How to Allocate your Label Budget? Choosing between Active Learning and Learning to Reject in Anomaly Detection

https://people.cs.kuleuven.be/~lorenzo.perini/

@LorenzoPerini95

Lorenzo Perini, Daniele Giannuzzi, Jesse Davis

*1st AAAI Workshop on Uncertainty Reasoning and Quantification in Decision Making (**UDM23**)*

Anomaly detection is the task of detecting the examples that do not follow an expected behaviour

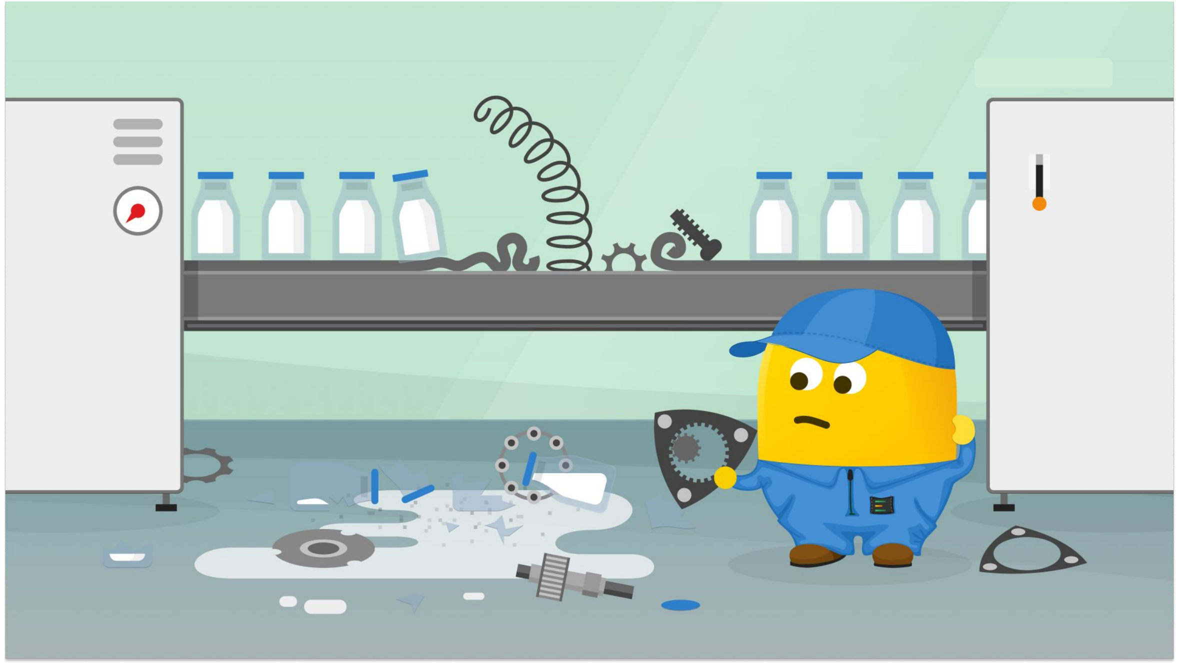

Anomalies are critical adverse events associated with monetary costs.

When using an anomaly detector for decision-making, it is crucial that the user trusts the system

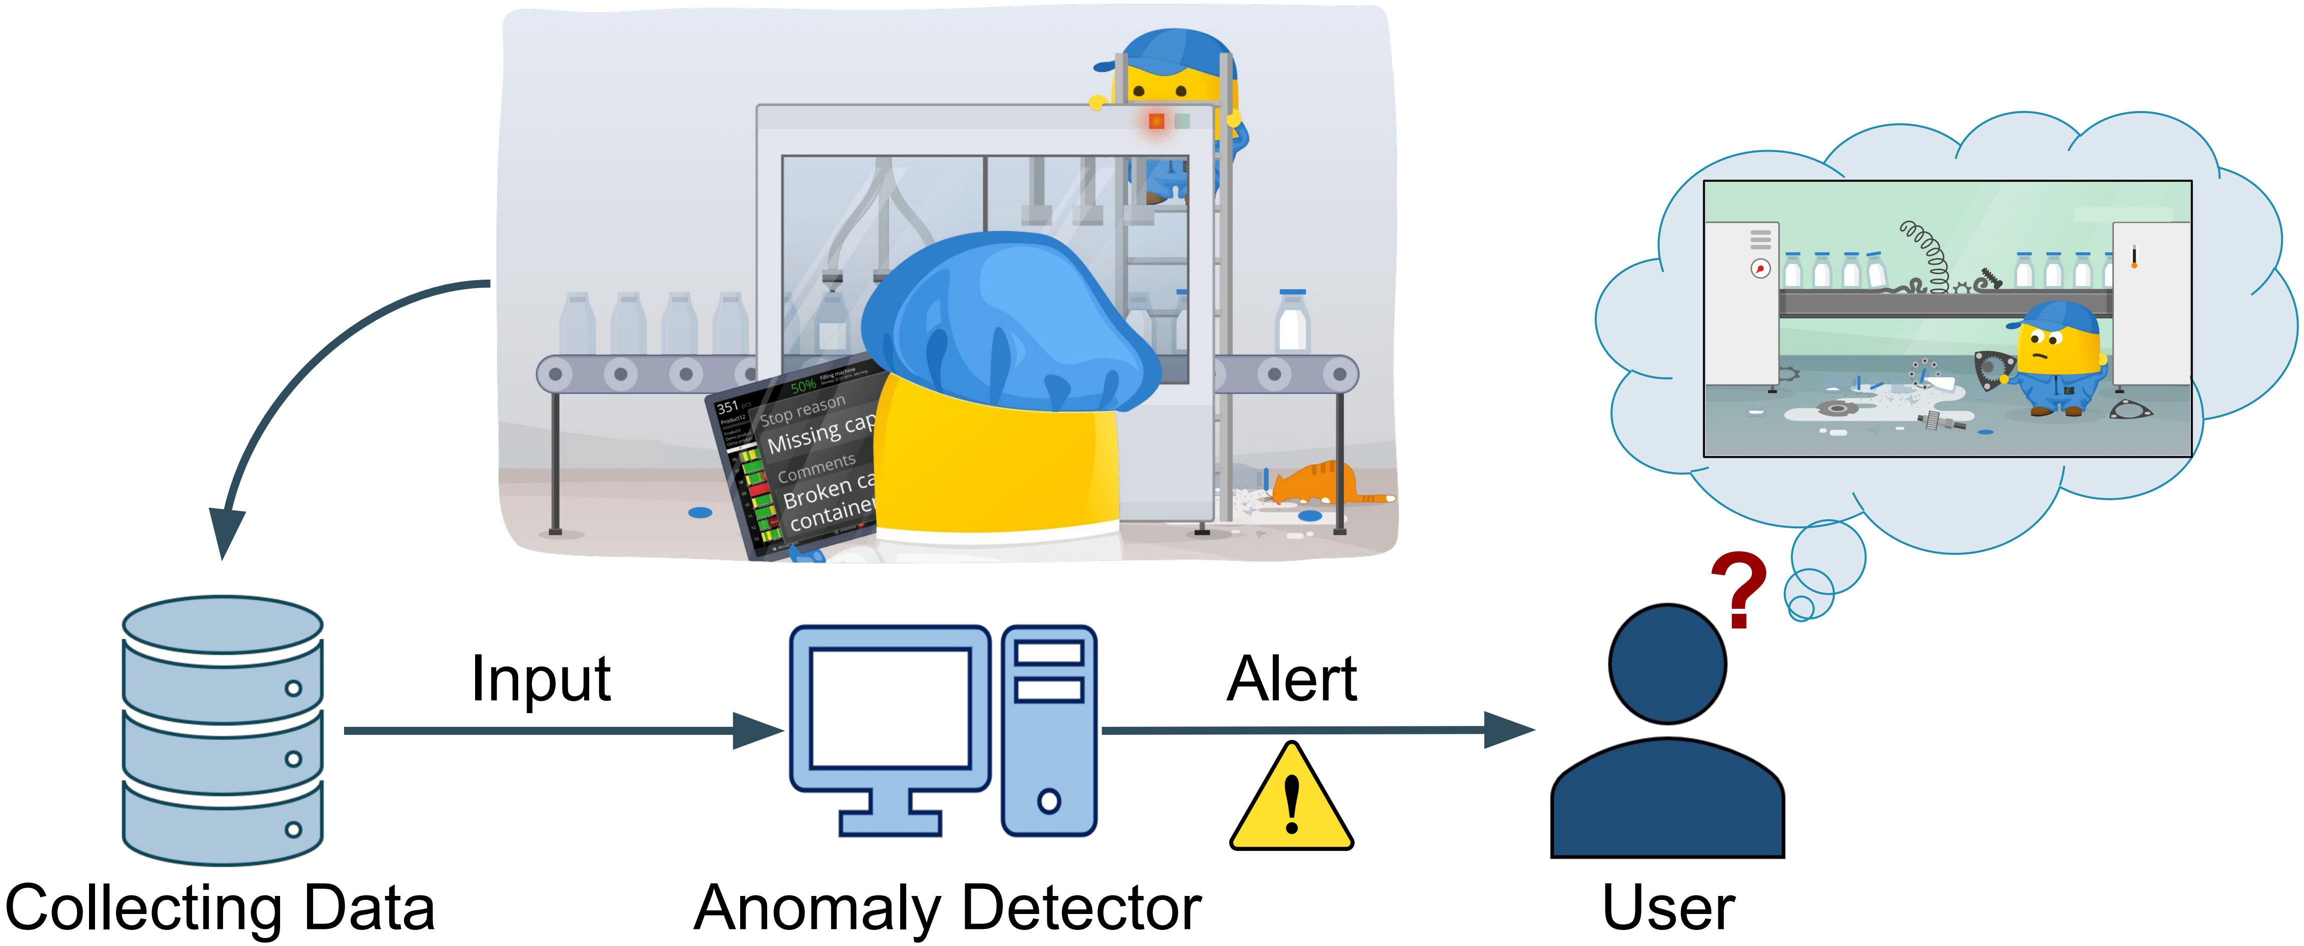

Unfortunately, Anomaly Detection is often tackled as an <u>unsupervised</u> task because anomalies are rare

