Karlsruhe Institute of Technology

Validating One-Class Active Learning with User Studies – a Prototype and Open Challenges

Holger Trittenbach, Adrian Englhardt, and Klemens Böhm

3rd International Workshop on Interactive Adaptive Learning (IAL2019)

September 16, 2019 – Würzburg, Germany

User studies are necessary to validate the value proposition of active learning

Problem Statement

- **Active Learning:** Involves users in classifier training to improve classification results through feedback
- **Value proposition:** Obtain helpful information from users that is not yet contained in training data

➔ Currently no validation whether one can realize this value in an actual application

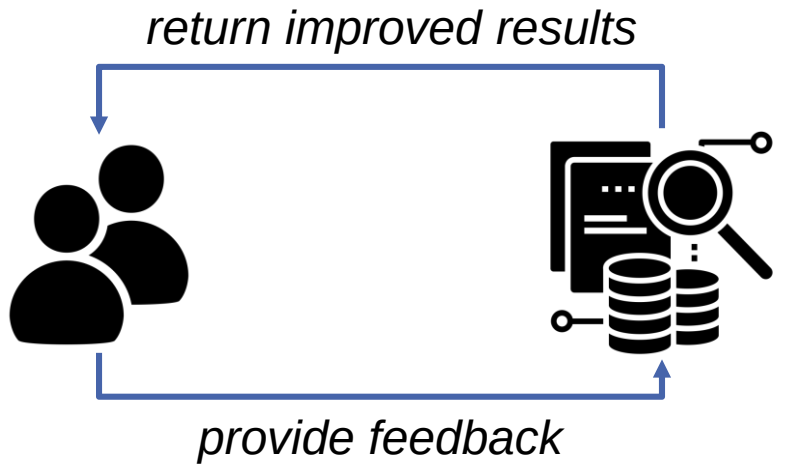

Our Contributions

- Architectural Sketch of an one-class active learning system
- Systematic overview of conceptual and technical challenges
- Roadmap towards validating active learning with user studies

Icons by Chanut is Industries from Noun Project

Our use case: High-Resolution Industrial Production Energy Data (HIPE)

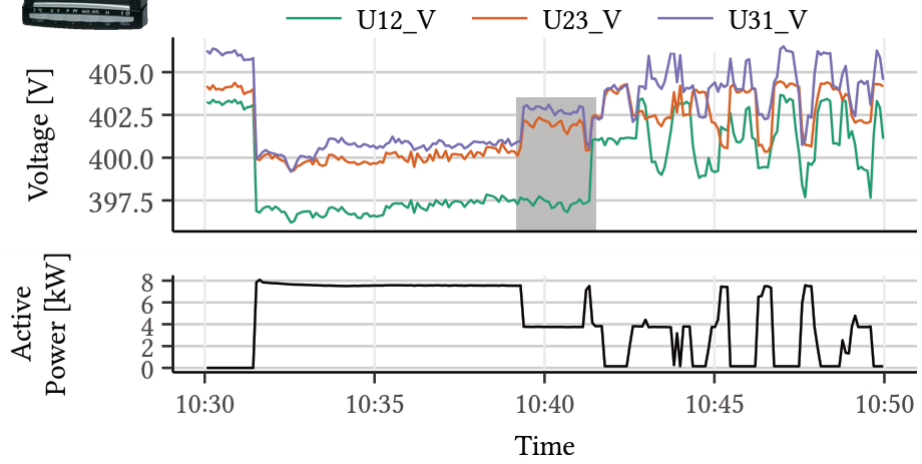

Smart-Meter Data Analysis

- Detect unusual sequences (Anomaly/Outlier Detection)
- Ground truth not available
- Learn from user feedback (Active Learning)

Source: S. Bischof, H. Trittenbach, M. Vollmer, D. Werle, T. Blank, and K. Böhm, “HIPE--An Energy-Status-Data Set from Industrial Production,” Proceedings of ACM e-Energy (e-Energy 2018)

