

ASK-the-Expert: Active learning based knowledge discovery using the expert

Kamalika Das¹, Ilya Avrekh², Bryan Matthews²
Manali Sharma³, and Nikunj Oza⁴

¹ USRA, NASA Ames Research Center, Moffett Field, CA, USA

² SGT Inc., NASA Ames Research Center, Moffett Field, CA, USA

³ Samsung Semiconductor Inc., San Jose, CA, USA**

⁴ NASA Ames Research Center, Moffett Field, CA, USA

{kamalika.das, ilya.avrekh-1, bryan.l.matthews, nikunj.c.oza}@nasa.gov
manali.s@samsung.com

Abstract. Often the manual review of large data sets, either for purposes of labeling unlabeled instances or for classifying meaningful results from uninteresting (but statistically significant) ones is extremely resource intensive, especially in terms of subject matter expert (SME) time. Use of active learning has been shown to diminish this review time significantly. However, since active learning is an iterative process of learning a classifier based on a small number of SME-provided labels at each iteration, the lack of an enabling tool can hinder the process of adoption of these technologies in real-life, in spite of their labor-saving potential. In this demo we present ASK-the-Expert, an interactive tool that allows SMEs to review instances from a data set and provide labels within a single framework. ASK-the-Expert is powered by an active learning algorithm for training a classifier in the backend. We demonstrate this system in the context of an aviation safety application, but the tool can be adopted to work as a simple review and labeling tool as well, without the use of active learning.

Keywords: active learning, graphical user interface, review, labeling

1 Introduction

Active learning is an iterative process that requires feedback on instances from a subject matter expert (SME) in an interactive fashion. The idea in active learning is to bootstrap an initial classifier with a few examples from each class that have been labeled by the SME. Traditional active learning approaches select an informative instance from the unlabeled data set and ask SMEs to review the instance and provide a label. This process continues iteratively until a desired level of performance is achieved by the classifier or when the budget (allotted resources) for the SME is exhausted. Much of the research in active learning simulates this interaction between the learner and the SME. In particular, all

** This work was done when the author was a student at Illinois Institute of Technology

labels are collected from the SME a priori and during the active learning process, the relevant labeled instances are revealed to the learner, based on its requests at each iteration. The problem of using such retrospective evaluation of an active learning algorithm is twofold. Firstly, the lack of availability of an interactive interface is largely responsible for the generally low adoption of active learning algorithms in practical scenarios. Secondly, the simulated environment fails to achieve the biggest benefit associated with the use of active learning: reduction of SME review time. This is because the SME has to review and label all examples a priori. Therefore, for utilizing active learning frameworks in situations of low availability of labeled data, it is important to have an interactive tool that allows SMEs to review and label instances only when asked by the learner.

2 Application and demo scenario

A major focus of the commercial aviation community is discovery of unknown safety events in flight operational data through the use of unsupervised anomaly detection algorithms. However, anomalies found using such approaches are abnormal only in the statistical sense, i.e., they may or may not represent an operationally significant event (e.g. represent a real safety concern). After an algorithm produces a list of statistical anomalies, an SME must review the list to identify those that are operationally relevant for further investigation. Usually, less than 1% of the hundreds or thousands of statistical anomalies turn out to be operationally relevant. Therefore, substantial time and effort is spent examining anomalies that are not of interest and it is essential to optimize this review process in order to reduce SME labeling efforts (man hours spent in investigating results). A recently developed active learning method [2] incorporates SME feedback in the form of rationales for classification of flights to build a classifier that can distinguish between uninteresting and operationally significant anomalies with 70% fewer labels compared to manual review and comparable accuracy.

To the best of our knowledge, there exists no published work that describes such software tools for review and annotation of numerical data using active learning. There are some image and video annotation tools that collect labels, such as LabelMe from MIT CSAIL [1]. Additionally there are active learning powered text labeling tools, such as Abstrackr [3] designed specifically for medical experts for citation review and labeling. The major difference between these annotator tools and our tool is the absence of context in our case. Unlike in the case of image or text data where the information is self-contained in the instance being reviewed, in our case, we have to enable the tool to obtain additional contextual information and visualize the feature space on demand. Other domains plagued by label scarcity can also benefit from the adaptation of this tool, with or without the use of an active learning algorithm.

3 System description

In this demo the goal of our 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 1a consists of two components, viz. the coordinator and the annotator.

