

Processing Scenarios: Stream

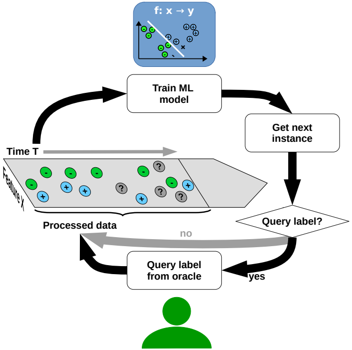

Recommended literature

- Cacciarelli and Kulahci, “A survey on online active learning”, 2023 (survey)
- Zliobaitė et al., “Active Learning With Drifting Streaming Data”, 2013 (concept drift)
- Kottke, Krempl, and Spiliopoulou, “Probabilistic Active Learning in Data Streams”, 2015 (budget management)
- Pham et al., “Stream-Based Active Learning for Sliding Windows Under Verification Latency”, 2022 (verification latency)

Processing Scenarios: Stream

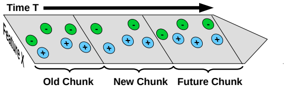

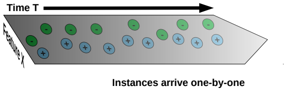

Chunk-based processing

versus

Instance-wise processing

IAL

Processing Scenarios: Stream

Chunk-based processing

- **Split data chronologically into chunks**
- AL on each chunk is similar to pool-based AL
- Often, ensemble with one new classfier per chunk is trained [a]
- Alernative: Clustering-based approaches [b]

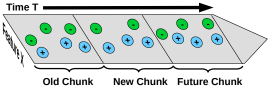

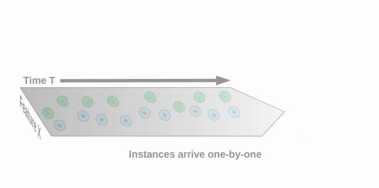

[a] E.g., Ryu et al., “An Efficient Method of Building an Ensemble of Classifiers in Streaming Data”, 2012; Zhu et al., “Active Learning From Stream Data Using Optimal Weight Classifier Ensemble”, 2010; Zhu et al., “Active Learning from Data Streams”, 2007

[b] E.g., Krempl, Ha, and Spiliopoulou, “Clustering-Based Optimised Probabilistic Active Learning (COPAL)”, 2015; Ienco et al., “Clustering Based Active Learning for Evolving Data Streams”,

Processing Scenarios: Stream

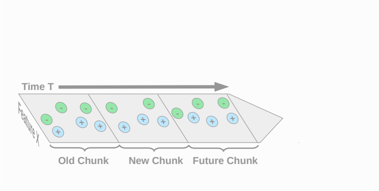

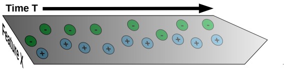

Instances arrive one-by-one

Instance-wise processing

- **Instances arrive one-by-one**
- Decision to query or not must be taken at once
- **Budget:** Trade-off between spatial and temporal usefulness [a]

[a] See Kottke, Krempl, and Spiliopoulou, “Probabilistic Active Learning in Data Streams”, 2015

IAL

Scenarios: Stream: Concept Drift

IAL

Scenarios: Stream: Concept Drift

Categorizing Drift

See e.g., Krempl et al., “Open Challenges for Data Stream Mining Research”, 2014; Žliobaitė, Pechenizkiy, and Gama, “An Overview of Concept Drift Applications”, 2016; Webb et al., “Understanding Concept Drift”, 2017.

IAL

Scenarios: Stream: Concept Drift

Categorizing Drift

- **Affected distribution**: $P(X, Y)$, $P(X)$, $P(Y)$, $P(Y|X)$, $P(X|Y)$

See e.g., Krempl et al., "Open Challenges for Data Stream Mining Research", 2014; Žliobaitė, Pechenizkiy, and Gama, "An Overview of Concept Drift Applications", 2016; Webb et al., "Understanding Concept Drift", 2017.

IAL

Scenarios: Stream: Concept Drift

Categorizing Drift

- **Affected distribution**: $P(X, Y)$, $P(X)$, $P(Y)$, $P(Y|X)$, $P(X|Y)$
- **Smoothnes** of concept transition: sudden shift vs. gradual drift

See e.g., Krempl et al., “Open Challenges for Data Stream Mining Research”, 2014; Žliobaitė, Pechenizkiy, and Gama, “An Overview of Concept Drift Applications”, 2016; Webb et al., “Understanding Concept Drift”, 2017.

IAL

Scenarios: Stream: Concept Drift

Categorizing Drift

- **Affected distribution**: $P(X, Y)$, $P(X)$, $P(Y)$, $P(Y|X)$, $P(X|Y)$
- **Smoothnes** of concept transition: sudden shift vs. gradual drift
- **Recurring** or singular contexts

See e.g., Krempl et al., “Open Challenges for Data Stream Mining Research”, 2014; Žliobaitė, Pechenizkiy, and Gama, “An Overview of Concept Drift Applications”, 2016; Webb et al., “Understanding Concept Drift”, 2017.

IAL

Scenarios: Stream: Concept Drift

Categorizing Drift

- **Affected distribution**: $P(X, Y)$, $P(X)$, $P(Y)$, $P(Y|X)$, $P(X|Y)$
- **Smoothnes** of concept transition: sudden shift vs. gradual drift
- **Recurring** or singular contexts
- **Systematic** (change patterns) or unsystematic

See e.g., Krempl et al., “Open Challenges for Data Stream Mining Research”, 2014; Žliobaitė, Pechenizkiy, and Gama, “An Overview of Concept Drift Applications”, 2016; Webb et al., “Understanding Concept Drift”, 2017.

IAL

Scenarios: Stream: Concept Drift

Categorizing Drift

- **Affected distribution**: $P(X, Y)$, $P(X)$, $P(Y)$, $P(Y|X)$, $P(X|Y)$
- **Smoothnes** of concept transition: sudden shift vs. gradual drift
- **Recurring** or singular contexts
- **Systematic** (change patterns) or unsystematic
- **Real** or virtual

See e.g., Krempl et al., “Open Challenges for Data Stream Mining Research”, 2014; Žliobaitė, Pechenizkiy, and Gama, “An Overview of Concept Drift Applications”, 2016; Webb et al., “Understanding Concept Drift”, 2017.