Use Case: the pre-processing pipeline relies on features engineering

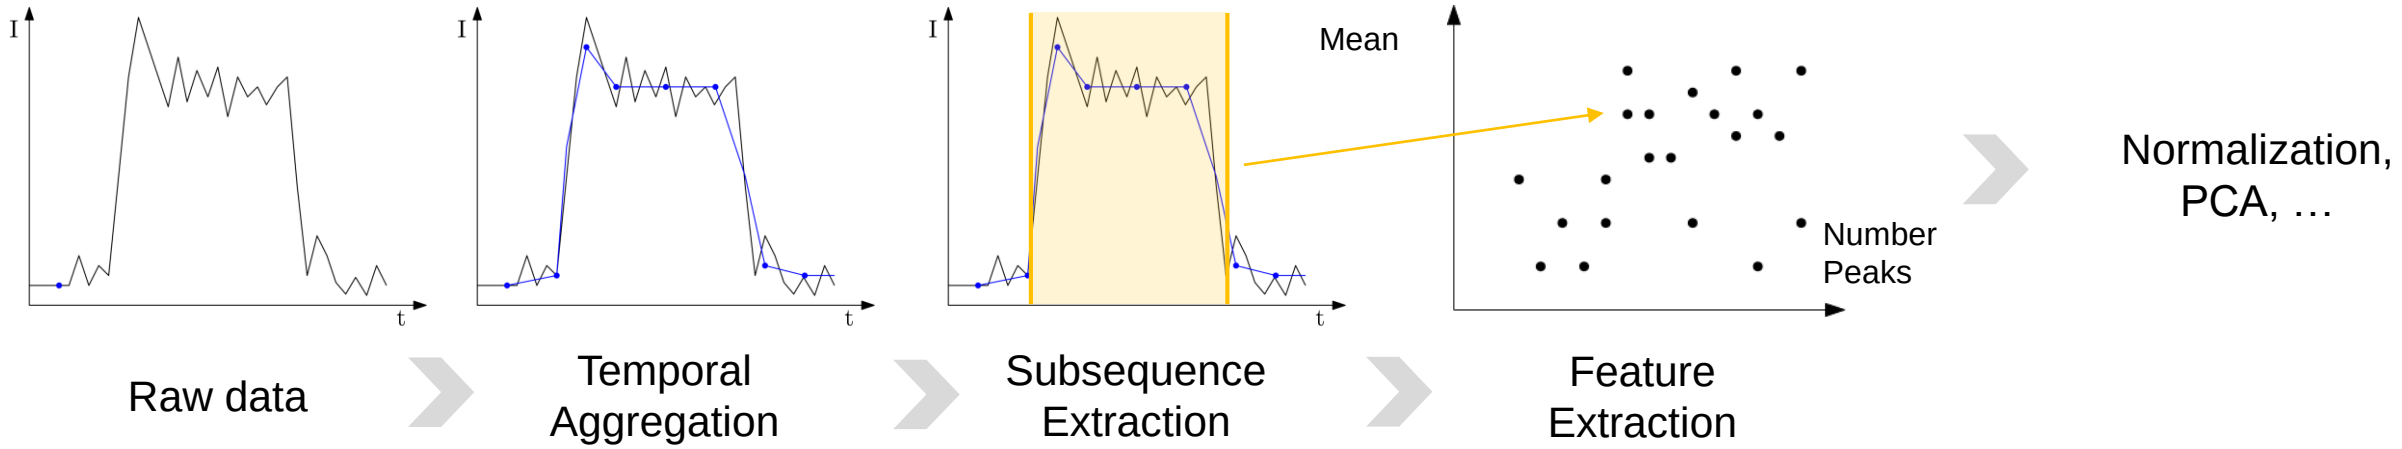

Source: M. Vollmer, A. Englhardt, H. Trittenbach, P. Bielski, S. Karrari, K. Böhm, "Energy time-series features for emerging applications on the basis of human-readable machine descriptions," Proceedings of ACM e-Energy (e-Energy 2019)

