

Sound, vibration recognition boost context-aware computing

15 October 2018



Carnegie Mellon University researchers have developed a method that enable smart devices to figure out where they are and what people are doing around them by analyzing sounds from their microphones Credit: Carnegie Mellon University


Smart devices can seem dumb if they don't understand where they are or what people around them are doing. Carnegie Mellon University researchers say this environmental awareness can be enhanced by complementary methods for analyzing sound and vibrations.

"A smart speaker sitting on a kitchen countertop cannot figure out if it is in a kitchen, let alone know what a person is doing in a kitchen," said Chris Harrison, assistant professor in CMU's Human-Computer Interaction Institute (HCII). "But if these devices understood what was happening around them, they could be much more helpful."

Harrison and colleagues in the Future Interfaces Group will report today at the Association for Computing Machinery's User Interface Software and Technology Symposium in Berlin about two approaches to this problem—one that uses the most ubiquitous of sensors, the microphone, and another that employs a modern-day version of eavesdropping technology used by the KGB in the

1950s.

In the first case, the researchers have sought to develop a sound-based activity recognition system, called Ubioustics. This system would use the existing microphones in smart speakers, smartphones and smartwatches, enabling them to recognize sounds associated with places, such as bedrooms, kitchens, workshops, entrances and offices.



Carnegie Mellon University researchers are using laser vibrometry -- a method similar to one once used by the KGB for eavesdropping -- to monitor vibrations and movements of objects, enabling smart devices to be aware of human activity. Credit: Carnegie Mellon University

"The main idea here is to leverage the professional sound-effect libraries typically used in the entertainment industry," said Gierad Laput, a Ph.D. student in HCII. "They are clean, properly labeled, well-segmented and diverse. Plus, we can transform and project them into hundreds of different variations, creating volumes of data perfect for training deep-learning models.

"This system could be deployed to an existing device as a software update and work

immediately," he added.

The plug-and-play system could work in any environment. It could alert the user when someone knocks on the front door, for instance, or move to the next step in a recipe when it detects an activity, such as running a blender or chopping.

The researchers, including Karan Ahuja, a Ph.D. student in HCII, and Mayank Goel, assistant professor in the Institute for Software Research, began with an existing model for labeling sounds and tuned it using sound effects from the professional libraries, such as kitchen appliances, power tools, hair dryers, keyboards and other context-specific sounds. They then synthetically altered the sounds to create hundreds of variations.

Laput said recognizing sounds and placing them in the correct context is challenging, in part because multiple sounds are often present and can interfere with each other. In their tests, Ubioustics had an accuracy of about 80 percent—competitive with human accuracy, but not yet good enough to support user applications. Better microphones, higher sampling rates and different model architectures all might increase accuracy with further research.

A video explaining Ubioustics:

In a separate paper, HCII Ph.D. student Yang Zhang, along with Laput and Harrison, describe what they call Vibrosight, which can detect vibrations in specific locations in a room using laser vibrometry. It is similar to the light-based devices the KGB once used to detect vibrations on reflective surfaces such as windows, allowing them to listen in on the conversations that generated the vibrations.

"The cool thing about vibration is that it is a byproduct of most human activity," Zhang said. Running on a treadmill, pounding a hammer or typing on a keyboard all create vibrations that can be detected at a distance. "The other cool thing is that vibrations are localized to a surface," he added. Unlike microphones, the vibrations of one activity don't interfere with vibrations from another. And unlike microphones and cameras, monitoring

vibrations in specific locations makes this technique discreet and preserves privacy.

This method does require a special sensor, a low-power laser combined with a motorized, steerable mirror. The researchers built their experimental device for about \$80. Reflective tags—the same material used to make bikes and pedestrians more visible at night—are applied to the objects to be monitored. The sensor can be mounted in a corner of a room and can monitor vibrations for multiple objects.

Zhang said the sensor can detect whether a device is on or off with 98 percent accuracy and identify the device with 92 percent accuracy, based on the object's [vibration](#) profile. It can also detect movement, such as that of a chair when someone sits in it, and it knows when someone has blocked the sensor's view of a tag, such as when someone is using a sink or an eyewash station.

A video explaining Vibrosight:

Provided by Carnegie Mellon University

APA citation: Sound, vibration recognition boost context-aware computing (2018, October 15) retrieved 14 October 2019 from

<https://techxplore.com/news/2018-10-vibration-recognition-boost-context-aware.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.