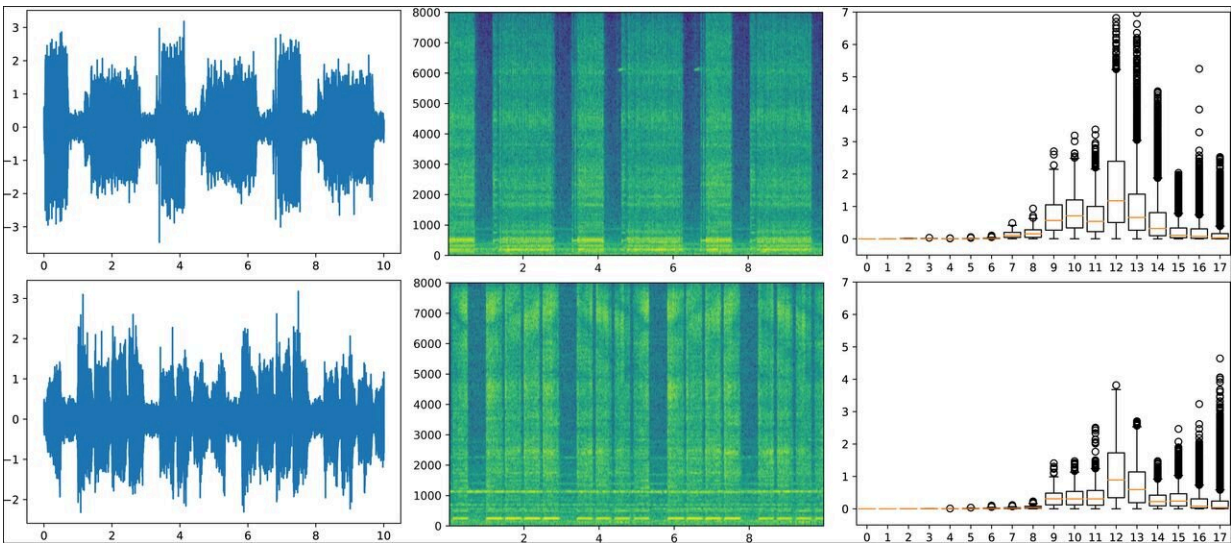


Artificial intelligence listens to the sound of healthy machines

February 25 2022, by Florian Meyer



Examples of abnormal signals. Shown are raw data, log-spectrogram, and obtained coefficient distribution per level for two abnormal measurements of the slider rail 0 at SNR 0 dB. Credit: DOI: [10.1073/pnas.2106598119](https://doi.org/10.1073/pnas.2106598119)

Sounds provide important information about how well a machine is running. ETH researchers have now developed a new machine learning method that automatically detects whether a machine is "healthy" or requires maintenance.

Whether railway wheels or generators in a power plant, whether pumps or valves—they all make sounds. For trained ears, these noises even have

a meaning: devices, machines, equipment or rolling stock [sound](#) differently when they are functioning properly compared to when they have a defect or fault.

The sounds they make, thus, give professionals useful clues as to whether a machine is in a good—or "healthy"—condition, or whether it will soon require maintenance or urgent repair. Those who recognize in time that a machine sounds faulty can, depending on the case, prevent a costly defect and intervene before it breaks down. Consequently, the monitoring and analysis of sounds have been gaining in importance in the operation and maintenance of technical infrastructure—especially since the recording of tones, noises and acoustic signals is made comparatively cost-effective with modern microphones.

To extract the required information from such sounds, proven methods of signal processing and data analysis have been established. One of them is the so-called wavelet transformation. Mathematically, tones, sounds or noise can be represented as waves. Wavelet transformation decomposes a function into a set of wavelets which are wave-like oscillations localized in time. The underlying idea is to determine how much of a wavelet is in a signal for a defined scale and location. Although such frameworks have been quite successful, they still can be a time-consuming task.