Active Learning for Outlier Detection asks users for class-label annotations

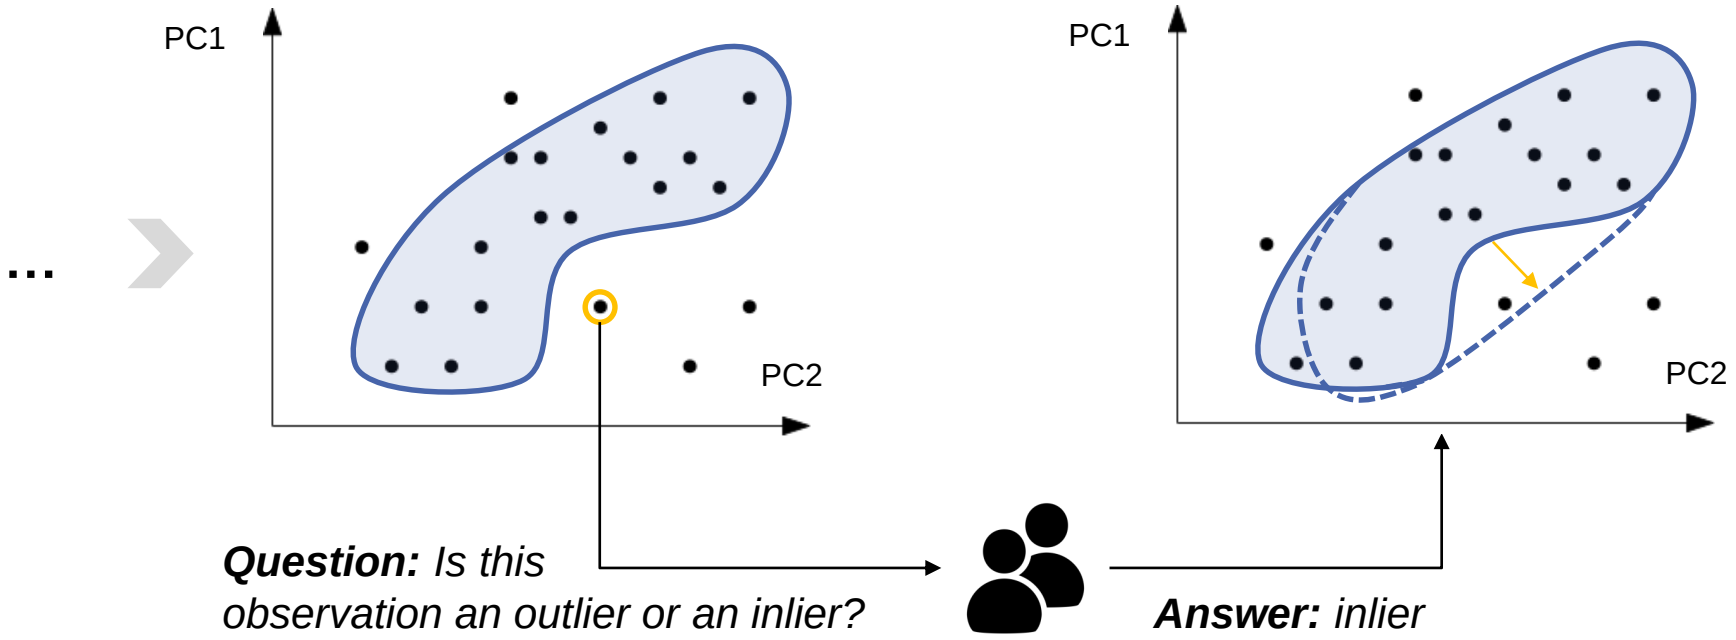

Fundamental Active Learning Assumptions

- **Feedback:** Users provide accurate feedback
- **Acceptance:** Users do not question how feedback changes the classifier

The fundamental active learning assumptions do not hold in practice

$$x_{42} = < 0.123, 14.23, -12.334, 0.01, 123, 42.42, 77, 55.31111, \dots, -3.1415 >$$

Question: Is this observation an outlier or an inlier?