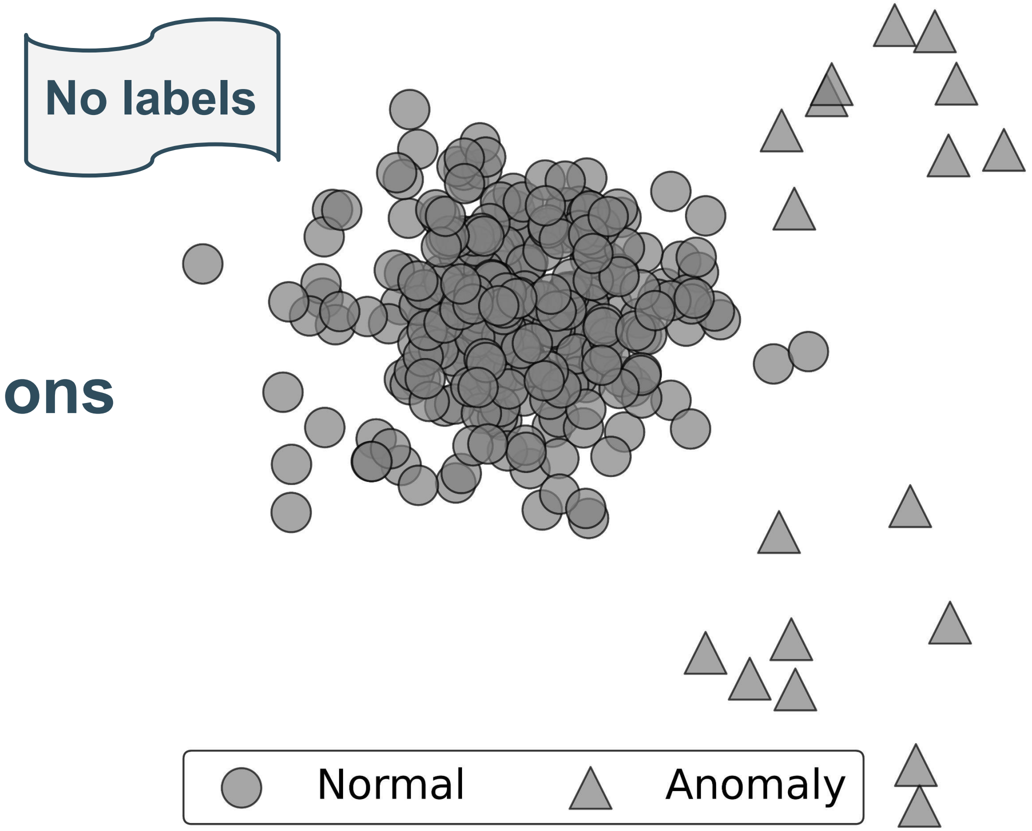

➔ Anomaly detectors employ **heuristic intuitions** to learn a decision boundary without labels.

➔ Intuitions are **hard to verify** and **may not hold** in some cases.

Some predictions may have <u>high uncertainty</u>, especially for the examples close to the decision boundary

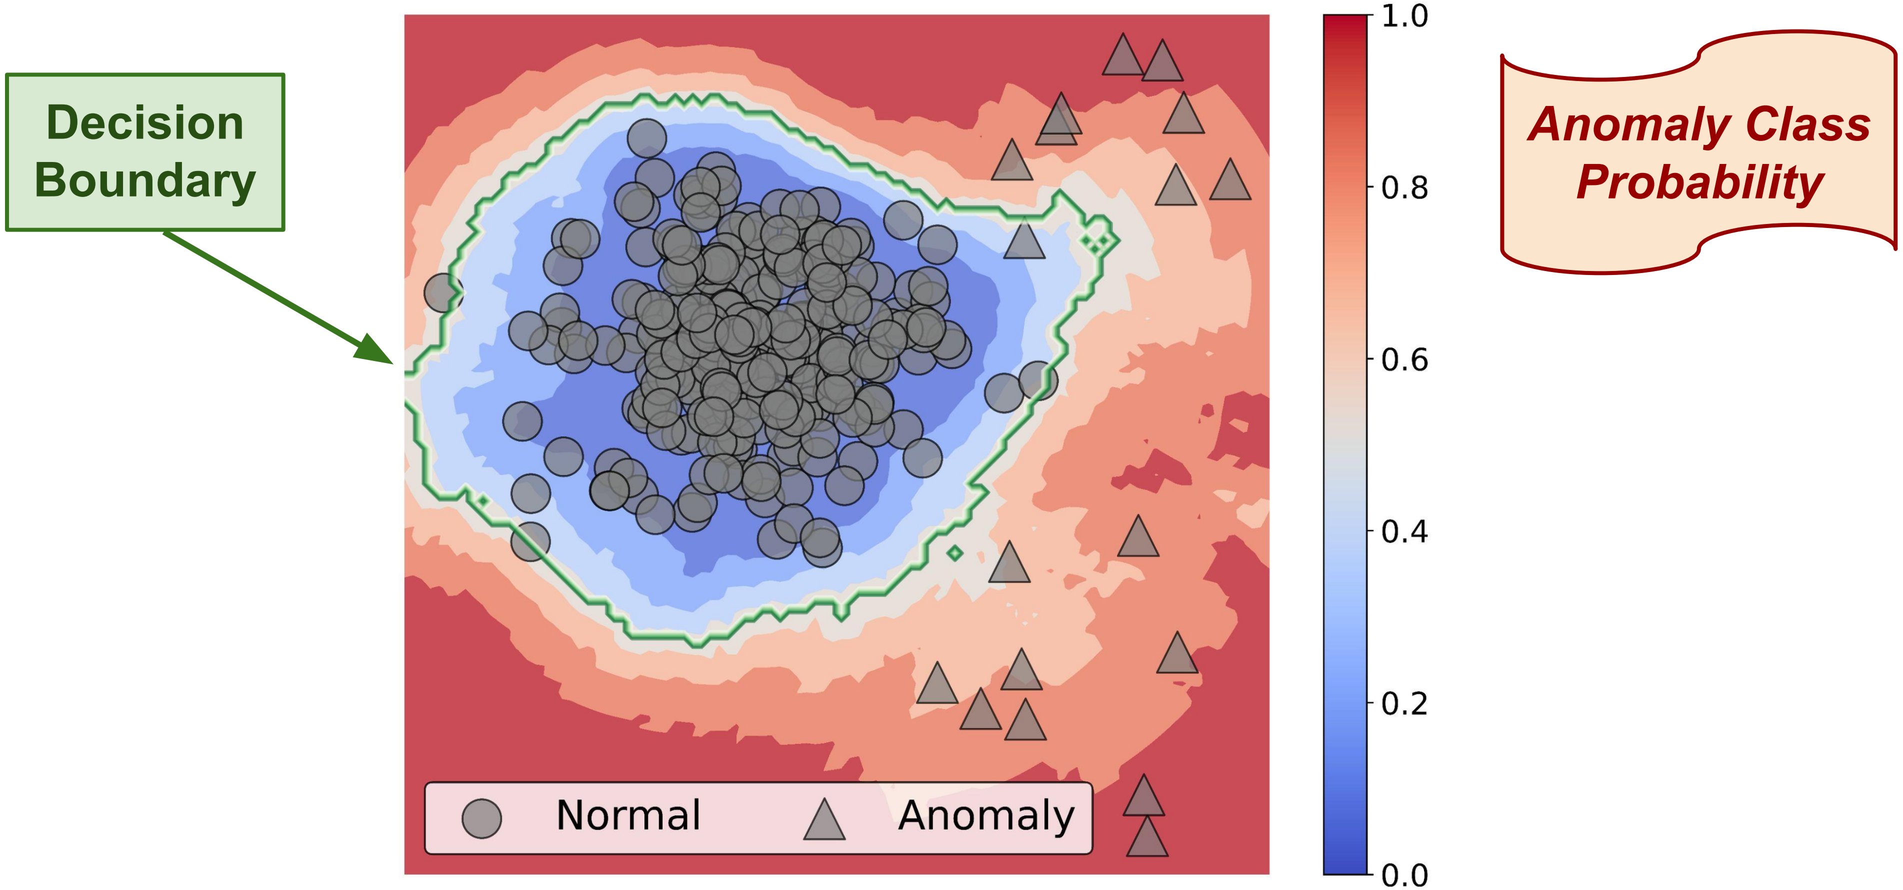

How can we reduce such uncertainty?

