

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

Holger Trittenbach, Adrian Englhardt, and Klemens Böhm

Background and Motivation



Smart Meter Data Analysis

- Detect unusual sequences (anomaly/outlier detection)
- Ground truth not available
- ➔ Learn from feedback (One-Class Active Learning OCAL)


Problem Statement

- Value proposition: AL obtains information from users
- ➔ Currently no validation in an actual application

Contributions

- Architectural Sketch of an OCAL system
- Overview of conceptual and technical challenges
- Roadmap towards validating AL with user studies

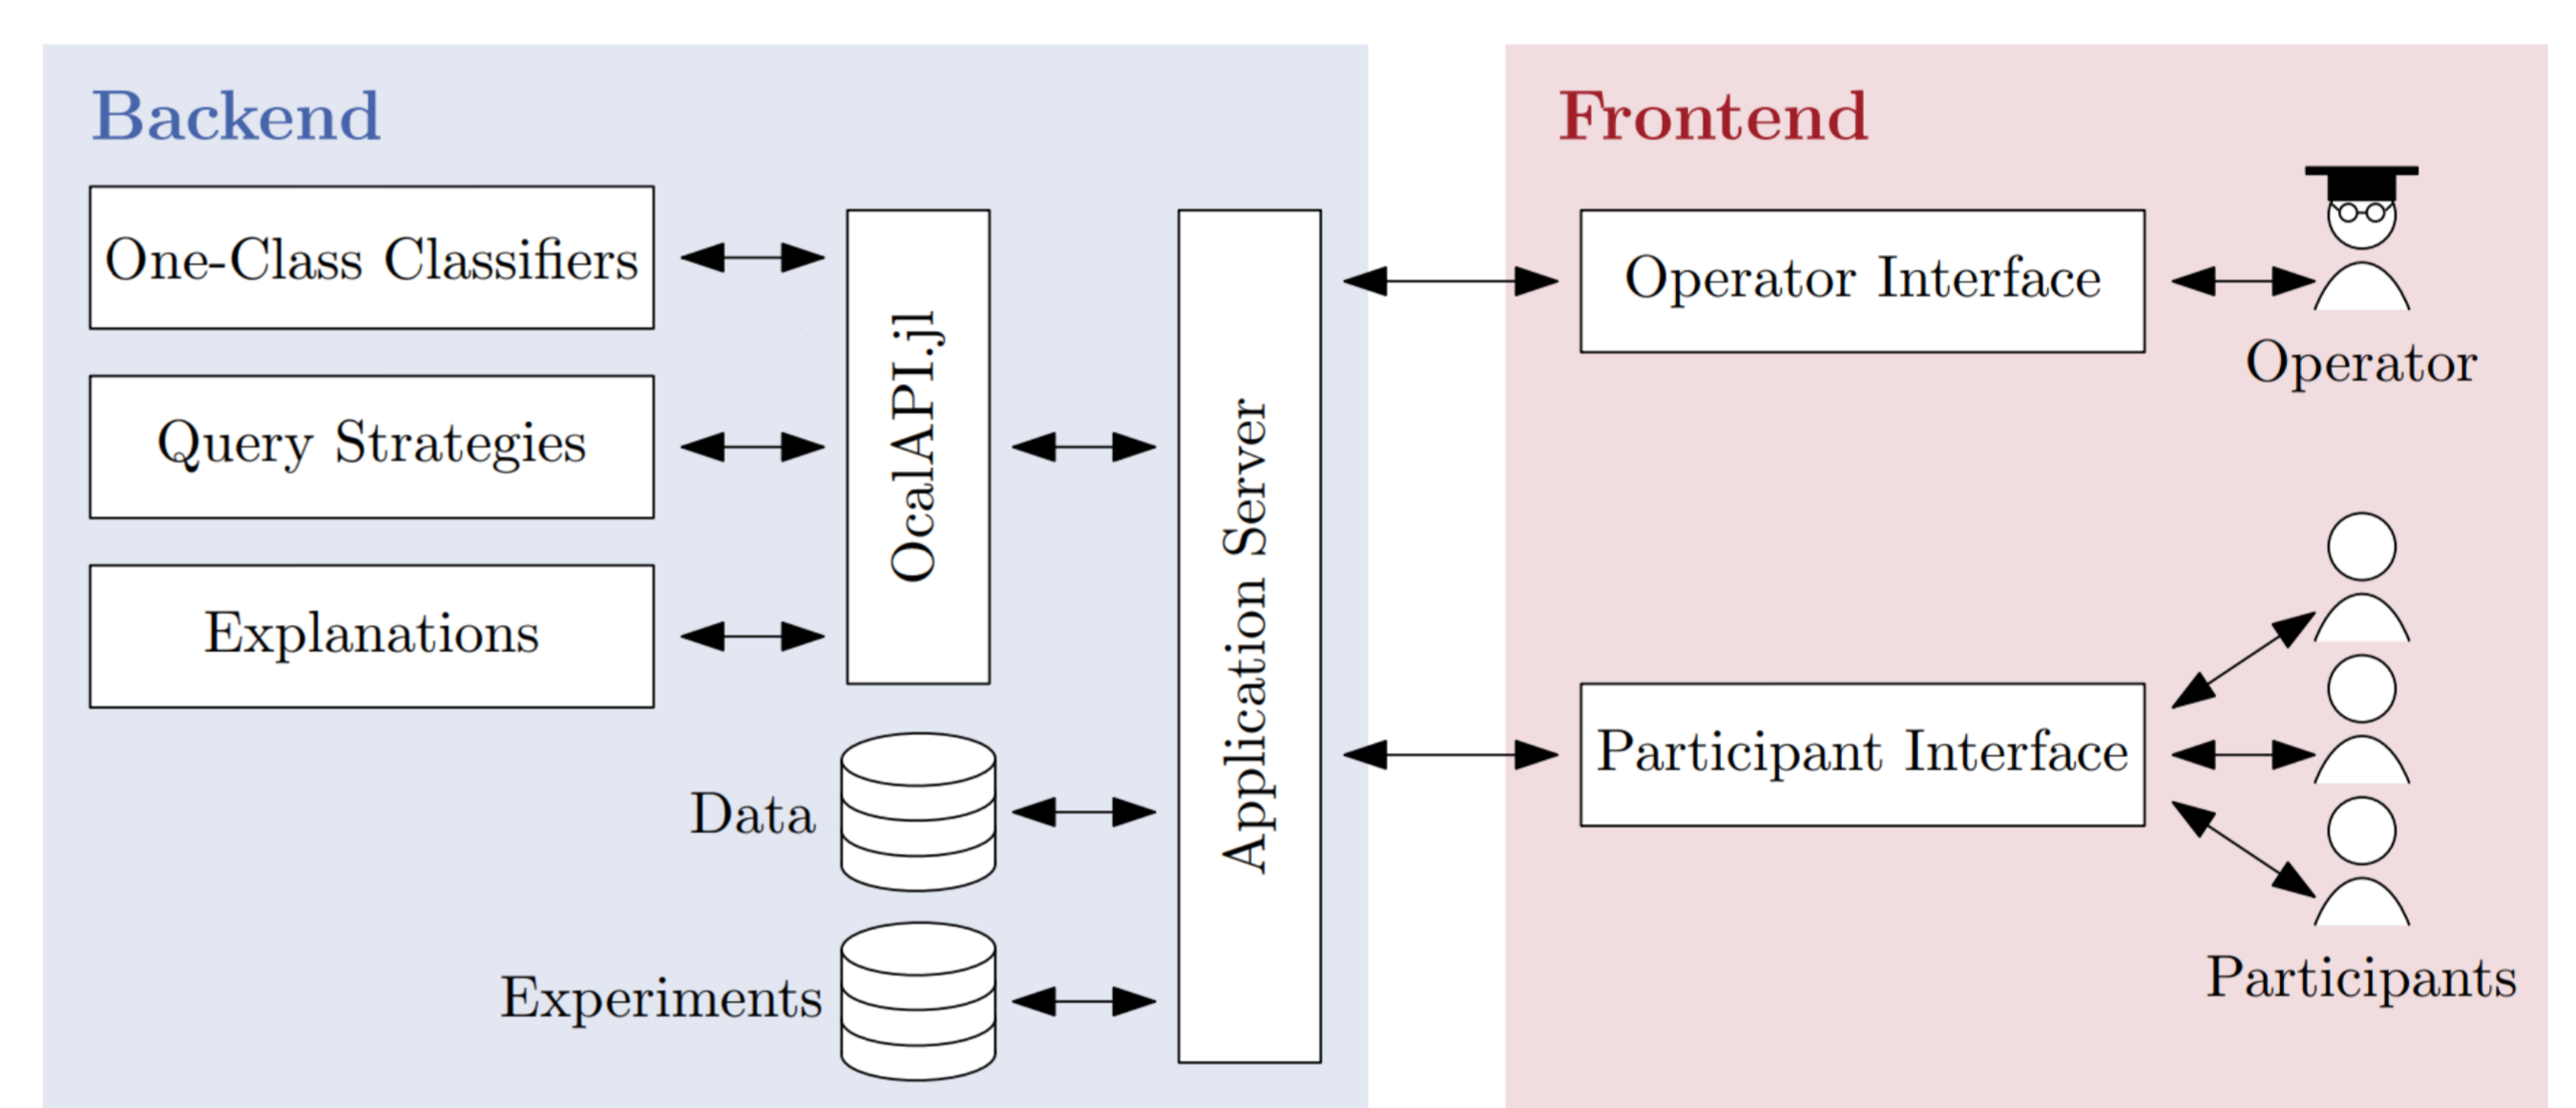
Scenario



Challenges

- **Conceptual:** Type of interaction, huge design space
- **Technical:** Open problems with realization of OCAL

Architectural Sketch



Conceptual Challenges

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)
- Base Learner
- Query strategies
- Hyperparameters
- ➔ How to select a good configuration?

Preparation of Information

- Query presentation (raw/features)
- Result presentation (subspaces)
- Black-Box Explanations
- Contextual information (feedback of other users)
- ➔ Which information supports users to provide accurate feedback?

Technical Challenges

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?
- ➔ 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

References

- [1] S. Bischof et al., “HIPE – An Energy-Status-Data Set Industrial Production,” (e-Energy 2018)
[2] M. Vollmer et al., “Energy time-series features for emerging applications on the basis of human-readable machine descriptions,” (e-Energy 2019)