Icon by Thomas Deckert from Noun Project

Our architectural sketch to evaluate one-class active learning with user studies

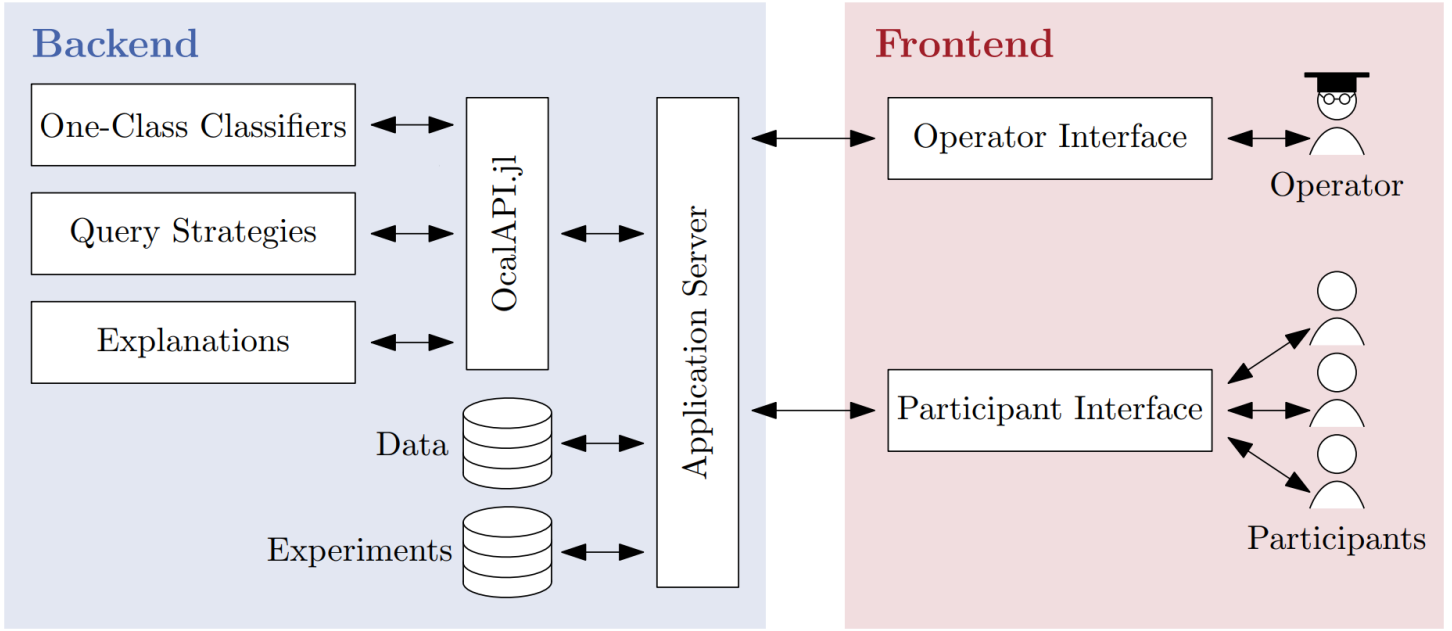

Challenges

- **Conceptual:** Different types of interaction, huge design space
- **Technical:** Open problems with realization of one-class active learning

The conceptual challenges complicate the design of OCAL systems

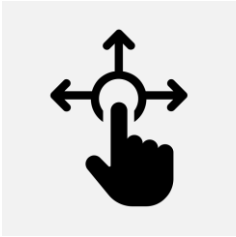

Type of Interaction

- System selects observations
- User selects observations *(VIAL)*

➔ Do insights from VIAL transfers to outlier detection?

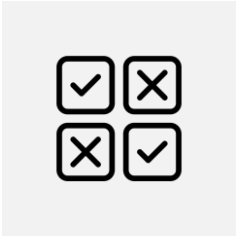

Type of Feedback

- Binary *(inlier/outlier)*
- Negative Feedback *(multi-class)*
- Feature-level

➔ Does OCAL benefit from different types of feedback?

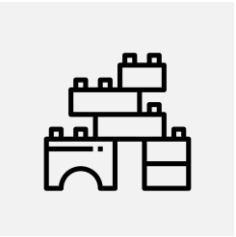

Design Space

- Learning Scenario *(sequential/batch, budget, assumptions on data distribution, initial setup)*
- Base Learner
- Query strategies
- Hyperparameters

➔ How to select a good configuration?

Preparation of Information

- Query presentation *(raw/feature space, context)*
- Result presentation *(contour plots, subspaces)**
- Black-Box Explanations *(local, interactive)*
- Contextual information *(feedback of other users, accuracy on labeled data)*

➔ Which information supports users to provide accurate feedback?

*) H. Trittenbach, K. Böhm, "One-Class Active Learning for Outlier Detection with Multiple Subspaces", CIKM 2019 (to appear)
Icons by Marksu Loritas, pongsakorn, Adrian Coquet, and Becris from Noun Project

There are several technical challenges in the way of implementing an OCAL System

Cold Start

- Applicability of query strategy and classifier

➔ Query strategy switching?

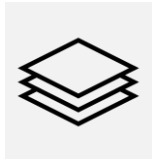

Batch Queries

- Avoid frequent retraining
- May support intuition

➔ New query selection schemes?

Runtime

- Interactivity not guaranteed

➔ Incremental Learning?

➔ Subsampling?

➔ Speculative execution?

Evaluation at Runtime

- Evaluate if feedback is good/sufficient
- Quality estimate may fluctuate with increasing feedback

➔ Query strategies to improve quality estimate?

Management of Data Flows

- Architectural challenge: where to retain data, classifier, predictions?
- Where do computations take place?

➔ Several trade-offs that are not well understood

Icons by Landan Lloyd, Nono Martinez Alonso, Creaticca Creative Agency, Creative Mahira and SBTS from Noun Project

Validating One-Class Active Learning with users studies as a step-wise process

Simplify Use Case

- Identify relatable domain *(e.g., image data)*
- Fix details *(e.g., no initial labels, sequential feedback)*
- Small data *(avoid scalability and data-flow difficulties)*

Validate Information Presented

- Find optimal classifier/hyperparameters/query strategy with ground truth
- Experiment with different query and result representations, explanations and contextual information

Implemented in our prototype

Validate Classifier/ Learning Strategy

- Fix presentations, vary classifier and query strategy
- Extensions *(e.g., batch queries)*

Generalize

- Increase number of observations *(scalability)*
- Different data sets

Conclusions

Problem Statement

- Currently no validation whether one can realize the value proposition of active learning in an actual application
- There are several challenges in the design and implementation of an interactive learning system

Contributions

- Systematic overview of conceptual and technical challenges
- Architectural Sketch of an one-class active learning system
- Prototype implementation
- Roadmap towards validating active learning with user studies

Gefördert durch DFG Deutsche Forschungsgemeinschaft

https://github.com/englhardt/SVDD.jl
https://github.com/englhardt/OneClassActiveLearning.jl
https://github.com/englhardt/OcalAPI.jl