1) Learning a more accurate detector by acquiring a limited number of labels using <u>Active Learning</u>

Active Learning selects strategic training examples ✖ to query their label to a user.

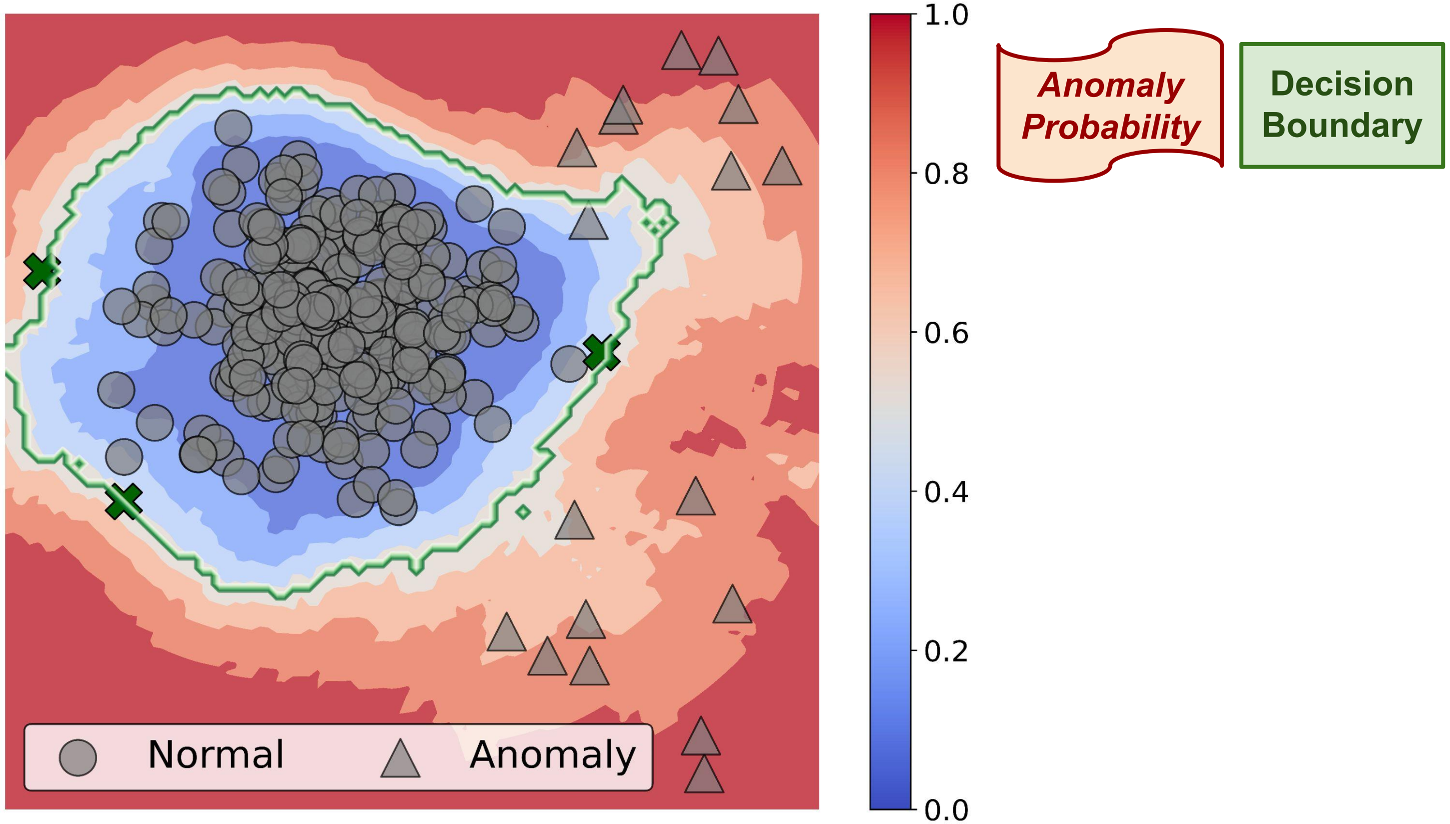

1) Learning a more accurate detector by acquiring a limited number of labels using Active Learning

Semi-Supervised detectors use the training labels ● to reduce the uncertainty.

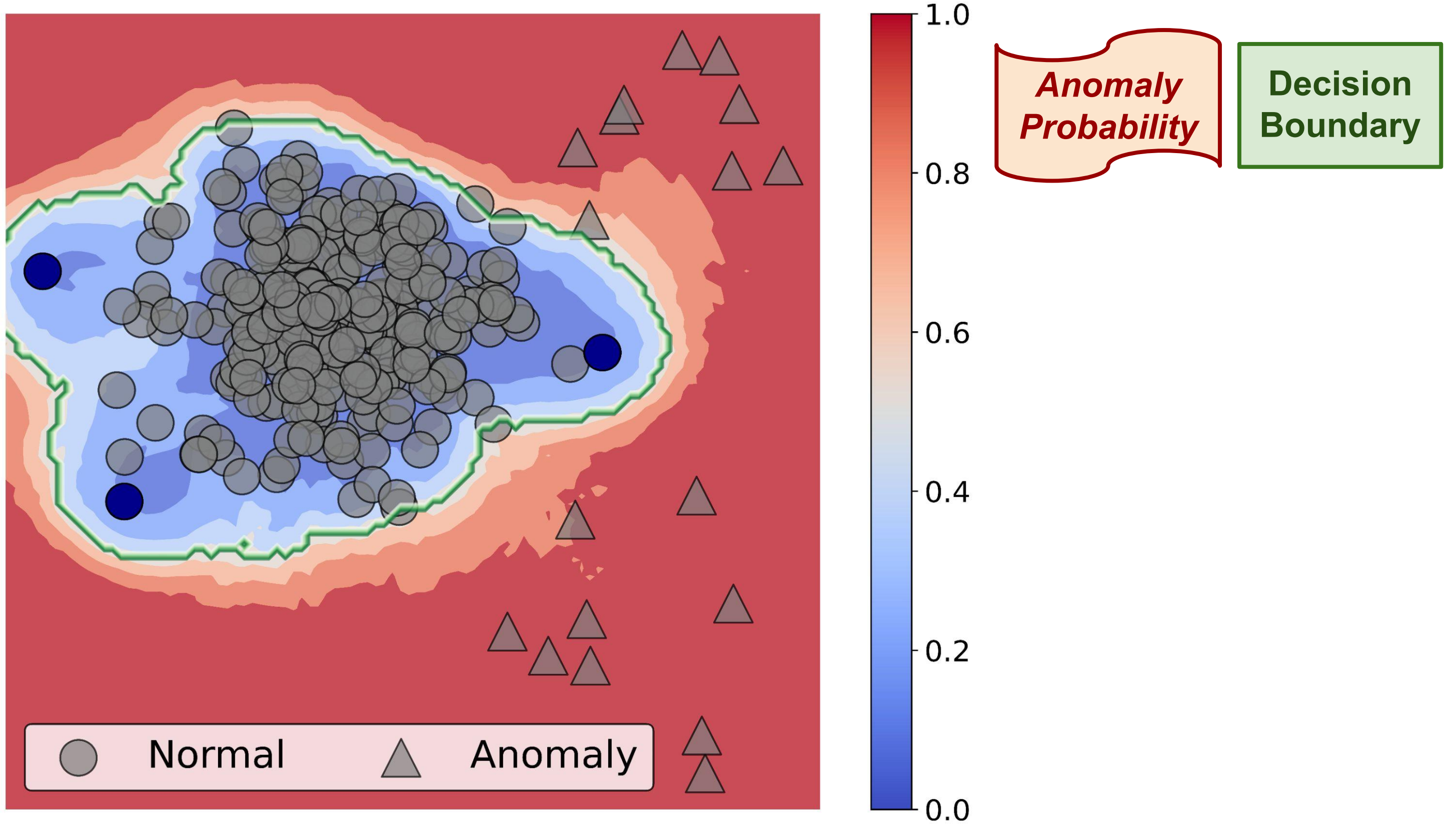

2) Allowing the detector to reject high-uncertainty predictions using Learning to Reject

Learning to Reject uses i.i.d. validation labels to learn the rejection region.