IAL

Scenarios: Stream: Self Lock-In Problem

Motivation

Simply using static (*iid*) strategies fails

IAL

Scenarios: Stream: Self Lock-In Problem

Motivation

Simply using static (*iid*) strategies fails

- Example: Uncertainty sampling

Scenarios: Stream: Self Lock-In Problem

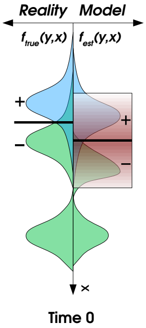

Motivation

Simply using static (*iid*) strategies fails

- Example: Uncertainty sampling

IAL

Scenarios: Stream: Self Lock-In Problem

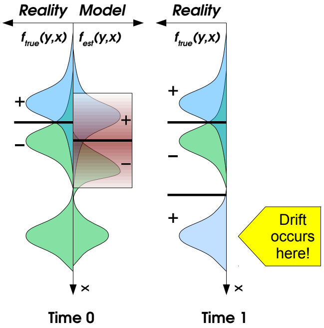

Motivation

Simply using static (*iid*) strategies fails

- Example: Uncertainty sampling

IAL

Scenarios: Stream: Self Lock-In Problem

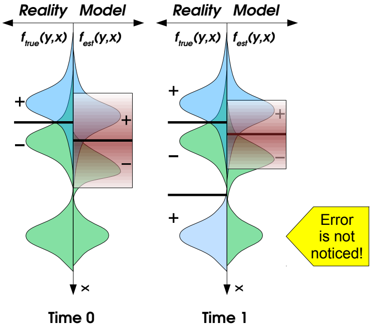

Motivation

Simply using static (*iid*) strategies fails

- Example: Uncertainty sampling

Scenarios: Stream: Self Lock-In Problem

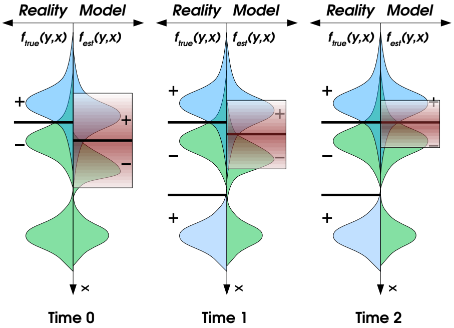

Motivation

Simply using static (*iid*) strategies fails

- Example: Uncertainty sampling

IAL

Scenarios: Stream: Self Lock-In Problem

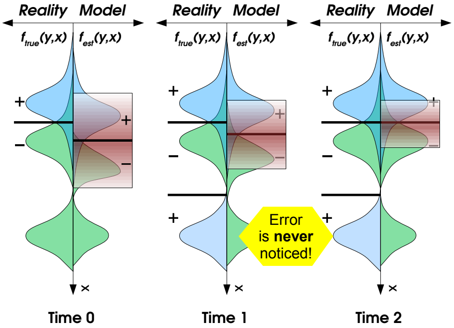

Motivation

Simply using static (*iid*) strategies fails

- Example: Uncertainty sampling

Scenarios: Stream: Self Lock-In Problem

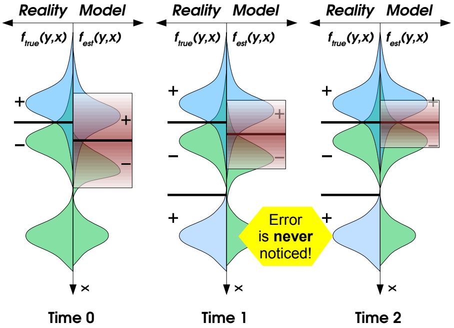

Motivation

Simply using static (*iid*) strategies fails

- Example: Uncertainty sampling
- Error is *never* even noticed!

IAL

Scenarios: Stream: Self Lock-In Problem

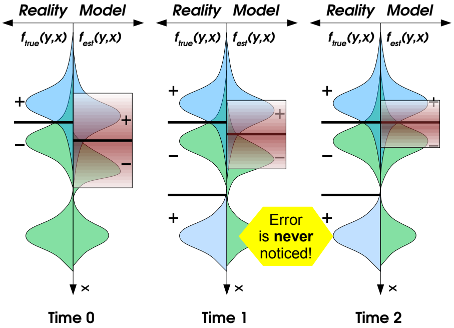

Motivation

Simply using static (*iid*) strategies fails

- Example: Uncertainty sampling
- Error is *never* even noticed!
- **Active learner (self) lock-in** on an outdated hypothesis

Scenarios: Stream: Self Lock-In Problem

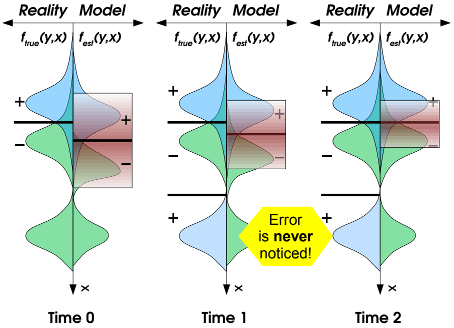

Motivation

Simply using static (*iid*) strategies fails

- Example: Uncertainty sampling
- Error is *never* even noticed!
- **Active learner (self) lock-in** on an outdated hypothesis
- **Anywhere, anytime** drift can occur
  Zliobaitė et al., "Active Learning with Evolving Streaming Data", 2011

IAL

Scenarios: Stream: Challenges

Pool Active Learning

- Where to buy instances (spatial usefulness)?
  - Balance Exploration and Exploitation in the dataspace

IAL

Scenarios: Stream: Challenges

Pool Active Learning

- Where to buy instances (spatial usefulness)?
  - Balance Exploration and Exploitation in the dataspace

Stream Active Learning

- Where to buy labels (**spatial usefulness**)?

IAL

Scenarios: Stream: Challenges

Pool Active Learning

- Where to buy instances (spatial usefulness)?
  - Balance Exploration and Exploitation in the dataspace

Stream Active Learning

- Where to buy labels (**spatial usefulness**)?
- **Consider Drift**
  - Labels might change over time and have to be validated
  - Lifetime of labels

IAL

Scenarios: Stream: Challenges

Pool Active Learning

- Where to buy instances (spatial usefulness)?
  - Balance Exploration and Exploitation in the dataspace

Stream Active Learning

- Where to buy labels (**spatial usefulness**)?
- **Consider Drift**
  - Labels might change over time and have to be validated
  - Lifetime of labels
- When to buy labels (**temporal usefulness**)?
  - Balance Exploration and Exploitation in time

IAL

Scenarios: Stream: Spatial Usefulness

Where to buy labels?

- Use scores from pool-based methods like
  - Uncertainty sampling
  - Query by committee
  - Probabilistic active learning

IAL

Scenarios: Stream: Spatial Usefulness

Where to buy labels?

- Use scores from pool-based methods like
  — Uncertainty sampling
  — Query by committee
  — Probabilistic active learning

Approach

Find best instances spatially (based on feature vectors) balancing:

- exploration (observe unsampled regions)
- exploitation (acquire labels in regions near decision boundaries to elaborate the decision)

IAL

Scenarios: Stream: Budget in Streams

- Pools: absolute number (e.g. stop after 40 labels)
- Streams: relative definition necessary (e.g. buy 10%)
- How to distribute the budget over time?

IAL

Scenarios: Stream: Budget in Streams

- Pools: absolute number (e.g. stop after 40 labels)
- Streams: relative definition necessary (e.g. buy 10%)
- How to distribute the budget over time?
  - constantly (every 10th label $\rightarrow$ no spatial selection necessary)

IAL

Scenarios: Stream: Budget in Streams

- Pools: absolute number (e.g. stop after 40 labels)
- Streams: relative definition necessary (e.g. buy 10%)
- How to distribute the budget over time?
  - constantly (every 10th label $\rightarrow$ no spatial selection necessary)
  - almost constantly (with a small tolerance window) Kottke, Krempl, and Spiliopoulou, “Probabilistic Active Learning in Data Streams”, 2015

IAL

Scenarios: Stream: Budget in Streams

- Pools: absolute number (e.g. stop after 40 labels)
- Streams: relative definition necessary (e.g. buy 10%)
- How to distribute the budget over time?
  - constantly (every 10th label $\rightarrow$ no spatial selection necessary)
  - almost constantly (with a small tolerance window) Kottke, Krempl, and Spiliopoulou, “Probabilistic Active Learning in Data Streams”, 2015
  - bounded (budget should not exceed 10%) Zliobaitė et al., “Active Learning With Drifting Streaming Data”, 2013

IAL

Scenarios: Stream: Budget in Streams

- Pools: absolute number (e.g. stop after 40 labels)
- Streams: relative definition necessary (e.g. buy 10%)
- How to distribute the budget over time?
  - constantly (every 10th label $\rightarrow$ no spatial selection necessary)
  - almost constantly (with a small tolerance window) Kottke, Krempl, and Spiliopoulou, "Probabilistic Active Learning in Data Streams", 2015
  - bounded (budget should not exceed 10%) Žliobaitė et al., "Active Learning With Drifting Streaming Data", 2013
  - dynamic (budget changes over time)

Scenarios: Stream: Temporal usefulness (When to buy?)

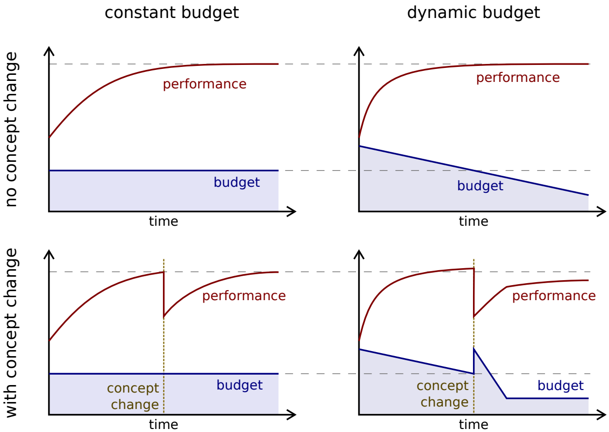

IAL

Scenarios: Stream: Temporal usefulness (When to buy?)

- Labels in the beginning are more beneficial as they affect more future decisions (resp. after changes)

Scenarios: Stream: Temporal usefulness (When to buy?)

- Labels in the beginning are more beneficial as they affect more future decisions (resp. after changes)
- But: one does not know when change take place

IAL

Scenarios: Stream: Temporal usefulness (When to buy?)

- Labels in the beginning are more beneficial as they affect more future decisions (resp. after changes)
- But: one does not know when change take place
- Standard technique: constant budget

IAL

Scenarios: Stream: Temporal usefulness (When to buy?)

