

**Title Here:** Ask-the-Expert: Minimizing human review for big data analytics through active learning

**PI:** Nikunj C. Oza, Ph.D.

**Co-Is:** Kamalika Das, Ph.D., Bryan Matthews

**Activity Type/Phase:** New Start

**Primary TA:** 11.4.2

**Start TRL:** 3

**End TRL:** 4

### **Goal/Gap**

In this CIF project, we worked toward semi-automating knowledge discovery from anomaly detection algorithms through the use of active learning. Active learning is an area of research within machine learning that uses an “expert in the loop” to learn from large data sets that have very few annotations or labels available, and where providing such labels is expensive. In our case, the task can be defined as the identification of safety events from flight operational data. Since traditional anomaly detection algorithms cannot differentiate between operationally relevant and irrelevant statistical anomalies, Subject Matter Experts (SMEs) have a lengthy and expensive burden of investigating every example identified by the detection algorithm, classifying and labeling them as relevant or irrelevant. Active learning identifies the unlabeled example for which a label would most improve the classifier, asks the domain expert for a label, and repeats this process until there are no more resources (time, budget) available for labeling or a minimum required performance is reached. A positive label indicates an operationally significant safety event whereas a negative label indicates otherwise. Based on these few labels we propose to build an active learning system that utilizes the SME's time in the most effective manner by iteratively asking for labels for as few informative instances as possible. Our work was proposed to be a stepping stone toward implementation and deployment of the system with user interface to be pursued by the Aviation Operations and Safety Program (AOSP) given its interest in safety monitoring and discovery of safety incidents.

### **Approach/Innovation**

Active learning has been an area of research in the machine learning community for almost a decade. Although new algorithms are regularly being proposed for active learning in major machine learning conferences and journals and has been a favorite among researchers in academia, it has not seen much adoption by the industry (practitioners). A primary reason is that, due to its iterative nature, it does not scale gracefully with the growing size of data sets that most real applications deal with. To overcome this weakness, we do not plan to use active learning on our aviation data sets directly, but in conjunction with other detection algorithms. Since, in our proposed approach, active learning takes place on the output of the detection algorithm, it has to handle less than 1% of the original data set. In the machine learning literature, this approach is new. Also, in the detection literature the state-

of-the-art is manual review of results by SME. Active learning has never been used for this purpose.

The key technical challenge for the proposed work was the incorporation of SME feedback in the classification algorithm. During active learning when the SME identifies an example to be operationally relevant, he also provides his rationale behind the classification. If the rationale can be summarized in a machine-understandable way, then it can be used to improve classifier accuracy much faster and with very few labeled instances. In this work, we developed a method to translate and incorporate SME rationales along with labels for incorporating into the active learning framework.

### **Results/Knowledge Gained**

Tell what was planned as well as what was actually accomplished. Describe the outcome and knowledge gained (this includes lessons learned). Insert or append any images or charts that add context to the results. (1 to 3 paragraphs)

The goal of this work was 2-fold: (i) significantly reduce the SME review time for anomaly detection through the use of active learning for differentiating between operationally significant and not operationally significant (only statistically) anomalous results found by an unsupervised anomaly detection algorithm, (ii) build a software interface that easily allows us to run this active learning system iteratively through the cloud to facilitate expert involvement in the labeling process.

For accomplishing the first goal, we developed an active learning algorithm that incorporates SME feedback in the form of rationales to build a classifier that can distinguish between uninteresting and operationally significant anomalies. In this process, first an unsupervised anomaly detection algorithm is run on all the flight data to generate a ranked list of statistically significant anomalies. A very small percentage of these are presented to SMEs to bootstrap the active learning process. The SME provides labels for each of these instances along with an explanation about the label. A positive label indicates an operationally significant safety event whereas a negative label indicates otherwise. Based on these few labels we build an active learning system that (i) utilizes the SME's time in the most effective manner by iteratively asking for labels for few of the most informative instances, (ii) elicits rationales/explanations from the SME for why s/he assigns a certain label to an instance, and (iii) constructs new features, based on rationales, that are incorporated in future iterations of active learning and classifier training. We evaluate all strategies using precision @k measure which can be defined as the number of positive instances in top k instances ranked by the classifier. This measure is most suitable for our application because the SMEs go through a list of anomalies to identify those that are operationally significant for further investigation, and improving precision @k means that

the SMEs would analyze more of the OS flights compared to the NOS flights. We chose precision @5 for evaluation. We bootstrap the classifier using an initially labeled set containing one OS flight and one NOS flight, and at each round of active learning the learner picks a new flight for labeling. We evaluate our strategy using 2-fold cross validation and repeat each experiment 10 times per fold starting with a different bootstrap, and present average results over 20 different runs. We set the budget (B) in our experiments to 45 flights.