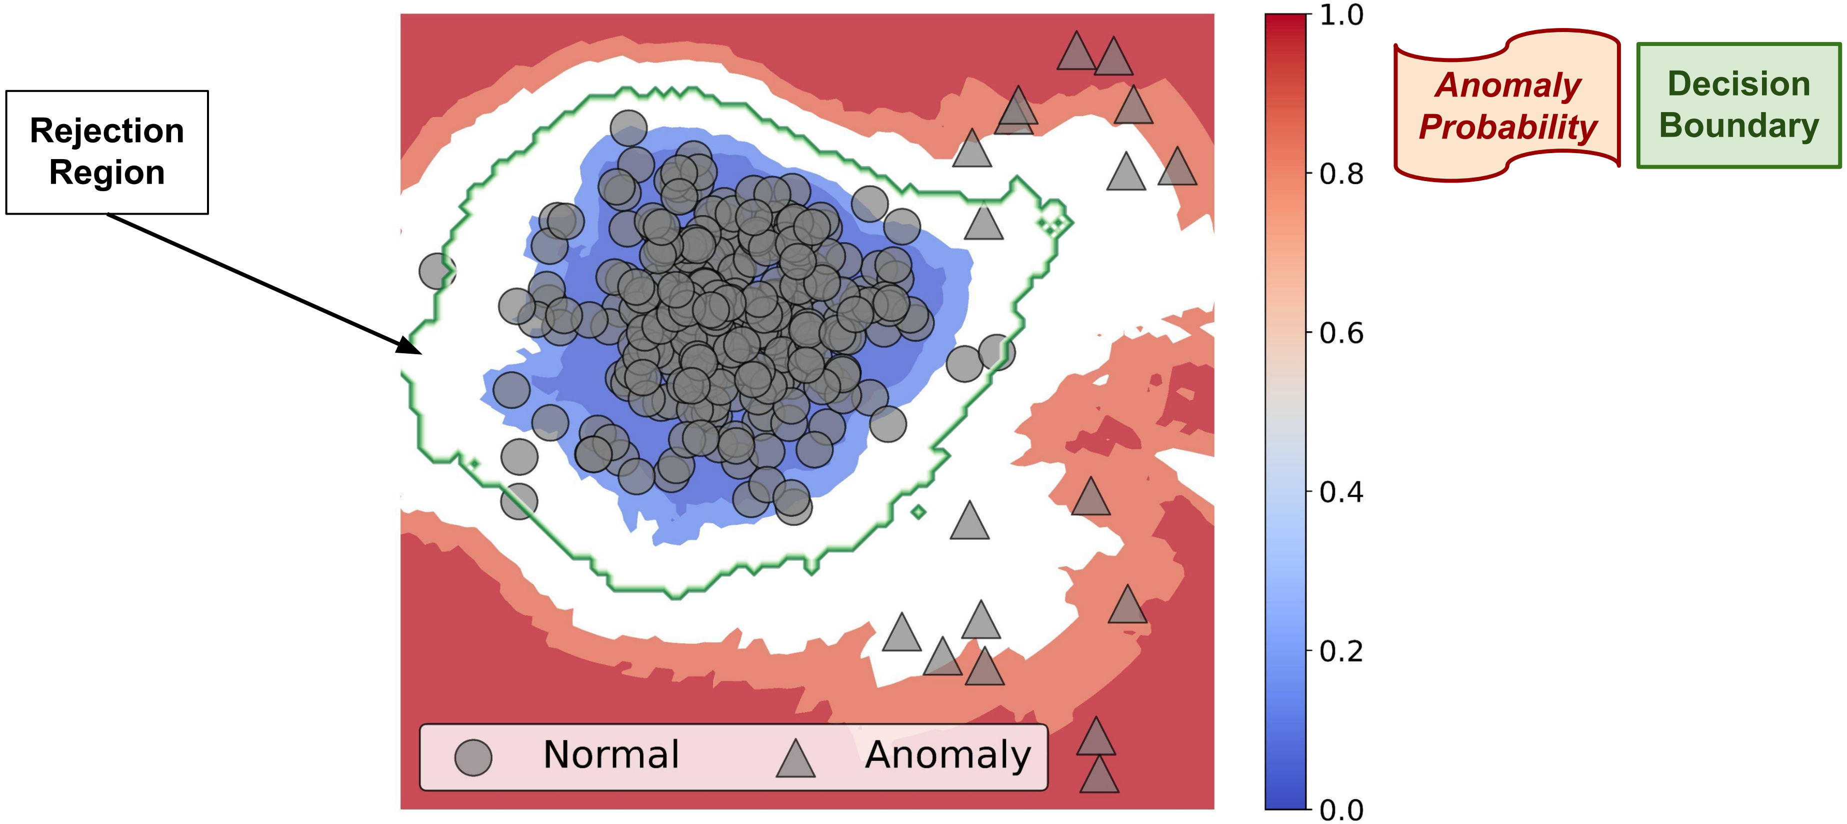

2) Allowing the detector to reject high-uncertainty predictions using Learning to Reject

Detectors with rejection have lower uncertainty at test time.