- Labels in the beginning are more beneficial as they affect more future decisions (resp. after changes)
- But: one does not know when change take place
- Standard technique: constant budget

Exploration vs. Exploitation

- Exploration: Sample randomly to be able to detect change

IAL

Scenarios: Stream: Temporal usefulness (When to buy?)

- Labels in the beginning are more beneficial as they affect more future decisions (resp. after changes)
- But: one does not know when change take place
- Standard technique: constant budget

Exploration vs. Exploitation

- Exploration: Sample randomly to be able to detect change
- Exploitation: Sample the most promising labels

Scenarios: Stream: Temporal usefulness (When to buy?)

- Labels in the beginning are more beneficial as they affect more future decisions (resp. after changes)
- But: one does not know when change take place
- Standard technique: constant budget

Exploration vs. Exploitation

- Exploration: Sample randomly to be able to detect change
- Exploitation: Sample the most promising labels
- How to cope with gradual drifts?

IAL

Scenarios: Stream: Temporal usefulness (When to buy?)

- Labels in the beginning are more beneficial as they affect more future decisions (resp. after changes)
- But: one does not know when change take place
- Standard technique: constant budget

Exploration vs. Exploitation

- Exploration: Sample randomly to be able to detect change
- Exploitation: Sample the most promising labels
- How to cope with gradual drifts?
- High budgets after change might cause problems due to less spatial usefulness

Processing Scenarios: Stream with Latency

Verification Latency

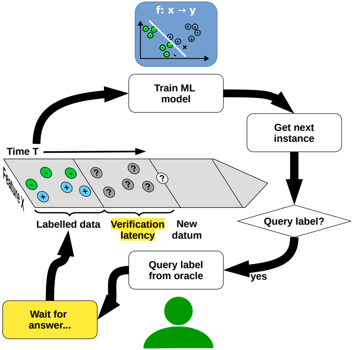

- **Delay** between query and answer
- **Achronologic**: new unlabelled instances might arrive before previously queried labels

IAL

Processing Scenarios: Stream with Latency

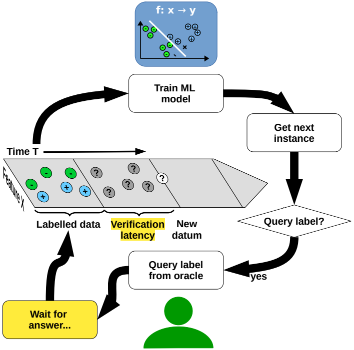

Verification Latency

- **Delay** between query and answer
- **Achronologic**: new unlabelled instances might arrive before previously queried labels
- **Redundancy in queries**, if new instance is similar to already queried (but not yet obtained) one
- **Knowledge gaps**, if no new instance was queried in time, before old labelled one got forgotten

Processing Scenarios: Stream with Latency

Verification Latency

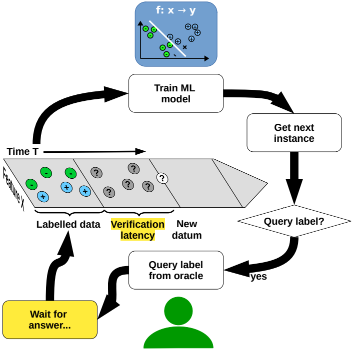

- **Delay** between query and answer
- **Achronologic**: new unlabelled instances might arrive before previously queried labels
- **Redundancy in queries**, if new instance is similar to already queried (but not yet obtained) one
- **Knowledge gaps**, if no new instance was queried in time, before old labelled one got forgotten
- Pham et al., "Stream-Based Active Learning for Sliding Windows Under Verification Latency", 2022 (first paper on verification latency and AL)

Processing Scenarios: Stream with Latency

Naive Approach

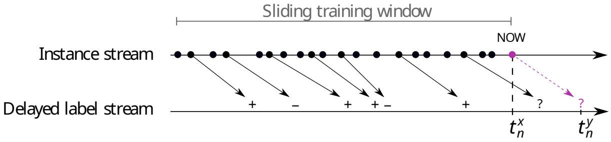

Figure: Naive (Latency-Ignorant) Approach

IAL

Processing Scenarios: Stream with Latency

Latency-Aware Approach

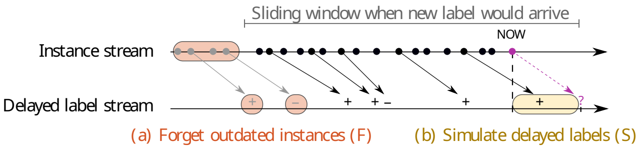

Figure: Verification Latency-Aware Approach suggested in Pham et al., "Stream-Based Active Learning for Sliding Windows Under Verification Latency", 2022

Active Learning: Learning Objective

processing scenarios		learning objective		initiation of interaction		selected information
pool-based	✖	**inductive**	✖	active learning	✖	of labels (labelling)
stream-based		**transductive**		machine teaching		of features (feature acquistion)
query synthesis						of instances (class selection)
						of classification features / instances

Active Learning: Learning Objective

processing scenarios		learning objective		initiation of interaction		selected information
pool-based	✖	inductive	✖	active learning	✖	of labels (labelling)
stream-based		transductive		machine teaching		of features (feature acquistion)
query synthesis						of instances (class selection)
						of classification features / instances

Learning Objective: Inductive vs. Transductive

skip

IAL

Learning Objective: Inductive vs. Transductive

skip

Inductive

- Training and test data are different
- Objective: Generalising to unseen data

IAL

Learning Objective: Inductive vs. Transductive

skip

Inductive

- Training and test data are different
- Objective: Generalising to unseen data

Transductive

- Same data used for training needs to be classified
- Objective: Mastering given (training) data set

IAL

Learning Objective: Inductive vs. Transductive

skip

Particularities of Transductive AL

- **Evaluation data is known beforehand**, as test and train set are identical, no need to build a generalised model
- **Excluding** instances from being predicted by the classifier is possible by querying them from the oracle

Implications

- Ignore high aleatoric uncertainty for inductive setting
- Remove such instances by labelling for transductive setting
- See Kottke et al., “A Stopping Criterion for Transductive Active Learning”, 2022

IAL

Learning Objective: Inductive vs. Transductive

Transductive Gain

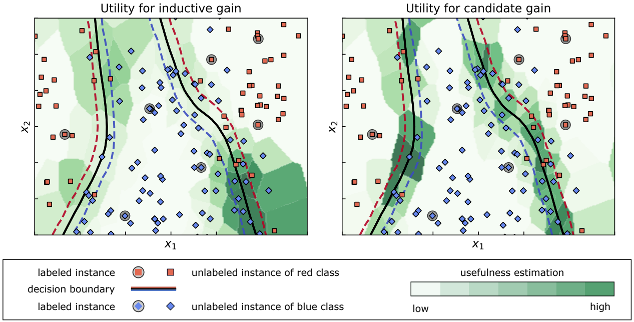

Figure: Transductive gain as sum of the utilities of inductive gain (left), and of candidate gain (right) Kottke et al., “A Stopping Criterion for Transductive Active Learning”, 2022, Fig 1

IAL

Active Learning: Initiatior of Interaction

processing scenarios		learning objective		initiation of interaction		selected information
pool-based	✖	inductive	✖	active learning	✖	of labels (labelling)
stream-based		transductive		machine teaching		of features (feature acquistion)
query synthesis						of instances (class selection)
						of classification features / instances

IAL

Active Learning: Initiatior of Interaction

processing scenarios		learning objective		initiation of interaction		selected information
pool-based	✖	inductive	✖	active learning	✖	of labels (labelling)
stream-based		transductive		**machine teaching**		of features (feature acquistion)
query synthesis						of instances (class selection)
						of classification features / instances

IAL

Initiatior of Interaction: Machine (Active Learning)

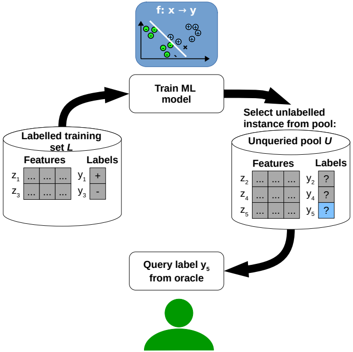

Active Learning

- **Machine** is proactive in the interaction

IAL

Initiatior of Interaction: Human (Machine Teaching)

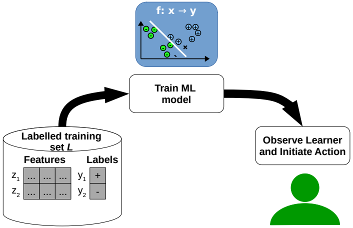

Machine Teaching

- **Human** is proactive in the interaction
- **No direct knowledge transfer** between teacher (human) and learner (machine)