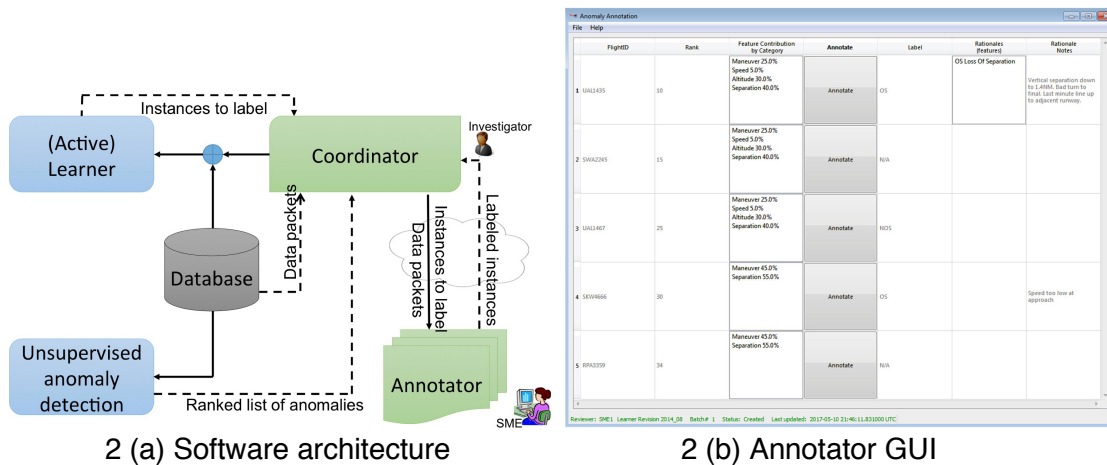


Figure 1. MLP w/Rationales (our algorithm). improves performance over MLP (basic active learning strategy that does not use explanations), random sampling, and ranking produced by unsupervised anomaly detection (MKAD) for precision@5

In the absence of our active learning framework, our SMEs took approximately 33 hours to review the entire set of 153 anomalies produced by MKAD. These 33 hours were spread over multiple weeks due to limited availability of SME time for such tasks, which is a standard problem in the industry. As Fig. 1 shows, the learning curve flattens out after labeling 35 flights. This would reduce the SME review time to less than one-third of the original time. This has implications on both man-hours and monetary savings. Moreover, basic active learning strategies could only achieve precision @5 of 0.57 whereas our method of MLP w/Rationales achieves precision @5 of 1 (75: 4% improvement over state-of-the-art)

As part of this project we also built an annotation interface is to facilitate review of a set of anomalies detected by an unsupervised anomaly detection algorithm and allow labeling of those anomalies as either operationally significant (OS) or not operationally significant (NOS). Our system, as shown in Figure 2a consists of two components, viz. the coordinator and the annotator.



The coordinator has access to the data repository and accepts inputs in the form of a ranked list of anomalies from the unsupervised anomaly detection algorithm. The coordinator is the backbone of the system communicating iteratively with the active learner, gathering information on instances selected for annotation and packing information for transmission to the annotator. Once the annotator collects and sends the labeled instances, the coordinator performs two tasks: (i) resolve labeling conflicts across multiple SMEs through the use of a majority voting scheme or by invoking an investigator review, and (ii) automate the construction of new rationale features as conjunctions and/or disjunctions of raw data features based on the rationale notes entered by the SME in the annotation window. All data exchange between the coordinator and the annotator happens through cloud based storage. The annotator, shown in Figure 2b is the graphical user interface that the SMEs work with and needs to be installed at the SME end. When the annotator is opened, it checks for new data packets (to be labeled) on the cloud. If new examples need annotation, the annotator window displays the list of examples ranked in the order of importance along with the features identified to be the most anomalous. Clicking on the annotate button next to each example, the SME can delve deeper into that example in order to provide a label for that instance.

Publication: Kamalika Das, Ilya Avrekh, Bryan Matthews, Manali Sharma, and Nikunj Oza, ASK-the-Expert: Active Learning Based Knowledge Discovery Using the Expert, *Proceedings of the European Conference on Machine Learning*, 2017.

## Technology Maturation Opportunities

Ricky Howard suggested that we work with Patricia Parsons-Wingerter, Jennifer Fogarty, and Molly Anderson to see if this work may be useful in the context of VESGEN 2D's identification of vascular structure. I discussed this with Patricia and she feels that this work would be useful. Members of my CIF Ask-the-Expert team, along with a member of Patricia's VESGEN CIF team, are preparing a brief proposal that we will send to Patricia, and then to Jennifer and Molly. We respectfully ask CIF to consider funding a CIF-style award for this to-be-proposed work.

We have proposed to directly extend the work of the current CIF within the Aviation Operations and Safety Program's (AOSP) System-Wide Safety (SWS) project with the ultimate goal of deployment at the FAA, MITRE, and others who perform aviation safety analysis.

### **Partnerships**

Our collaborator, Prof. Mustafa Bilgic, is an active researcher in the area of active learning. His former Ph.D. student, Manali Sharma, did her dissertation in the area of active learning and, as an intern at NASA Ames, did some algorithmic work that went into this CIF award. This CIF represents a significant step toward operationalizing this algorithmic work.