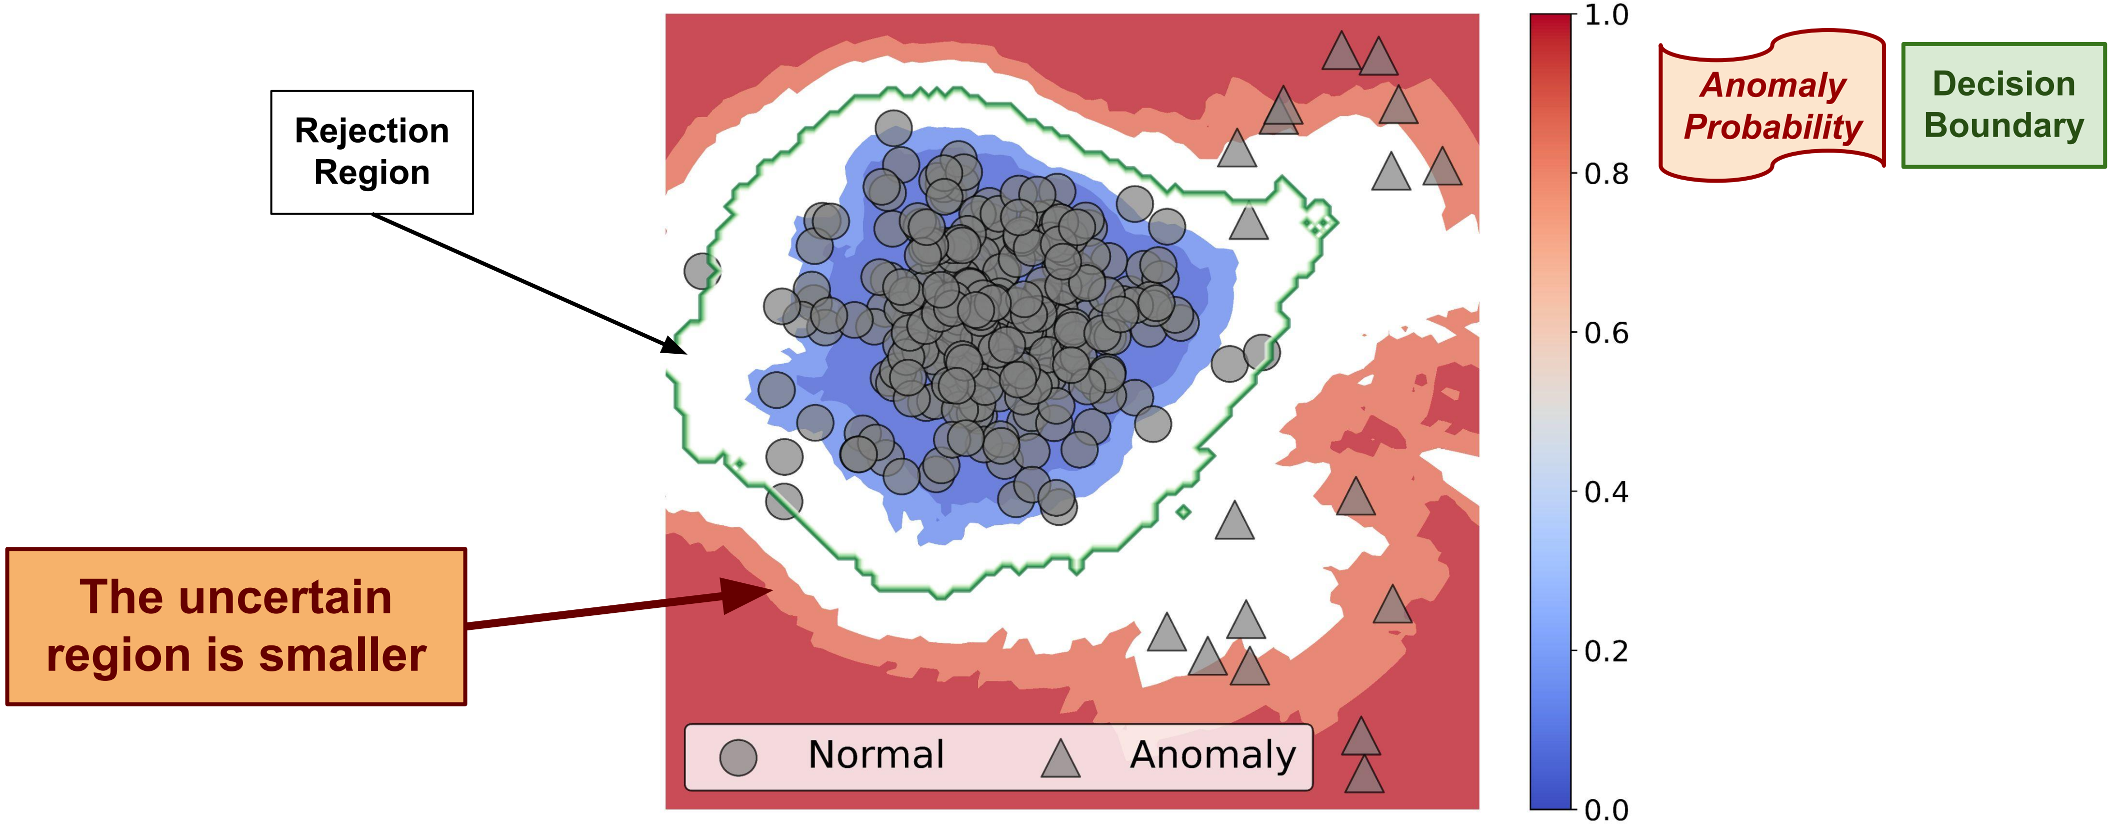

Problem: AL and LR rely on different type of labels

Active Learning uses

➔ Training biased labels
➔ Task: improve training phase

Learning to Reject uses

➔ Validation i.i.d. labels
➔ Task: improve test phase

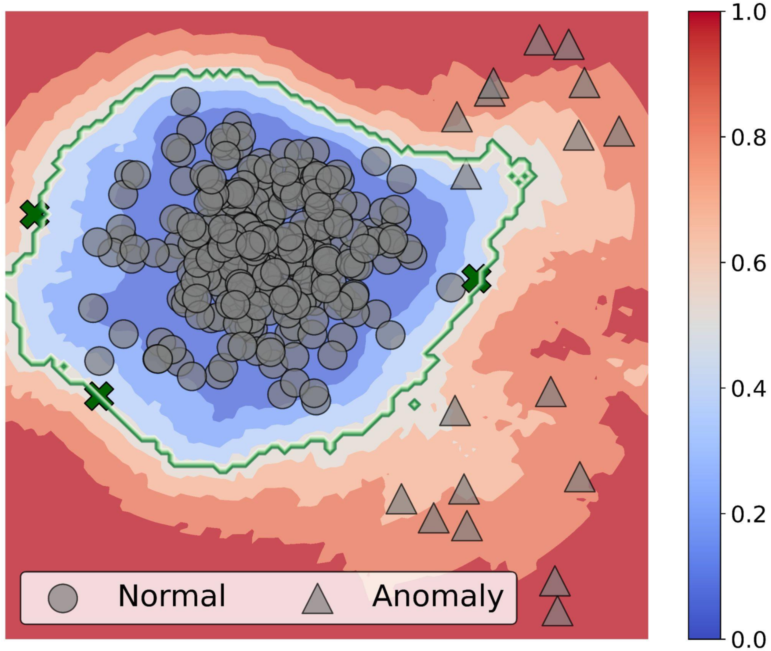

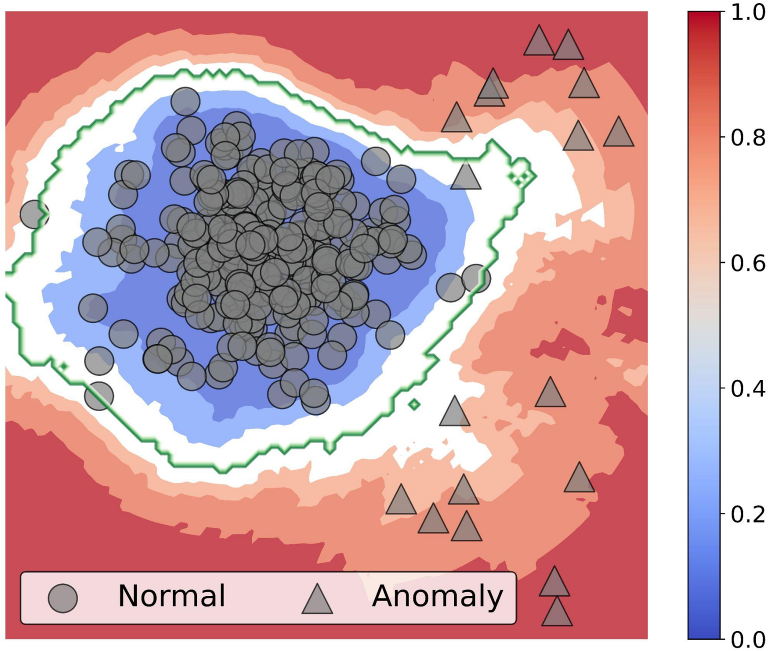

Anomaly Probability

Decision Boundary

Rejection Region

Usually the label budget is limited:

How do you choose between Active Learning and Learning to Reject?

➔ **Our goal is to find a strategy to decide how to allocate the label budget, i.e. how to split it between AL and LR.**

↳ **We measure the reward of allocating the label budget to either side.**

BALLAD: Budget allocation for Active Learning and Learning to reject in Anomaly Detection

Budget *B* to be used in *k* rounds by querying *g* examples at the time;

➔ Initialization:

a. train the detector with no labels;

b. collect g random labels for LR and for AL;

c. compute the initial rewards by measuring how the detector varies (**a** - **b**).

➔ Allocation loop:

a. Allocate the budget to the option with the highest reward;

b. Update the rewards.

Because we do not know how beneficial is the next allocation round, we look at the past reward

Challenge: designing a reward function that reflects the gain when querying the labels