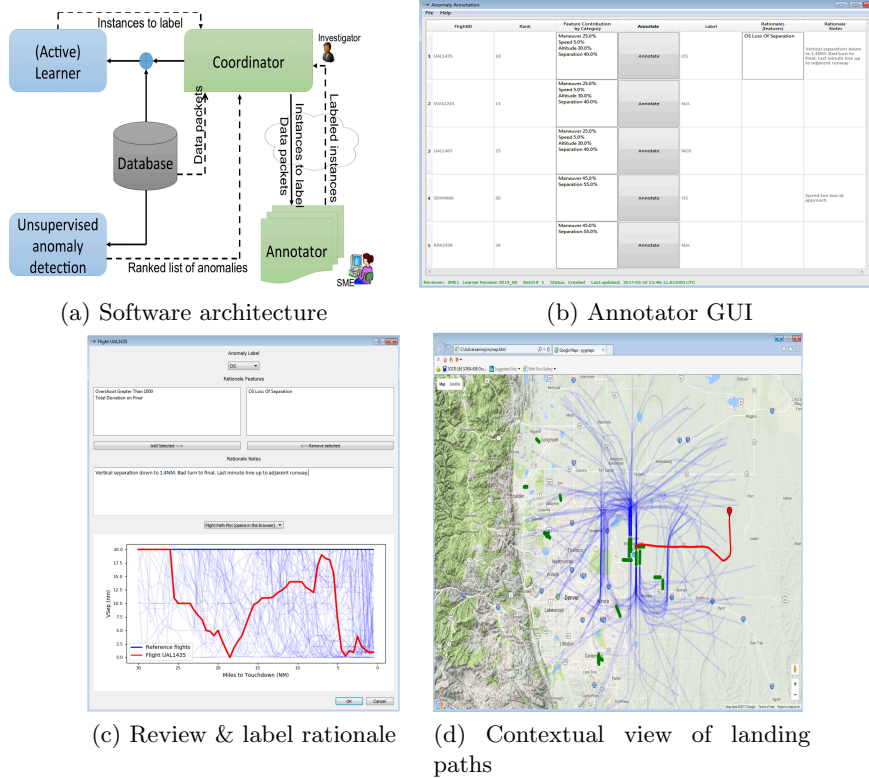


Fig. 1: Software architecture and snapshots of ASK-the-Expert

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 annota-

tor happens through cloud based storage. The annotator, shown in Figure 1b 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. The functions of the annotator include (i) obtaining examples to be labeled from the cloud and displaying them to the SME, (ii) allowing review of individual features as well as feature interactions (shown in Figure 1c), and (iii) occasionally providing additional context information by looking at additional data sources (for example, plotting flight paths in the context of other flights landing on the same runway at a certain airport using geographical data from maps, as shown in Figure 1d). Multiple annotators can be used simultaneously by different SMEs to label the same or different sets of examples. Once the labeled examples are submitted by the annotator, the coordinator collects and consolidates them and sends them back to the learner.

Demo plan: We will demonstrate the ASK-the-Expert tool for an aviation safety case study. Since the data cubes for normal and anomalous flights are proprietary information, the database will be hosted in our laptop. The coordinator tool will be live and running at NASA, gathering the latest set of flights that need to be labeled and uploading them on the cloud. We will demonstrate how the SMEs can review new examples in the context of other flights and provide labels. Their feedback will be sent back to the learner through the coordinator for the next iteration of classifier learning after incorporating new rationale features.

Acknowledgments. This work is supported in part by a Center Innovation Fund (CIF) 2017 grant at NASA Ames Research Center and in part by the NASA Aeronautics Mission Directorate. Manali Sharma was supported by the National Science Foundation CAREER award no. IIS-1350337.

References

1. Russell, B., Torralba, A., Murphy, K., Freeman, W.: Labelme: a database and web-based tool for image annotation. *Inter. J. of Comp. Vision* 77(1), 157–173 (2007)
2. Sharma, M., Das, K., Bilgic, M., Matthews, B., Nielsen, D., Oza, N.: Active learning with rationales for identifying operationally significant anomalies in aviation. In: *Proc. of European Conference on Machine Learning and Knowledge Discovery in Databases*. pp. 209–225. *ECML-PKDD 2016* (2016)
3. Wallace, B., Small, K., Brodley, C., Lau, J., Trikalinos, T.: Deploying an interactive machine learning system in an evidence-based practice center: Abstrackr. In: *Proc. of the 2nd ACM SIGHIT Int. Health Informatics Symposium*. pp. 819–824 (2012)