Noise pollution represents a genuine health hazard. Here, Sabri Pllana (Linnaeus University) explains how his team have developed an effective machine learning solution to help combat this problem.

Sound effects

A machine learning-driven IoT solution for noise classification

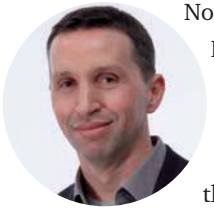

Noise is any undesired environmental sound. Recent studies have found that exposure to noise pollution may increase the risk for health issues, such as heart attack, obesity, impaired sleep or depression. Following the Environmental Noise Directive (END) 2002/49/EC, every European Union (EU) member state must assess environmental noise and develop noise maps every five years.

As sources of noise, including traffic, construction sites, music and sport events, may change over time, there is a need for continuous monitoring of noise. Health-damaging noise often occurs for only few minutes or hours, and it is not enough to measure the noise level every five years. Furthermore, sounds at the same decibel (dB) level may be perceived either as annoying noise or as a pleasant music. Therefore, it is necessary to go beyond the state-of-the-art approaches that measure only the dB level and actually identify the type of the noise. The internet of things (IoT) is a promising technology for improving many domains, such as e-health, and it may be also used to address the issue of noise pollution.

We have developed an approach for noise classification in smart cities using machine learning on a low-power and inexpensive IoT unit that comprises a Raspberry Pi Zero W and a ReSpeaker 2-Mic Pi HAT (Figure 1). The Raspberry Pi Zero W is the size of a credit card and costs only about €10; it may be used as an affordable computer to learn programming or to build smart devices, such as our solution for noise classification. We use ReSpeaker 2-Mic Pi HAT, a dual-mic array expansion board for Raspberry Pi, for sound sensing. Mel-frequency cepstral coefficients (MFCC) are extracted as audio features and applied to two classifiers: support vector machine (SVM) and k-nearest neighbours (KNN). We have implemented all algorithms used for our solution of the noise classification problem in Python.

To investigate the performance of the system, we conducted experiments with eight different classes of environmental sounds: quietness, silence, car horn, children playing, gunshot, jackhammer, siren, and street music. For the training of the system we selected noise-relevant environmental sound clips from popular sound datasets, such as UrbanSound8K and Sound Events. Our dataset contains more than 3,000 sound excerpts with lengths of up to four seconds. We observed in our experiments with various environmental sounds (such as car horns, jackhammer, or street music) that our solution provides high noise classification accuracy in the range of 85%-100%. Future work will investigate the usefulness of our solution for a large number of Raspberry Pi devices.

"Exposure to noise pollution may increase the risk for health issues, such as heart attack, obesity, impaired sleep or depression"

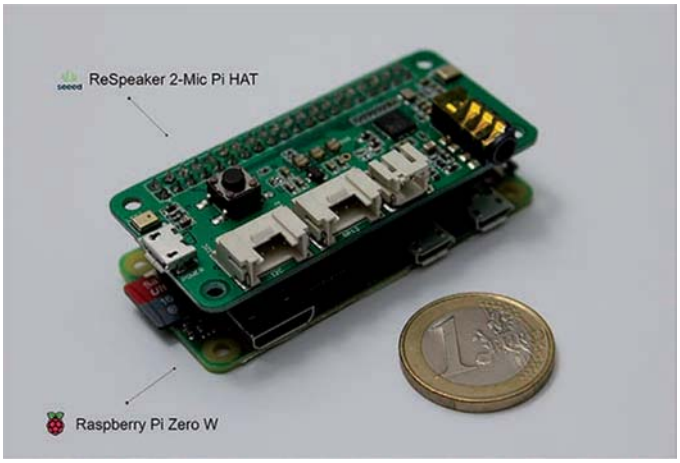

The proposed noise classification hardware platform consists of a Raspberry Pi Zero W and a ReSpeaker 2-Mic Pi Hat