Entropy

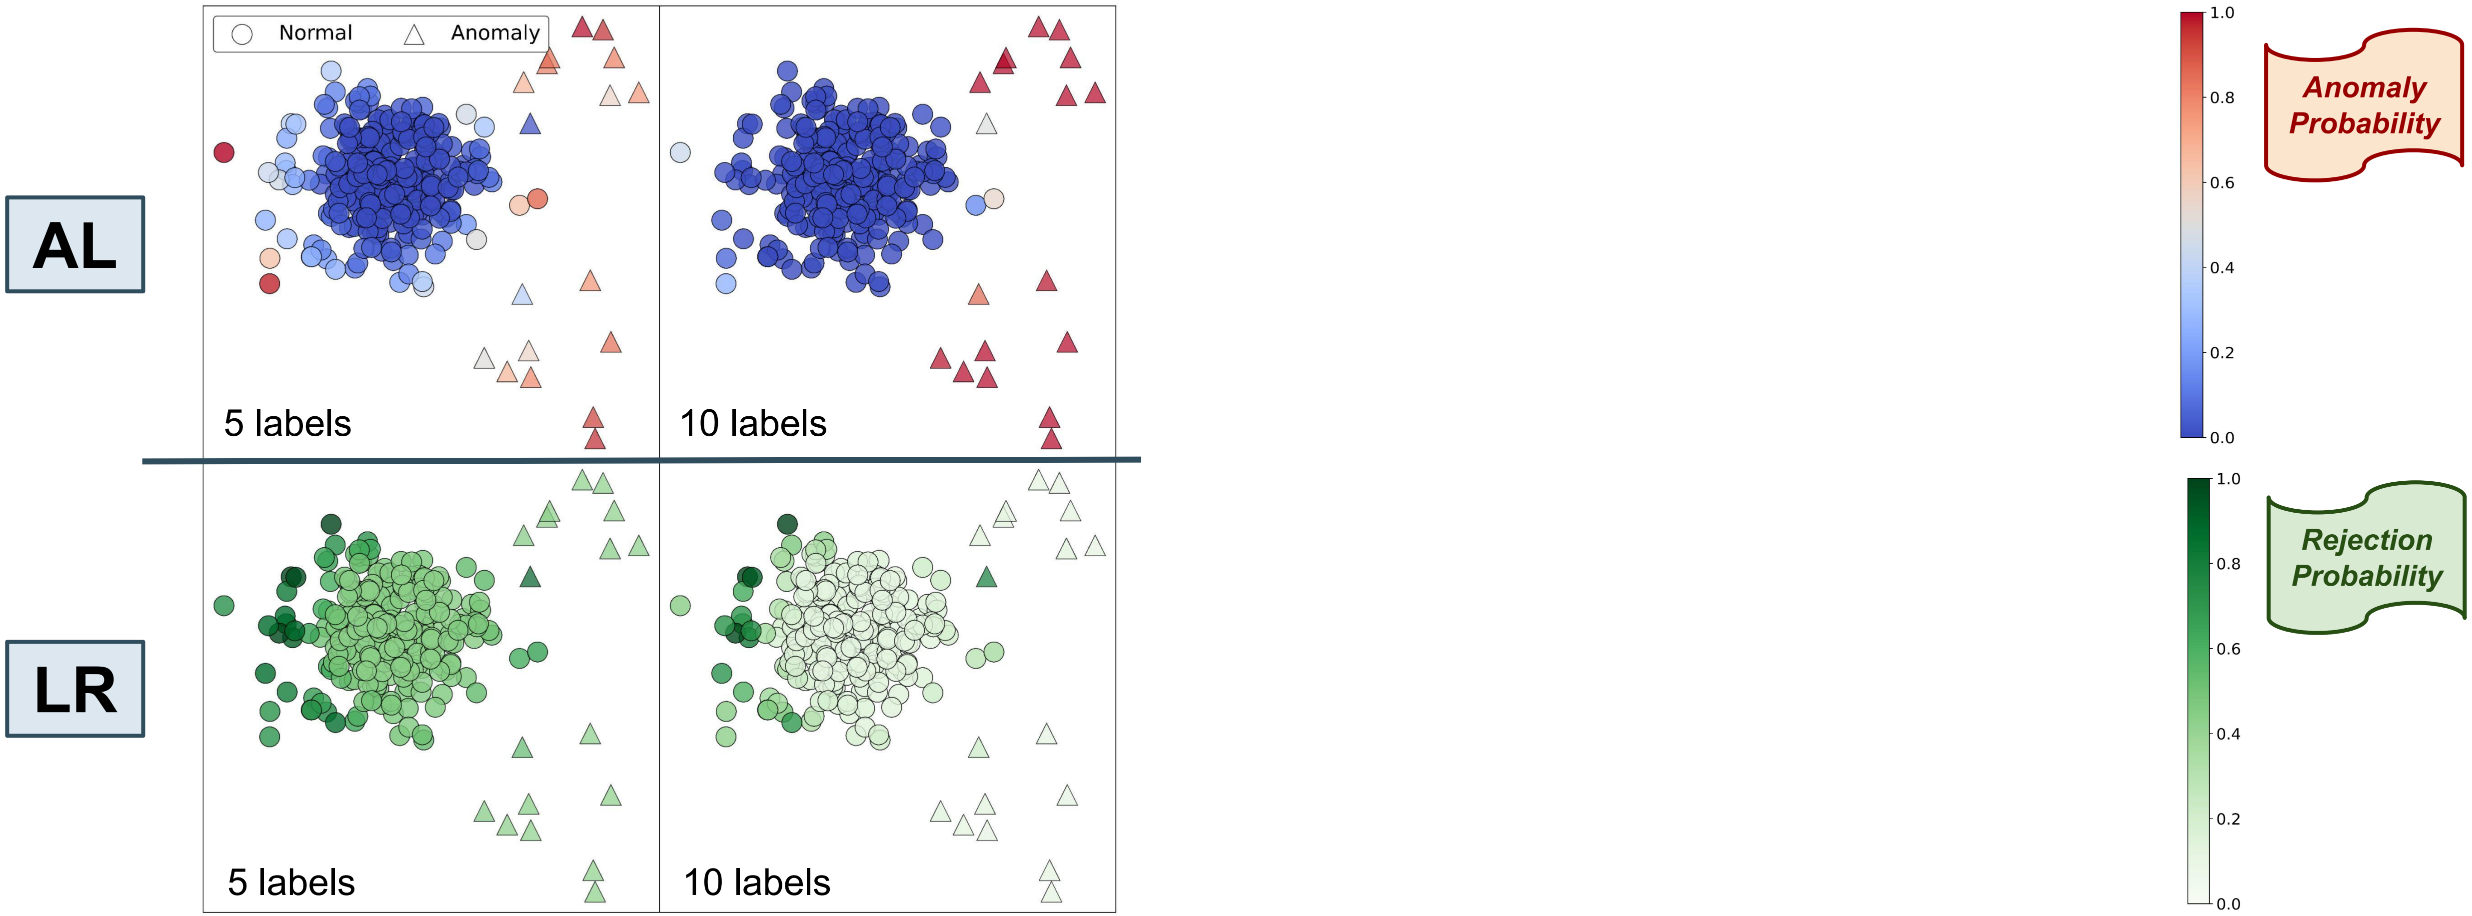

Challenge: designing a reward function that reflects the gain when querying the labels

Cosine Similarity

AL

LR

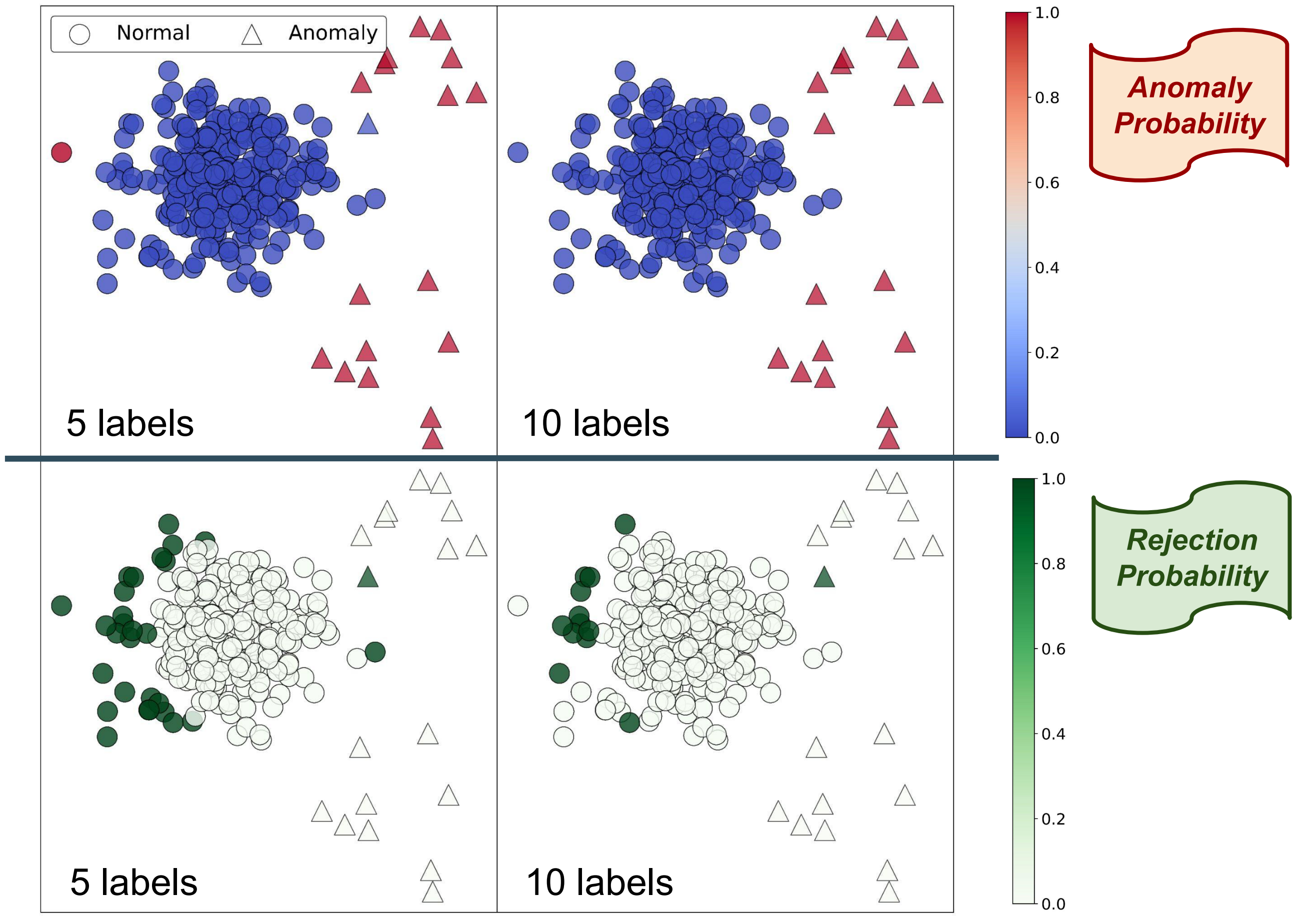

Detectors with different reject options have different test sets: how can we compare them?

We use a **cost-based** evaluation metric that penalizes:

1. False Positives by $\mathbf{c_{fp}}$;
2. False Negatives by $\mathbf{c_{fn}}$;
3. Rejections by $\mathbf{c_r}$;

Costs depend on the **application domain**, but they need to satisfy:

$$c_r \leq min\{c_{fp} \times (1 - \gamma), c_{fn} \times \gamma\}$$

where gamma **γ** is the proportion of anomalies in the dataset.

Experiments

Empirical evaluation on 18 benchmark datasets (Campos et al. 2016)

Experimental setup

- ➔ Baselines:
  - ★ All-in-AL only uses strategic labels by **Active Learning**;
  - ★ All-in-LR only uses random labels by **Learning to Reject**;
- ➔ Fixed costs $\mathbf{c_{fp} = c_{fn} = 1}$;
- ➔ We run 15 allocation rounds with 2% (of dataset size) labels each;
- ➔ We experimentally investigate <u>four</u> research questions.

Q1. Does BALLAD result in lower costs when compared to All-in-AL and All-in-LR ?

The lower the better!