Initiatior of Interaction: Human (Machine Teaching)

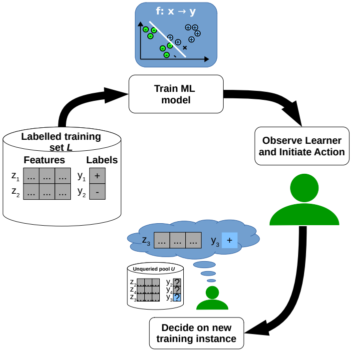

Machine Teaching

- **Human** is proactive in the interaction
- **No direct knowledge transfer** between teacher (human) and learner (machine)
- **Aim is designing an optimal training set**

IAL

Initiatior of Interaction: Human (Machine Teaching)

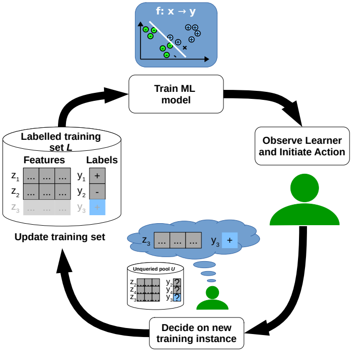

Machine Teaching

- **Human** is proactive in the interaction
- **No direct knowledge transfer** between teacher (human) and learner (machine)
- ***Aim is designing an optimal training set***
- See Tegen, "Interactive Online Machine Learning", 2022 (PhD thesis) and Tegen, Davidsson, and Persson, "A Taxonomy of Interactive Online Machine Learning Strategies", 2021 (review)

IAL

Initiatior of Interaction: Human (Machine Teaching)

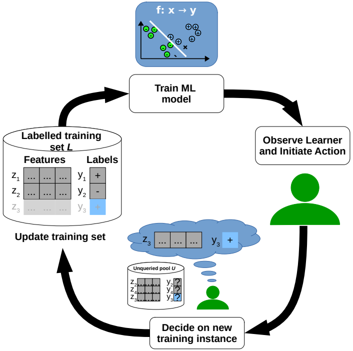

Triggers for human to add instances to training set might be

- Trigger by **error**
- Trigger by **state change**
- Trigger by **time**
- Trigger by **user factors**

IAL

Active Learning: Selected Information

processing scenarios		learning objective		initiation of interaction		**selected information**
pool-based	✖	inductive	✖	active learning	✖	of labels (labelling)
stream-based		transductive		machine teaching		of features (feature acquistion)
query synthesis						of instances (class selection)
						of classification features / instances

IAL

Active Learning: Selected Information

processing scenarios		learning objective		initiation of interaction		**selected information**
pool-based	✖	**inductive**	✖	**active learning**	✖	**of labels (labelling)**
stream-based		**transductive**		**machine teaching**		**of features (feature acquistion)**
query synthesis						**of instances (class selection)**
						of classification features / instances

IAL

Active Learning: Selected Information

- We will continue with this after the poster session and coffee break
- Questions, comments, suggestions?

Time	Program	Presenter / Author
09:00–11:00	**Session 1: Tutorials & Poster Session**	
09:00–09:30	Tutorial Part I: Foundations of Active Learning	A. Tharwat
09:30–10:30	Tutorial Part II: Beyond Pool-Based Scenarios	G. Krempl
11:30–11:00	Poster Session	
	Coffee Break (11:00–11:30)	
11:30–13:00	**Session 2: Tutorials**	
11:30–12:00	Tutorial Part III: Beyond Active Labelling	M. Bunse
12:00–12:30	Tutorial Part IV: Towards Explainable Active Learning using Meta-Learning	A. Saadallah
12:30–13:00	Tutorial Part V: Practical Challenges and New Research Directions	A. Tharwat
	Lunch Break (13:00–14:00)	

Lunch Break (13:00–14:00)

14:00–16:00	**Session 3: Keynote & Workshop Contributions**	
14:00–14:40	Keynote: From Insights to Impact: A Metrics-Driven Active Learning Journey	A. Abraham
14:40–15:00	Active Learning for Survival Analysis with Incrementally Disclosed Label Information	K. Dedja, F.K. Nakano & C. Vens
15:00–15:15	Towards Enhancing Deep Active Learning with Weak Supervision and Constrained Clustering	M. Aßenmacher, L. Rauch, J. Goschenhofer, A. Stephan, B. Bischl, B. Roth & B. Sick
15:15–15:30	Who knows best? A Case Study on Intelligent Crowdworker Selection via Deep Learning	M. Herde, D. Huseljic, B. Sick, U. Bretschneider & S. Oeste-Reiß
15:30–15:45	Role of Hyperparameters in Deep Active Learning	D. Huseljic, M. Herde, P. Hahn & B. Sick
15:45–16:00	Challenges for Active Feature Acquisition and Imputation on Data Streams	C. Beyer, M. Büttner & M. Spiliopoulou

Coffee Break (16:00–16:30)

Coffee Break (16:00–16:30)

16:30–17:40	**Session 4: Workshop Contributions & Closing**	
16:30–16:50	Active Learning with Fast Model Updates and Class-Balanced Selection for Imbalanced Datasets	Z. Huang, Y. He, M. Herde, D. Huseljic & B. Sick
16:50–17:10	Interpretable Meta-Active Learning for Regression Ensemble Learning	O. Saadallah & Z. Rouissi
17:10–17:30	Look and You Will Find It: Fairness-Aware Data Collection through Active Learning	H. Weerts, R. Theunissen & M. Willemsen
17:30–17:40	Closing	

IAL

Beyond active labeling

Mirko Bunse

IAL

Beyond active labeling

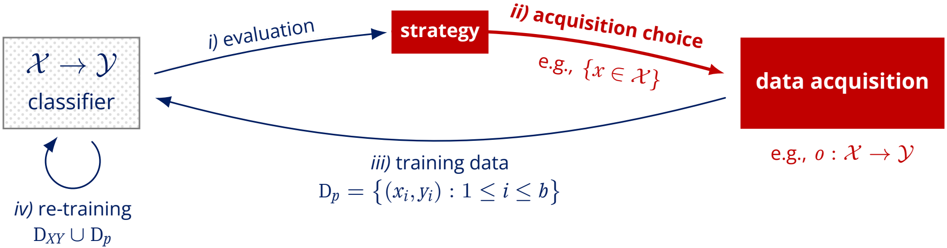

IAL

Beyond active labeling

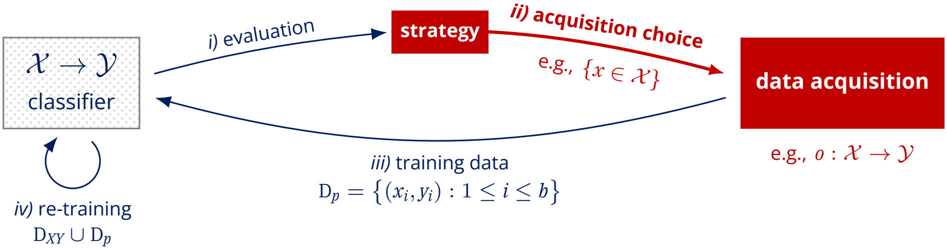

We often assume an oracle $o : \mathcal{X} \to \mathcal{Y}$, **but what if there is none?**

- lack of (human) expertise / lack of data interpretability
- extreme data volumes

Also, **labels aren't the only cost factor.**

IAL

Beyond active labeling

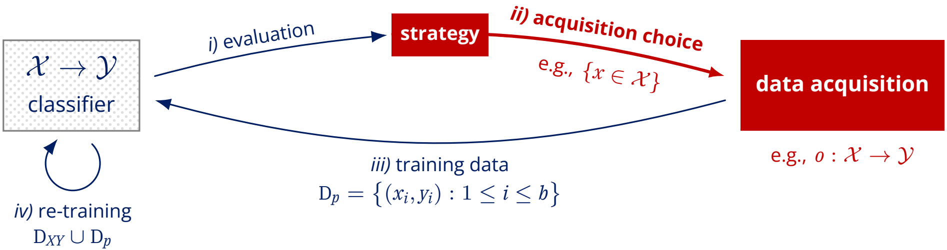

	data acquisition	**acquisition choice**	
active labeling	oracle $o : \mathcal{X} \to \mathcal{Y}$	feature vectors $\{x \in \mathcal{X}\}$	
active class selection	generator $g : \mathcal{Y} \to \mathcal{X}$	class proportions $p \in \mathbb{R}^{\lvert\mathcal{Y}\rvert}$	**this talk**
active feature acquisition	feature oracle $f : \mathcal{I} \times \mathcal{J} \to \mathbb{R}$	sample $\times$ feature indices $\{(i,j) \in \mathcal{I} \times \mathcal{J}\}$	**this talk**

IAL

Active class selection

Applications

ACS applications provide a generator $g : \mathcal{Y} \to \mathcal{X}$ that is costly.

