Demo Abstract: Acoustic Anomaly Detection System

Jae-wook Ahn, Keith Grueneberg, Bong Jun Ko, Wei-Han Lee, Eduardo Morales, Shiqiang Wang, Xiping Wang, David Wood*
{jaewook.ahn,kgruen,bko,wangshiq,xiping,dawood}@us.ibm.com,{wei-han.lee1,eduardo.morales}@ibm.com
IBM T. J. Watson Research Center, Yorktown Heights, NY, USA

ABSTRACT

Acoustic signals contain rich information of the environment. They can be used for detecting anomalous events such as in automated machine monitoring. In this demonstration, we present our acoustic anomaly detection system that captures acoustic signals and classifies them using machine learning techniques. Our system includes a server for sound management and model training, a mobile client for sound capturing and real-time classification, and a workbench that acts as a user interface. We will show the full operational pipeline of our system in this demonstration.

CCS CONCEPTS

• **Computing methodologies** → **Supervised learning by classification**; • **Computer systems organization** → **Client-server architectures**.

KEYWORDS

Acoustic signal analytics, anomaly detection, machine learning

ACM Reference Format:
Jae-wook Ahn, Keith Grueneberg, Bong Jun Ko, Wei-Han Lee, Eduardo Morales, Shiqiang Wang, Xiping Wang, David Wood. 2019. Demo Abstract: Acoustic Anomaly Detection System. In *The 17th ACM Conference on Embedded Networked Sensor Systems (SenSys '19), November 10–13, 2019, New York, NY, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3356250.3361963

1 INTRODUCTION

Acoustic signals carry rich information of our environment. They can be used to detect environmental events through acoustic signal analytics. For example, when a washing machine starts making a rattling sound, it may stop working soon. If a car makes a bad-bearing sound on a wheel, there is something wrong on it. If we use acoustic sensors to automatically monitor the sound of equipment and detect functional anomalies, maintenance services could be in place as soon as the equipment shows an unusual sign, thus preventing the damage or loss of valuable equipment. In another scenario, patients with different medical conditions, such as pneumonia, bronchiolitis, and asthma make different coughing sounds.

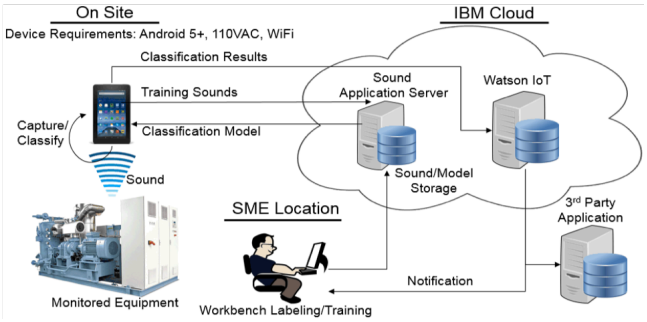

Figure 1: System Architecture.

By analyzing the coughing sounds, we could separate pneumonia from other diseases and hence provide patients with proper treatment in a timely manner.

In recent years, mobile sensing is a very active area of research [1, 2]. With the increase of mobile device's computing power, mobile acoustic sensing, analytics, and sense-making have been proven as a promising approach for automated machine monitoring for various applications. We have designed and built an acoustic monitoring system using machine learning techniques that can operate on mobile devices, aiming at enabling automated machine monitoring for different environments at a large scale. The system is fast enough to process large volumes of acoustic data and can be applied to a broad range of applications, such as home appliance anomaly detection, vehicle fault diagnosis, manufacturing equipment failure detection, building maintenance, health care, etc. In this demonstration, we present the architecture of our system and a live demonstration of ambient sound detection and classification using mobile phones. We also demonstrate how sounds are labeled, how acoustic models are trained using machine learning techniques, and how the trained models are tuned to improve classification accuracy.

2 SYSTEM ARCHITECTURE

The system architecture is shown in Fig. 1. A mobile client connects to a cloud-based server that provides functionalities for storing labeled sounds and trained models, as well as a workbench (web-based user interface) for managing the sounds and models. The mobile client is equipped with a software to perform sound capturing, sound labeling, communication with server, on-device feature extraction, and on-device classification. The server performs model training and (server-side) sound classification. The mobile client can perform local on-device classification by retrieving a trained model from the server. After acquiring a model, the client can perform continuous monitoring using either a local model or a server-based model. Optionally, the client can publish the classification results to a Watson IoT instance[1] to enable custom notification and alerting when an anomaly is detected.

*Authors listed alphabetically by last name.

SenSys '19, November 10–13, 2019, New York, NY, USA

ACM ISBN 978-1-4503-6950-3/19/11…$15.00
https://doi.org/10.1145/3356250.3361963

[1] https://developer.ibm.com/tutorials/cl-mqtt-bluemix-iot-node-red-app/