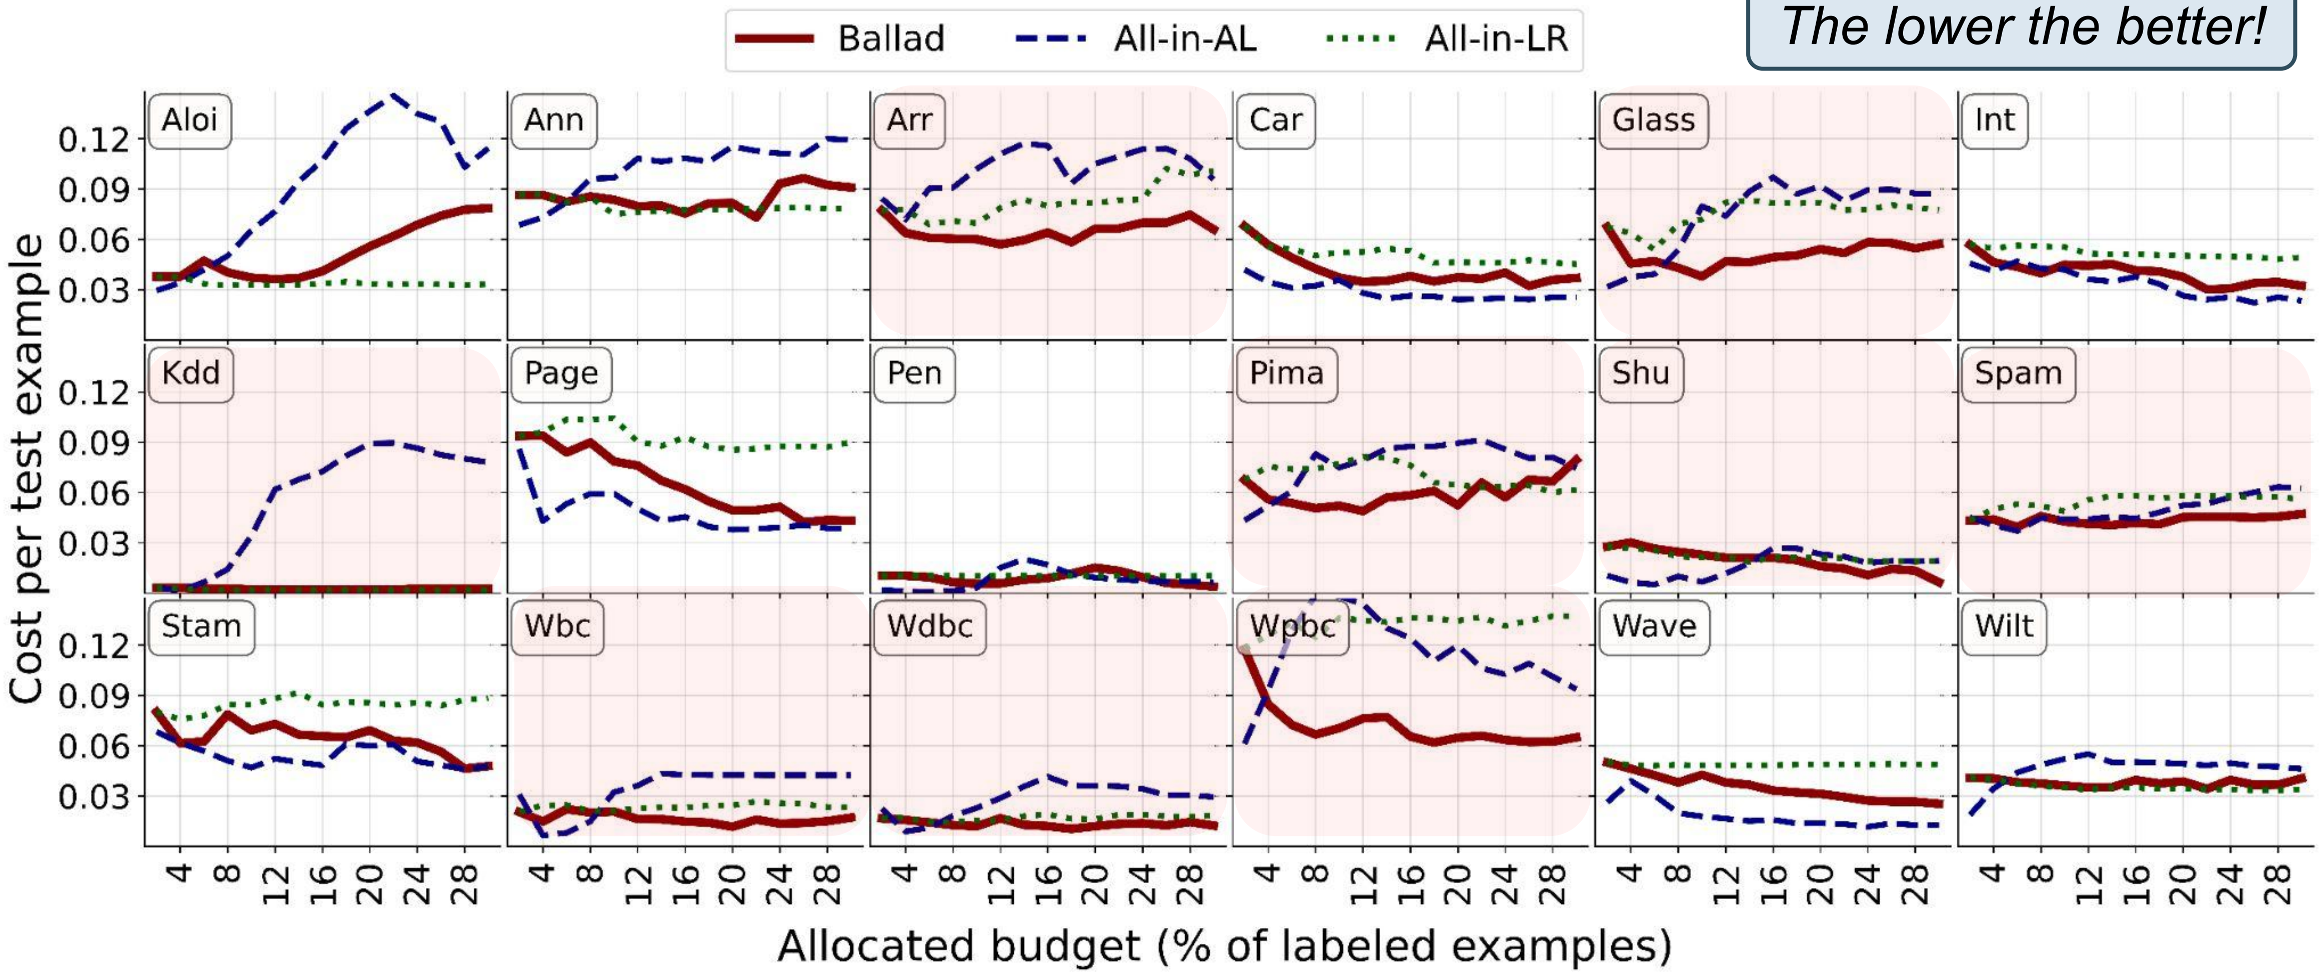

Q2. Which Reward Metric Is Better Between Entropy and Cosine Similarity ?

The lower the better!