- Particle detectors: accelerate a particle (Y) before it can be recorded (X)
- Gas sensors: inject a gas (Y) before it can be recorded (X)
- Brain-computer interaction: ask for an intent (Y) to record brain signal (X)
- Search engines for labeling: search for a concept (Y) to collect data (X)
- ...

This resulting data is called *"anti-causal"*[1] or *"intrinsically labeled"*[2].

[1] Schölkopf et al., "On causal and anticausal learning", 2012.
[2] Card and Smith, "The importance of calibration for estimating proportions from annotations", 2018.

IAL

Active class selection

Heuristic methods

Idea: acquire classes according to some utility measure $u : \mathcal{Y} \to \mathbb{R}$,

heuristic	**utility** $u(y)$	**intent**
uniform[3]	1	optimize AUROC or balanced accuracy
proportional[3]	$\mathbb{P}(Y = y)$	optimize accuracy if $\mathbb{P}(Y = y)$ is known
inverse[3]	$\mathrm{Accuracy}_h(y)^{-1}$	improve badly predicted classes
improvement[3]	$(\mathrm{Accuracy}_h(y) - \mathrm{LastAccuracy}_h(y))^{-1}$	exploit improvements
redistriction[3]	n_y, the number of changed predictions	stabilize volatile decision boundaries
ACS-PAL[4]	$\frac{1}{m_y^+} \sum_{i=1}^{m_y^+} u_{\mathrm{AL}}(x_i)$	avg. pseudo-instance utility
RF-Impurity[5]	$\frac{1}{m_y} \sum_{i=1}^{m_y} 1 - \mathbb{P}(y \mid x_i)$	avg. confusion

[3] Lomasky et al., “Active class selection”, 2007, .

[4] Kottke et al., “Probabilistic active learning for active class selection”, 2016.

[5] Bicego et al., “Active class selection for dataset acquisition in sign language recognition”, 2023.

Active class selection

PAC bounds

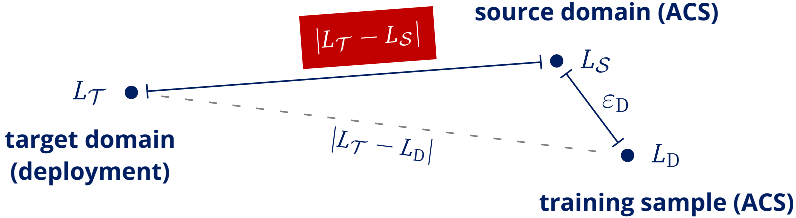

Label shift bound:[6] For any $\varepsilon_\mathrm{D} > 0$ and any fixed $h \in \mathcal{H}$, it holds with probability at least $1 - \delta$, where $\delta = 4e^{-2|\mathrm{D}|\varepsilon_\mathrm{D}^2}$, that

$$|L_\mathcal{T}(h) - L_\mathcal{S}(h)| \; - \; \varepsilon_\mathrm{D} \quad \leq \quad |L_\mathcal{T}(h) - L_\mathrm{D}(h)| \quad \leq \quad |L_\mathcal{T}(h) - L_\mathcal{S}(h)| \; + \; \varepsilon_\mathrm{D}$$

[6] Bunse and Morik, “Certification of model robustness in active class selection”, 2021.

Active class selection

Certification

Certified hypothesis: Let $\varepsilon \in \mathbb{R}$ and let $\delta \geq 0$. A hypothesis $h \in \mathcal{H}$ is (ε, δ)-certified for a set of class proportions $\mathcal{P}$ if, with probability at least $1-\delta$,

$$L_{\mathcal{T}}(h) \ \leq \ L_{\mathcal{S}}(h) \ + \ \varepsilon \quad \forall \, \mathbf{p}_{\mathcal{T}} \in \mathcal{P}$$

IAL

Active class selection

Certification

Certified hypothesis: Let $\varepsilon \in \mathbb{R}$ and let $\delta \geq 0$. A hypothesis $h \in \mathcal{H}$ is (ε, δ)-certified for a set of class proportions $\mathcal{P}$ if, with probability at least $1-\delta$,

$$L_{\mathcal{T}}(h) \;\leq\; L_{\mathcal{S}}(h) \;+\; \varepsilon \quad \forall\, \mathbf{p}_{\mathcal{T}} \in \mathcal{P}$$

Distance certificate: Let $(p, q) \in \{(1, \infty), (2, 2), (\infty, 1)\}$ be two vector norms. $h \in \mathcal{H}$ is (p, ε, δ)-certified for a distance of $d > 0$ if it is certified for $\mathcal{P} = \{\mathbf{p}_{\mathcal{T}} : \|\mathbf{p}_{\mathcal{T}} - \mathbf{p}_{\mathcal{S}}\|_p \leq d\}$.

IAL

Active class selection

Certification

Certified hypothesis: Let $\varepsilon \in \mathbb{R}$ and let $\delta \geq 0$. A hypothesis $h \in \mathcal{H}$ is (ε, δ)-certified for a set of class proportions $\mathcal{P}$ if, with probability at least $1-\delta$,

$$L_{\mathcal{T}}(h) \;\leq\; L_{\mathcal{S}}(h) \;+\; \varepsilon \quad \forall\, \mathbf{p}_{\mathcal{T}} \in \mathcal{P}$$

Distance certificate: Let $(p, q) \in \{(1, \infty), (2, 2), (\infty, 1)\}$ be two vector norms. $h \in \mathcal{H}$ is (p, ε, δ)-certified for a distance of $d > 0$ if it is certified for $\mathcal{P} = \{\mathbf{p}_{\mathcal{T}} : \|\mathbf{p}_{\mathcal{T}} - \mathbf{p}_{\mathcal{S}}\|_p \leq d\}$.

For $d = \frac{\varepsilon}{\|\boldsymbol{\ell}\|_q}$ we have $\delta = 0$, but $\|\boldsymbol{\ell}\|_q \leq \|\hat{\boldsymbol{\ell}} + \boldsymbol{\varepsilon}\|_q$ requires

$$\boldsymbol{\varepsilon}^* \;=\; \arg\min_{\boldsymbol{\varepsilon} > 0} \|\hat{\boldsymbol{\ell}} + \boldsymbol{\varepsilon}\|_q \quad \text{subject to} \quad \sum_{i=1}^{N} \delta_i \;\leq\; \delta$$

IAL

Active class selection

A strategy for uncertain deployment class proportions

Consequence: we need prior assumptions about deployment class proportions.

Strategy: acquire data through gradient descent steps $-\nabla_{\mathbf{m}}\varepsilon^*$, where

$$\varepsilon^*(\mathbf{m}) = \int \underbrace{\widehat{\mathbb{P}}(\mathbf{p}_{\mathcal{T}} = \mathbf{p})}_{\text{prior}} \cdot \underbrace{\|\mathbf{p}_{\mathcal{S}}(\mathbf{m}) - \mathbf{p}\|_p \cdot \|\boldsymbol{\ell}(\mathbf{m})\|_q^*}_{\text{upper loss bound}} \, \mathrm{d}\mathbf{p},$$

IAL

Active class selection

A strategy for uncertain deployment class proportions

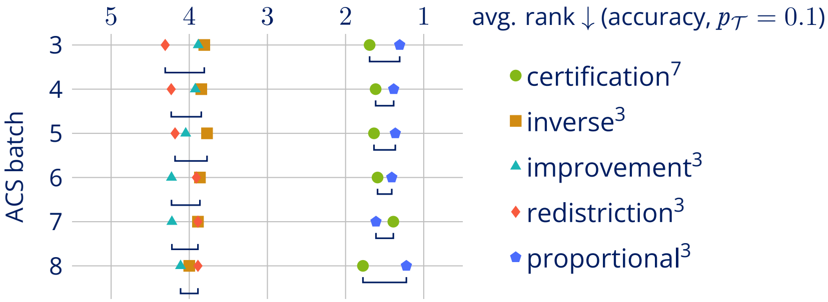

Outlook: non-decomposable loss functions, like F_1 score.

