



future

What is BBC Future? Latest In Depth Video

OUR PICKS
Best of Future

NEW
Future Now

Future Now | Artificial intelligence

Why an AI firm is busy smashing thousands of windows

future now

More articles >

Breaking glass, barking dogs and crying babies – to computers it's all just noise. But a new wave of artificial intelligence is being taught how to listen in.

By Paul Marks
11 January 2017

In a sound-proofed hangar on an RAF airbase just north of Cambridge, UK, Chris Mitchell and his colleagues are busy using sledgehammers to teach their computers a lesson.

The team has gathered thousands of window panes and doors, all of different shapes and sizes, which they then smash, one by one, recording the distinctive shattering sound of each type of glass. Sometimes they swing sledgehammers or garden spades, sometimes they throw bricks. "We completely underestimated the mess it would make," says Mitchell. "And how tiring it would be."

Welcome to the latest frontier of artificial intelligence. Mitchell is CEO and founder of Audio Analytic, a Cambridge-based start-up that is training a machine learning system to recognise the sound of breaking glass.



The team worked with vets to track down and record the barks of as many different breeds of dog as they could, from Pekinese to Great Danes (Credit: Getty Images)

And it's not just glass: the company is also teaching computers to pick out other sounds that are important to humans, like smoke alarms, bawling babies and barking dogs. The idea is to build this ability to recognise sounds – without confusing a dropped glass with a smashed

window, say - into smart-home systems that will alert you when an intruder breaks in or your child starts to cry.

In the last few years, computers have become very good at understanding the world by sight. AIs are now better than humans at recognising certain objects, especially faces. But apart from speech recognition - which is at the heart of services like Apple's Siri, Google Home and Amazon's Alexa – highly accurate sound recognition has been given little attention. Everyday noises are just background din to most machines.



Training AI on music is a good way to teach it to spot patterns in complex data

Mitchell wants to change that. "What we're working on is a new field of AI that we call artificial audio intelligence," says Mitchell. "It's not something that has been tackled before in any meaningful sense."

Audio Analytic is part of a new wave of companies training machine learning systems to spot patterns in sounds. Uberchord, based in Berlin, is developing an AI that can help people learn to play guitar. It listens to you strum and tells you when you have your fingering wrong. Uberchord is one of several AI companies working with sound that Abbey Road Studios – one-time recording home of the Beatles – is investing in.

Another company, Cambridge Consultants, has taught an AI to recognise different genres of piano music, like ragtime or baroque. The system, called Aficionado, was trained on just a few hundred hours of piano playing, including both professional recordings and amateur practice videos taken from YouTube. The training data was deliberately patchy, says Monty Barlow at Cambridge Consultants. "We were challenging the AI to handle the near infinite complexity of live music."

Aficionado's musical chops are not just for show, however. Training the system on music – and getting it to ignore irrelevant factors such as tempo, volume or tone – turns out to be a good way to teach it to spot

patterns in complex data in general, whatever it represents. Aficionado's first task will be to identify faults in telecommunications networks.



The ability to tell the difference between a dropped wine glass and a smashed window could let smart homes alert you if an intruder breaks in (Credit: Audio Analytic)

But Audio Analytic has bigger ambitions. "We want to create a taxonomy of all sounds, and that is a huge undertaking," says Mitchell. So far, the company's software can identify breaking windows, crying babies and smoke alarms. At the Consumer Electronics Show in Las Vegas last week, they added barking dogs to their repertoire.

They are also working on an anomaly detector, which will pick up sounds that seem out of the ordinary - a change from the normal background hubbub - like the clatter of someone falling over. Or the hiss of a leaking water pipe. Eventually, they want to add car alarms and perhaps - for the US market - gunshots. Audio Analytic then plans to license these sound-recognition systems to makers of smart-home gadgets.

The ability to recognise different sounds matters, says Nina Bhatia, managing director of Hive, a UK-based smart thermostat and lighting company. "It is fast becoming absolutely vital for smart home technology to detect and interpret a wide range of ambient sounds, so people can

respond easily and quickly to what's going on in their homes when they're not there," she says. "You could be alerted if your smoke alarm was going off while you're in a meeting at work, and not just when you're on your sofa."

“

A baby's cries could turn on a nightlight and make a lullaby play from a nearby speaker

As well as sending an alert to your phone, such systems could also take actions by themselves. A smashed window could make the lights turn on. A baby's cries could turn on a nightlight and make a lullaby play from a nearby speaker.

Indeed, Chinese electronics firm Sengled is using Audio Analytic's technology in a smart lamp with a speaker built into its base. Other smart-home firms are building it into their devices too, including thermostats, which as they are often installed in a central position in a home, are in a good place for eavesdropping.

The hard part is making sure the AI correctly identifies what it hears, because false alerts could cause havoc. Yet machine learning systems are only as good as the examples they are trained on. As Mitchell puts it: "AI is bloody useless unless you have data."

“

We smashed glass for weeks and weeks
– Chris Mitchell

Getting that data is hard work. "We smashed glass for weeks and weeks," says Mitchell. "Some of these windows were full floor-to-ceiling shopfront ones. Smash those and they have a chance of taking your foot or leg off as the glass comes down."





Bird or smoke alarm? Some species of parrot can confuse an AI with their remarkably good ability to imitate an alarm's beeps (Credit: Alamy)

To get enough recordings of crying babies, the firm worked with parents' groups in Cambridge. To catalogue what they were recording, they then had to come up with their own lexicon to describe the different types of crying, says Mitchell. "For example, there's a very raspy one that seems to come from the back of the throat that we called the 'vocal cry'."

Dogs were somewhat easier. Working with vets, they tracked down as many different breeds as they could and introduced their AI to barks from tiny Pekinese up to sofa-sized Great Danes.

“

The company's office is filled with hundreds of smoke alarms

To teach their system what a smoke alarm sounds like, Audio Analytic simply bought as many different models as they could online. Hundreds are now stacked on shelves in their offices. At first their AI had trouble telling the beeps of a smoke alarm from other household beeps, such as ringing phones, alarm clocks and oven timers. So they trained it to focus not only on the alarm's pitch and duration but on the signature gap

between the beeps.

But no matter how many windows you smash or smoke alarms you set off, there will always be surprises down the line. There is a parrot species that does an uncanny impression of a smoke alarm. So Audio Analytic has had to teach its system to ignore this feathery false alarm.

Another sound they want to teach their system to look out for is the pitch

and intonation changes of aggressive human shouts - somebody threatening violence, say. This doesn't vary much with language or culture, says Mitchell. Distinctive changes in vocal sounds come when adrenaline floods the body and affects the voice box.

Audio Analytic has had to put this one on hold, however. They found that the sounds of chickens and chainsaws in a neighbourhood would also trigger their aggression detector.

It's a noisy world out there – but AIs are starting to listen.