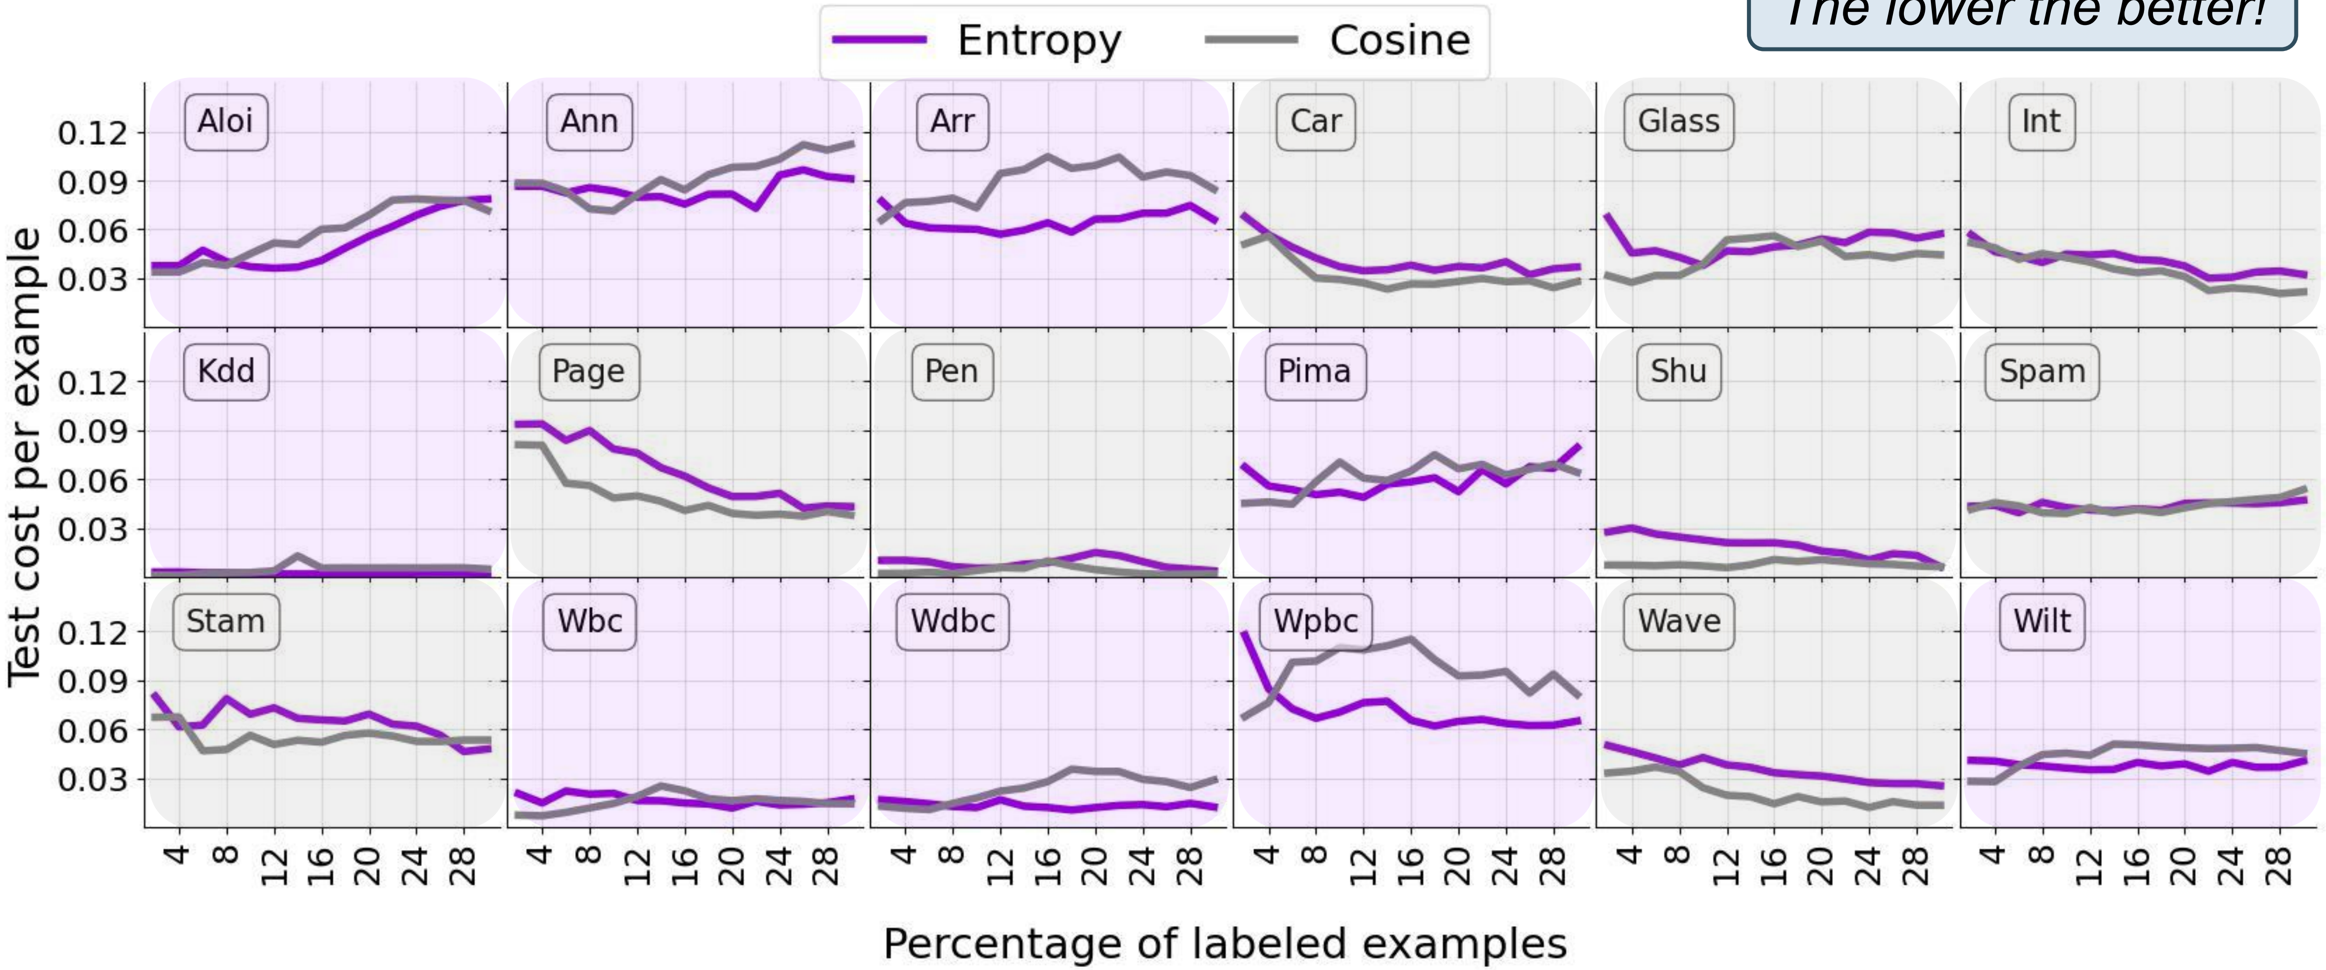

Q3. Is the reward function on a <u>similar scale</u> for AL and LR ?

We measure both rewards for each allocation round and compute the difference

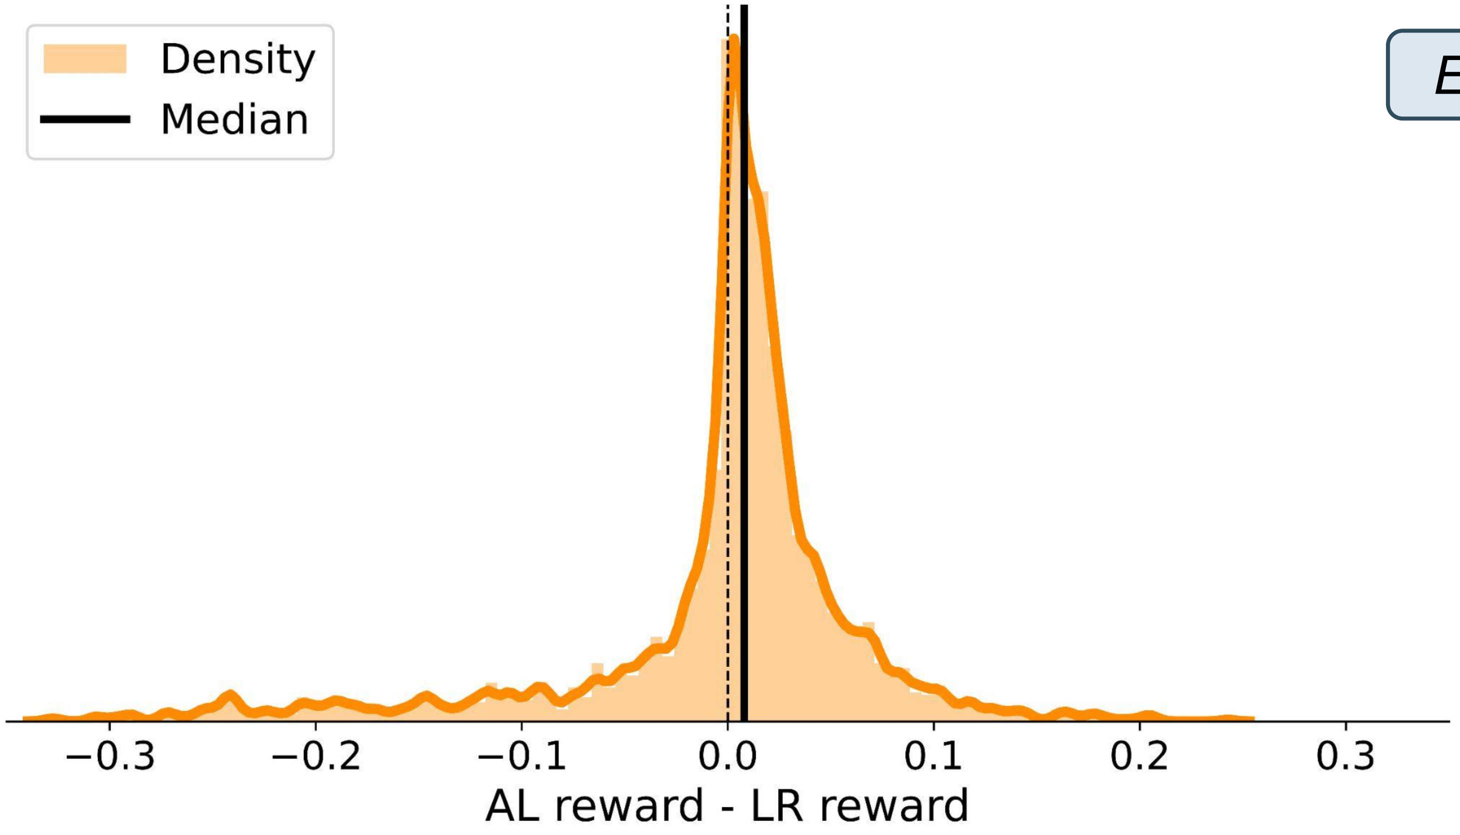

Entropy Reward

BALLAD: a reward-based strategy to decide how to allocate the label budget between AL and LR in Anomaly Detection

1. We introduced the challenge of how to allocate the labels (AL vs. LR);
2. We proposed BALLAD, which measures the reward of using the labels for either sides (AL, LR) and allocates the labels to the option with highest reward;
3. We investigated two types of reward functions: Entropy and Cosine Similarity;
4. Experimentally, we evaluated BALLAD on 18 benchmark datasets.

***Lorenzo Perini**, Daniele Giannuzzi, Jesse Davis*

lorenzo.perini@kuleuven.be

https://people.cs.kuleuven.be/~lorenzo.perini/

@LorenzoPerini95

1st AAAI Workshop on Uncertainty Reasoning and Quantification in Decision Making (UDM23)