[7] Bunse and Morik, “Active class selection with uncertain deployment class proportions”, 2021.

IAL

Active feature acquisition

	feature 1	feature 2	feature 3	label
instance 1	x_{11}	x_{12}	**?**	y_1
instance 2	**?**	x_{22}	x_{23}	y_2
instance 3	x_{31}	x_{32}	x_{33}	y_3
instance 4	x_{41}	**?**	**?**	y_4

Goal: select feature values x_{ij} to acquire

$$\max_{(i,j)\in\mathcal{I}\times\mathcal{J}} u(i,j)$$

This task might occur at **training** or at **test** time.

Active feature acquisition

	feature 1	feature 2	feature 3	label
instance 1	x_{11}	x_{12}	**?**	y_1
instance 2	**?**	x_{22}	x_{23}	y_2
instance 3	x_{31}	x_{32}	x_{33}	y_3
instance 4	x_{41}	**?**	**?**	y_4

Goal: select feature values x_{ij} to acquire

$$\max_{(i,j)\in\mathcal{I}\times\mathcal{J}} u(i,j)$$

This task might occur at **training** or at **test** time.

AFA applications provide an oracle $f:\mathcal{I}\times\mathcal{J}\to\mathbb{R}$

- Medical diagnosis: select examinations (x_{ij}) to take out
- Preprocessing: select features (x_{ij}) to compute from raw data
- ...

IAL

Active feature acquisition

Approaches

method	idea
matrix completion[8]	minimize classification & reconstruction error, omit well-reconstructed queries
confidence cascade[9]	sort features by cost, acquire each next feature for all uncertain instances
instance completion[10]	select instances for which to acquire all features
⋮	

Today at 15:45: Beyer, Büttner, Spiliopoulou, “AFA and imputation on data streams”.

[8] Huang et al., “Active feature acquisition with supervised matrix completion”, 2018.
[9] desJardins et al., “Confidence-based feature acquisition to minimize training and test costs”, 2010.
[10] Zheng and Padmanabhan, “On active learning for data acquisition”, 2002.

IAL

Beyond active labeling

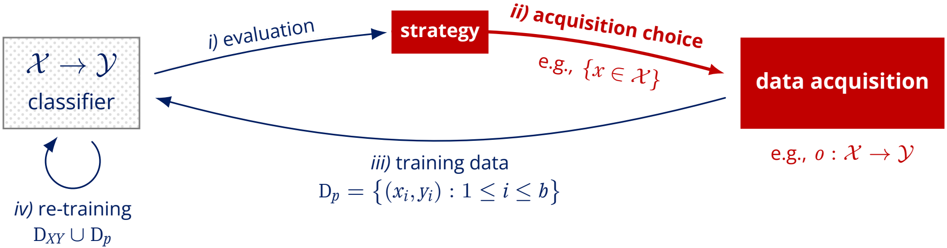

We often assume an oracle $o : \mathcal{X} \to \mathcal{Y}$, **but what if there is none?**

- lack of (human) expertise / lack of data interpretability
- extreme data volumes

Also, **labels aren't the only cost factor.**

IAL

Towards Explainable Active Learning using Meta-Learning

Amal Saadallah

IAL

Content

- Meta-Learning (Definition & Goal)
- Overview of Explainable Machine Learning
- Meta-Learning for Explainable Active Learning
- Example of Interpretable Active Sample Selection

IAL

Meta-Learning

Definition & Goal

Definition Learn over a series (distributions) of many different learning tasks.

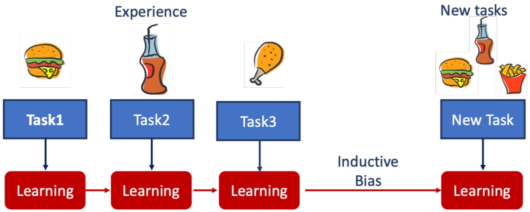

Meta-Learning
Definition & Goal

Definition Learn over a series (distributions) of many different learning tasks.

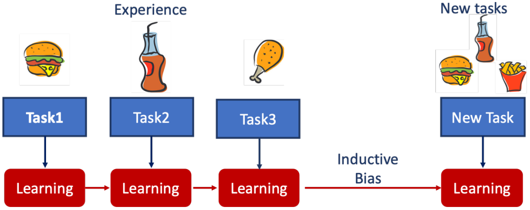

→ ***Learning-to-Learn***

Meta-Learning
Definition & Goal

Definition Learn over a series (distributions) of many different learning tasks.

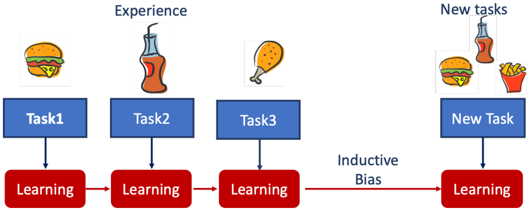

→ ***Learning-to-Learn***

Goal Enable models to acquire new knowledge or adapt quickly to new tasks with minimal data.

IAL

Meta-Learning

Definition Learn over a series (distributions) of many different learning tasks.

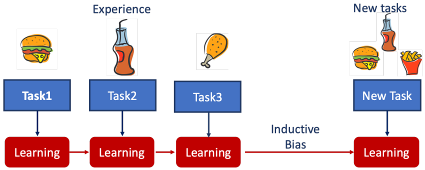

Goal Enable models to acquire new knowledge or adapt quickly to new tasks with minimal data.

Use Cases rare events, test-time constraints, data collection costs, etc.

IAL

Meta-Learning

Meta-Supervised Learning

- **Supervised Learning**
  Input: x, **Output:** y, $(x_i, y_i) \in \mathbb{D}$
  Goal: Learn a function $\hat{f}_\theta : \mathbf{X} \rightarrow \mathbf{Y}$ such that

$$\hat{f}_\theta(x_i, \theta) \approx f(x_i) = y_i, \forall x_i \in \mathbf{X}, y_i \in \mathbf{Y} \tag{1}$$

  where $\theta \in \mathbb{R}^n$ is an unknown (hyper)parameters vector learnt using $\mathbb{D}$.

- **Meta Supervised Learning?**

IAL

Meta-Learning

Meta-Supervised Learning

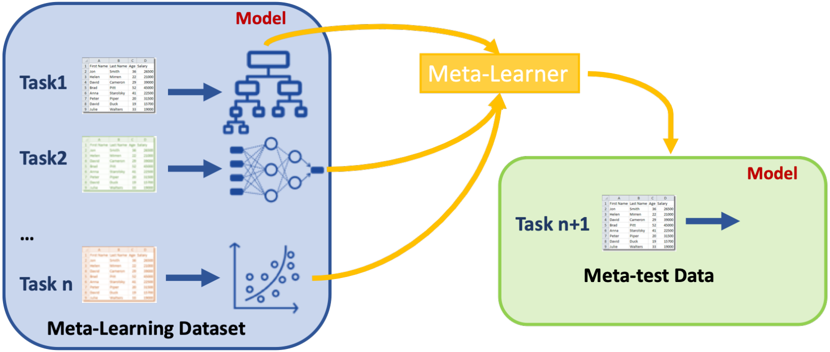

Meta-Learning for Active Learning

Taxonomy

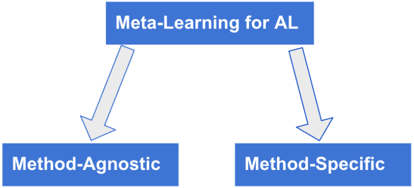

IAL

Meta-Learning for Active Learning

Taxonomy

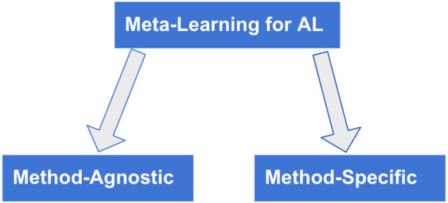

Instance={labelled and/or unlabelled data points, Budget, Active Learning method}

Finn, C., Xu, K., & Levine, S. (2018), Yoon, Jaesik, et al. (2019), Contardo,et al. (2017)

IAL

Meta-Learning for Active Learning

Taxonomy

Meta-Learning for AL

Method-Agnostic | **Method-Specific**

Instance={Labelled and/or unlabelled data points, Budget, Active Learning method}

Finn, C., Xu, K., & Levine, S. (2018), **Yoon, Jaesik, et al. (2019), Contardo, et al. (2017), …**

Instance={Labelled and/or unlabelled data points, Budget, Specific Active Learning Criterion: *online tuning of the Uncertainty Sampling threshold, Loss reduction* **}**

