


AI Can Identify a Person From a Single Footstep

Originally published at <https://www.unite.ai/ai-can-identify-a-person-from-a-single-footstep/>




A new research initiative has produced a low-cost system capable of identifying a person based on the sound of their footsteps, from as little as a single step.

In the [paper](#) *Passive mUlti-peRson idEntification via Deep Footstep Separation and Recognition (PURE)*, a collaboration between researchers from Nanyang Technological University and the University of Kentucky, among others, identification rates were established at up to 90%, from audio samples that are extremely brief.




Five distinctive footstep profiles captured in PURE.

The architecture for PURE relies on data from an array of commodity microphones, with the raw audio capture denoised via background [spectral subtraction](#). Where the signal to noise ratio is high, including conversation that is occurring at time of capture, a source separation algorithm is activated to perform discrete extraction of the footsteps.




The footstep audio is clarified and analyzed via [domain adversarial adaptation](#), with the framework comprising a feature extractor, an identity predictor, and a domain discriminator.



Hardware for PURE

The equipment used for PURE is a microphone array embedded in a customized rig based around the Raspberry Pi 4.




The microphones capture audio at the highest available rate for 'structure borne' signals (feet contacting the ground), since this data is of extremely short duration, and needs to be as detailed as possible. However, air-borne footsteps (the sound feet make in the arc towards the next contact with the ground) are downsampled to 16kHz in order to save local processing capacity for structure-borne steps.

The researchers synthesized a training data set from the [Footsteps Sound Effects Soundboard](#), as well as from [Footsteps Sound Effects](#) from Epidemic Sound. The audio component from various Ted Talks were used to produce training data for the process of exfiltrating footsteps from background conversation.

Preventing 'Replay Attacks' In Footstep Recognition

A system of this nature needs to be resilient to 'replay attacks', where a malefactor might record a particular footstep pattern and replay it in the hope that the system will identify the recording as a live user.

To thwart this, PURE analyzes the Time-of-Arrival (ToA) in 'contact' footsteps, and the Angle-of-Arrival (AoA) in airborne footsteps.



The lack of dynamic information in replayed footsteps reveals them fairly easily, though it is necessary to account for this when processing the data. By observing the natural irregularity of footsteps, and also their speed in the context of the environment (since it is unlikely that one would either run or dawdle for instance, in an office environment), it's possible to ensure that the data being received is authentic.

The project uses [beamforming](#) techniques to calculate ToA, but the extraction of AoA is more complex, requiring an R-Net neural network that, again, uses adversarial learning to calculate the range of a footstep. This is essentially the same model as the earlier neural network, except that the identity predictor is substituted with a range estimator.

Accuracy

PURE was tested in a wide range of acoustic environments, and using a variety of walking speeds over a range of distances. As the number of people creating footsteps increases, the accuracy naturally drops, as it also does when the speed of multiple footstep sources increases.

However, depending on domain adaptation, the results over 100 trials found that the system could identify a user from 3-5 footsteps with a range of accuracy from 90.73% to 96.53%; from 2-3 footsteps with a range of accuracy from 88.16% to 95.92%; and from a single footstep with a range of accuracy from 81.75% to 88.6%.

The researchers foresee wide applicability for PURE, due to the low cost of the commodity hardware involved, and the fact that it also outperforms similar systems in terms of latency and accuracy, while being robust to environmental interference and replay attacks.

The Growth Of Gait-Analysis

This particular sphere of machine learning research has centered primarily on computer vision over the last ten years, and received a cultural fillip when [used](#) as a plot device in *Mission Impossible: Rogue Nation* (2015).

To date, gait recognition technologies have been proposed for use in [elder care](#), post-surgical [rehabilitation](#), and more controversially for personalized ad serving in retail environments, although such a system obviously has potential uses for employee monitoring in secure environments.

In 2018 it was [reported](#) that Chinese authorities use vision-based gait analysis from AI development company [Watrix](#) as a facet in its closed circuit public surveillance systems.

Gait recognition has also been implemented by [monitoring](#) the reflectance of Wi-Fi signals.

However, all these approaches have inherent limitations, either requiring lighting conditions that cannot be guaranteed, unoccluded views, prohibitively expensive specialized equipment, overly specific local conditions, or body-worn equipment, among other hurdles.

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Portfolio site: martinanderson.ai

Contact: martin@martinanderson.ai



Discover More