Detecting defects at an early stage

Now ETH researchers have developed a machine learning method that makes the wavelet transformation fully learnable. This new approach is particularly suitable for high-frequency signals, such as sound and vibration signals. It enables to automatically detect whether a machine sounds "healthy" or not. The approach developed by postdoctoral researchers Gabriel Michau, Gaëtan Frusque, and Olga Fink, Professor of Intelligent Maintenance Systems, and now published in the journal

PNAS, combines signal processing and machine learning in a novel way. It enables an intelligent algorithm, i.e. a calculation rule, to perform acoustic monitoring and sound analysis automatically. Due to its similarity to the well-established wavelet transformation, the proposed machine learning approach offers a good interpretability of the results.

The researchers' goal is that in the near future, professionals who operate machines in industry will be able to use a tool that automatically monitors the equipment and warns them in time—without requiring any special prior knowledge—when conspicuous, abnormal, or "unhealthy" sounds occur in the equipment. The new machine [learning process](#) not only applies to different types of machines, but also to different types of signals, sounds, or vibrations. For example, it also recognizes sound frequencies that humans—such as high-frequency signals or ultrasound—cannot hear by nature.

However, the learning process does not simply beat all types of signals over a bar. Rather, the researchers have designed it to detect the subtle differences in the various types of sound and produce machine-specific findings. This is not trivial since there are no faulty samples to learn from.

Focused on 'healthy' sounds

In real industrial applications, it is usually not possible to collect many representative sound examples of defective machines, because defects only occur rarely. Therefore, it is not possible to teach the algorithm what noise data from faults might sound like and how they differ from the healthy sounds. The researchers, therefore, trained the algorithms in such a way that the machine learning algorithm learned how a machine normally sounds when it is running properly and then recognizes when a sound deviates from normal.

To do this, they used a variety of sound data from pumps, fans, valves, and slide rails and chose an approach of "unsupervised learning," where it was not them who "told" an algorithm what to learn, but rather the computer learned autonomously the relevant patterns. In this way, Olga Fink and her team enabled the learning process to recognize related sounds within a given type of machine and to distinguish between certain types of faults on this basis.

Even if a dataset with faulty samples would have been available, and the authors could have been able to train their algorithms with both the healthy and defect sound samples, they would never have been certain that such a labeled data collection contained all sound and fault variants. The sample might have been incomplete, and their learning method might have missed important fault sounds. Moreover, the same type of machine can produce very different sounds depending on the intensity of use or the ambient conditions, so that even technically almost identical defects might sound very different depending on a given machine.

Learning from bird songs

However, the [algorithm](#) is not only applicable to sounds made by machines. The researchers also tested their algorithms to distinguish between different bird songs. In doing so, they used recordings from bird watchers. The algorithms had to learn to distinguish between different bird songs of a certain species—ensuring also that the type of microphone the bird watchers used did not matter: "Machine learning is supposed to recognize the bird songs, not to evaluate the recording technique," says Gabriel Michau. This learning effect is also important for technical infrastructure: even with the [machines](#), the algorithms have to be agnostic to the mere background noise and to the influences of the recording technique when aiming to detect the relevant sounds.

For a future industrial application, it is important that the machine

learning will be able to detect the subtle differences between sounds: to be useful and trustworthy for the professionals in the field, it must neither alert too often nor miss relevant sounds. "With our research, we were able to demonstrate that our machine learning approach detects the anomalies among the sounds, and that it is flexible enough to be applied to different types of signals and different tasks," concludes Olga Fink. An important characteristic of their learning method is that it is also able to monitor the sound evolution so that it can detect indications of possible defects from the way the sounds evolve over time. This opens several interesting applications.

More information: Gabriel Michau et al, Fully learnable deep wavelet transform for unsupervised monitoring of high-frequency time series, *Proceedings of the National Academy of Sciences* (2022). [DOI: 10.1073/pnas.2106598119](https://doi.org/10.1073/pnas.2106598119)

Provided by ETH Zurich

Citation: Artificial intelligence listens to the sound of healthy machines (2022, February 25) retrieved 3 May 2024 from <https://techxplore.com/news/2022-02-artificial-intelligence-healthy-machines.html>

<p>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.</p>
--