Pang, Kunkun, et al. (2018), Ravi, S., & Larochelle, H. (2018), Martins, V. E., Cano, A., & Junior, S. B. (2023), Taguchi, et al. (2019), Saadallah, O., & Rouissi, Z. (2023),…

IAL

eXplainable Machine Learning XML

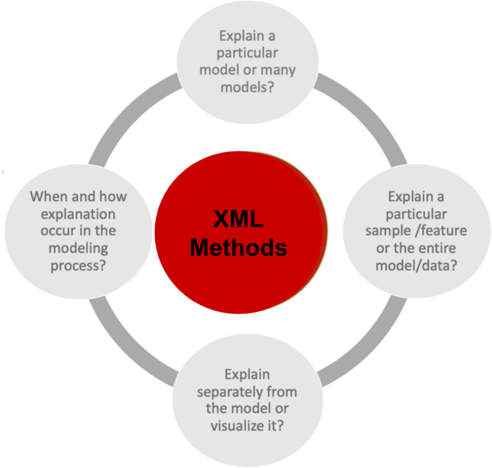

eXplainable Machine Learning XML

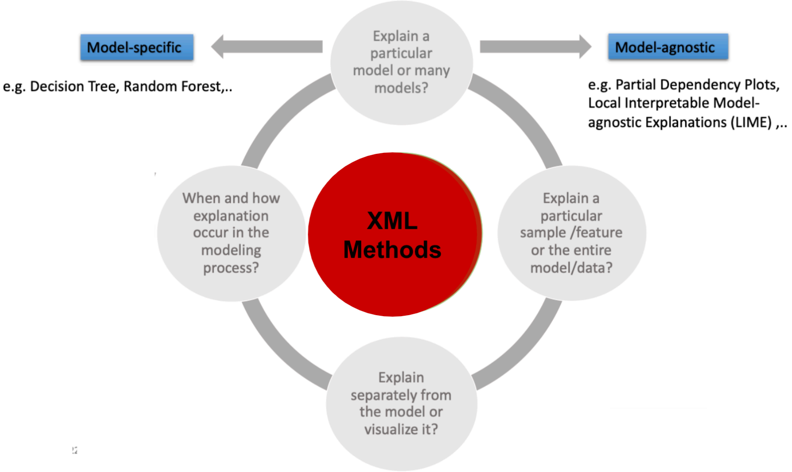

IAL

eXplainable Machine Learning XML

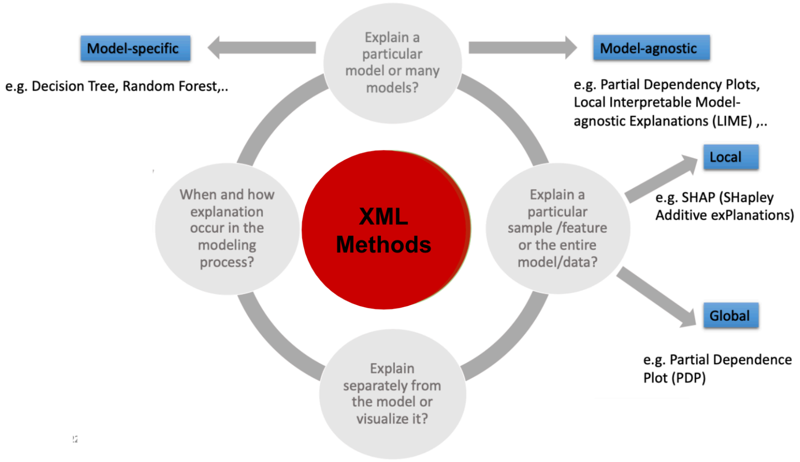

eXplainable Machine Learning XML

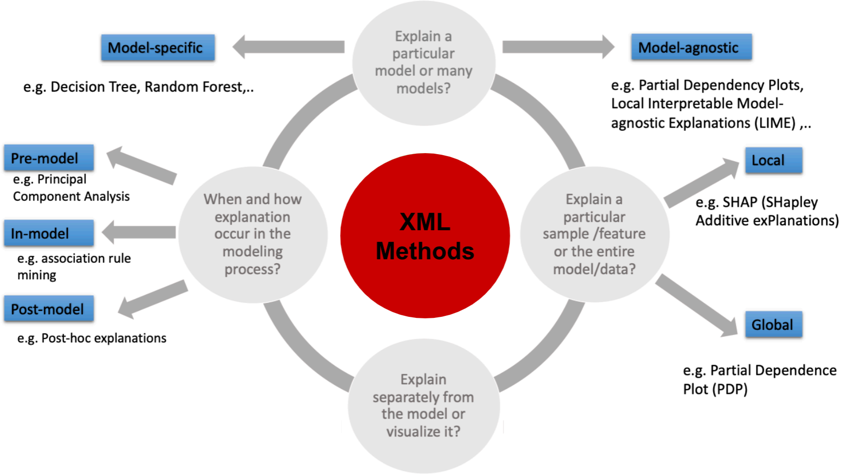

eXplainable Machine Learning XML

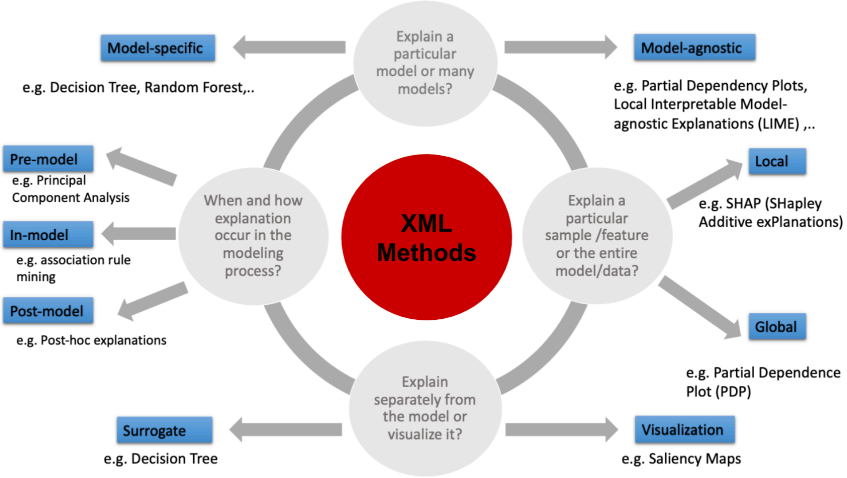

IAL

Meta-Learning for Explainable Active Learning

Interpretable Meta-Model

How?

- Use interpretable base models like decision trees, rule-based models, or linear models within the meta-learning framework.
- Integration of interpretable models that enhance the overall AL pipeline interpretability.

Consequences

- These models inherently offer transparency compared to complex, black-box architectures.
- Transparency in how the meta-model combines information from tasks leads to an interpretable AL system.

IAL

Meta-Learning for Explainable Active Learning

Explainable Active Sample Selection

How?

- Provide explanations for the model's selected samples during Active Learning (AL).
- Use techniques like uncertainty estimation, saliency maps, or gradient-based attribution to justify sample selection.

Consequences

- Explanations guide human annotators in understanding why certain samples are chosen for labeling:
  - — Importance to the model's decision
  - — Contribution to the input distribution

IAL

Meta-Learning for Explainable Active Learning

Attention Mechanisms

How?

- Employ neural networks/ reinforcement learning with neural networks
- Use attention mechanisms to highlight important input features, e.g., gradient-based ...
- Visualize attention maps to understand the model's focus.

Consequences

- Identify key factors influencing the model's decision for data instance selection.

IAL

Meta-Learning for Explainable Active Learning

Post-hoc Explanation Techniques

How?

- Integrate post-hoc explanation methods into meta-learning.
- Utilize LIME or SHAP to generate local explanations for individual data instances' predictions.

Consequences

- Gain insights into specific factors driving data point selection decisions.

IAL

Meta-Learning for Explainable Active Learning

Regularization with Explainability Constraints

How?

- Incorporate regularization terms into the meta-learning optimization process.
- Examples include discouraging complex decision boundaries or enforcing feature importance sparsity.

Consequences

- Encourage models to produce more interpretable decisions regarding the active sample selection process.

IAL

Meta-Learning for Explainable Active Learning

Human-in-the-Loop Feedback

How?

- Involve human annotators in the Active Learning process.
- Gather feedback from annotators to refine model explanations.

Consequences

- Explanations aligned with human understanding and preferences for improved interpretability.

IAL

Example of Interpretable Active Sample Selection

Interpretation of the sample selection for Bike dataset

feature	description
season	four seasons
yr	year (0: 2011, 1:2012)
mnth	month (1 to 12)
hr	hour (0 to 23)
holiday	whether day is holiday or not
wkday	day of the week
wkgday	if day is neither weekend nor holiday
wsit	variable encoding the weather situation
temp	normalized temperature
atemp	normalized feeling temperature
hum	normalized humidity
windspeed	normalized wind speed
cnt	[response] total number of rental bikes

Table: Feature OF Bike Dataset

Taguchi, Yusuke, Keisuke Kameyama, and Hideitsu Hino. ”Active Learning with Interpretable Predictor.” 2019 International Joint Conference on Neural Networks (IJCNN). IEEE, 2019.

Example of Interpretable Active Sample Selection

Interpretation of the sample selection for Bike dataset

feature	description
season	four seasons
yr	year (0: 2011, 1:2012)
mnth	month (1 to 12)
hr	hour (0 to 23)
holiday	whether day is holiday or not
wkday	day of the week
wkgday	if day is neither weekend nor holiday
wsit	variable encoding the weather situation
temp	normalized temperature
atemp	normalized feeling temperature
hum	normalized humidity
windspeed	normalized wind speed
cnt	[response] total number of rental bikes

Table: Feature OF Bike Dataset

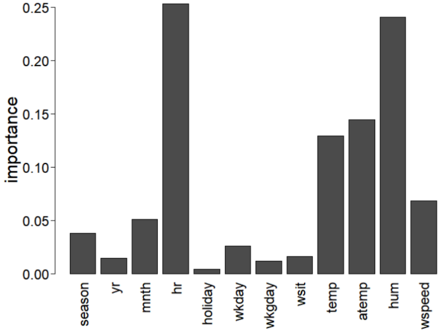

Figure: Variable importance for the main model before 14-th sample selection

Example of Interpretable Active Sample Selection

Analysis at the 14-th iteration of active sample selection:

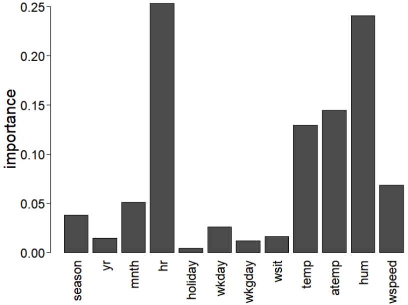

Figure: Variable importance for the main model before 14-th sample selection

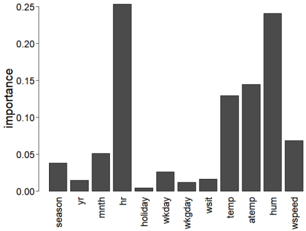

Figure: Variable importance for the Meta-model before 14-th sample selection

Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.1700	0.4500	0.5800	0.6012	0.7950	1.0000

Figure: Summary of Hum before the 14-th data point selection.

$\rightarrow$ The value of **Hum** in the actually selected sample at the 14-th iteration of the algorithm was **0.24**.

IAL

Applications and Practical Challenges, and Closing Discussion

Alaa Tharwat Othman

Practical Challenges of AL in Real Environments

- The Imbalanced Data Problem

[11] Tharwat and Schenck, “A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions”, 2023.

IAL

Practical Challenges of AL in Real Environments

- The Imbalanced Data Problem
- Low Query Budget

[11] Tharwat and Schenck, "A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions", 2023.

IAL

Practical Challenges of AL in Real Environments

- The Imbalanced Data Problem
- Low Query Budget
- Using Initial Knowledge for Training Learning Models

[11] Tharwat and Schenck, “A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions”, 2023.

IAL

Practical Challenges of AL in Real Environments

- The Imbalanced Data Problem
- Low Query Budget
- Using Initial Knowledge for Training Learning Models
- The Concept Drift Phenomenon in Data Streams

[11] Tharwat and Schenck, “A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions”, 2023.

IAL

Practical Challenges of AL in Real Environments

- The Imbalanced Data Problem
- Low Query Budget
- Using Initial Knowledge for Training Learning Models
- The Concept Drift Phenomenon in Data Streams
- Stopping Criteria

[11] Tharwat and Schenck, “A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions”, 2023.

IAL

Practical Challenges of AL in Real Environments

- The Imbalanced Data Problem
- Low Query Budget
- Using Initial Knowledge for Training Learning Models
- The Concept Drift Phenomenon in Data Streams
- Stopping Criteria
- Error prone and/or multiple oracles
- Noisy Labeled Data
- AL with Crowdsourcing Labelers

[11] Tharwat and Schenck, “A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions”, 2023.

IAL

Practical Challenges of AL in Real Environments

- The Imbalanced Data Problem
- Low Query Budget
- Using Initial Knowledge for Training Learning Models
- The Concept Drift Phenomenon in Data Streams
- Stopping Criteria
- Error prone and/or multiple oracles
- Noisy Labeled Data
- AL with Crowdsourcing Labelers
- AL with Outliers

[11] Tharwat and Schenck, “A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions”, 2023.

IAL

Practical Challenges of AL in Real Environments

- The Imbalanced Data Problem
- Low Query Budget
- Using Initial Knowledge for Training Learning Models
- The Concept Drift Phenomenon in Data Streams
- Stopping Criteria
- Error prone and/or multiple oracles
- Noisy Labeled Data
- AL with Crowdsourcing Labelers
- AL with Outliers
- AL in High-Dimensional Environments

[11] Tharwat and Schenck, “A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions”, 2023.

IAL

Practical Challenges of AL in Real Environments

- The Imbalanced Data Problem
- Low Query Budget
- Using Initial Knowledge for Training Learning Models
- The Concept Drift Phenomenon in Data Streams
- Stopping Criteria
- Error prone and/or multiple oracles
- Noisy Labeled Data
- AL with Crowdsourcing Labelers
- AL with Outliers
- AL in High-Dimensional Environments
- ML-Based Active Learners[11]

[11] Tharwat and Schenck, "A Survey on Active Learning: State-of-the-Art, Practical Challenges and Research Directions", 2023.

IAL

AL with Different Technologies (Research Areas)

- AL with Deep Learning

IAL

AL with Different Technologies (Research Areas)

- AL with Deep Learning
- AL with Optimization

IAL

AL with Different Technologies (Research Areas)

- AL with Deep Learning
- AL with Optimization
- AL with Simulation

IAL

AL with Different Technologies (Research Areas)

- AL with Deep Learning
- AL with Optimization
- AL with Simulation
- AL with Design of Experiments

IAL

AL with Different Technologies (Research Areas)

- AL with Deep Learning
- AL with Optimization
- AL with Simulation
- AL with Design of Experiments
- Few-Shot Learning with AL

Lunch Break (13:00–14:00)

14:00–16:00	**Session 3: Keynote & Workshop Contributions**	
14:00–14:40	Keynote: From Insights to Impact: A Metrics-Driven Active Learning Journey	A. Abraham
14:40–15:00	Active Learning for Survival Analysis with Incrementally Disclosed Label Information	K. Dedja, F.K. Nakano & C. Vens
15:00–15:15	Towards Enhancing Deep Active Learning with Weak Supervision and Constrained Clustering	M. Aßenmacher, L. Rauch, J. Goschenhofer, A. Stephan, B. Bischl, B. Roth & B. Sick
15:15–15:30	Who knows best? A Case Study on Intelligent Crowdworker Selection via Deep Learning	M. Herde, D. Huseljic, B. Sick, U. Bretschneider & S. Oeste-Reiß
15:30–15:45	Role of Hyperparameters in Deep Active Learning	D. Huseljic, M. Herde, P. Hahn & B. Sick
15:45–16:00	Challenges for Active Feature Acquisition and Imputation on Data Streams	C. Beyer, M. Büttner & M. Spiliopoulou

Coffee Break (16:00–16:30)

Coffee Break (16:00–16:30)

16:30–17:40	**Session 4: Workshop Contributions & Closing**	
16:30–16:50	Active Learning with Fast Model Updates and Class-Balanced Selection for Imbalanced Datasets	Z. Huang, Y. He, M. Herde, D. Huseljic & B. Sick
16:50–17:10	Interpretable Meta-Active Learning for Regression Ensemble Learning	O. Saadallah & Z. Rouissi
17:10–17:30	Look and You Will Find It: Fairness-Aware Data Collection through Active Learning	H. Weerts, R. Theunissen & M. Willemsen
17:30–17:40	Closing	

Thank you!