Detection and Localisation of Gunshots

Martin Chan Sofie Karlsson

Lund University

In this master's thesis we study the possibilities to detect and position sharp sounds, for example gunshots, using only microphones. Our system requires minimal maintenance and is easy to use. The performance is adequate for large environments and results are achieved within seconds.

Introduction

The frequent school shootings in the US has drawn a lot of attention; a system to detect gunshots are therefore interesting to develop. During an investigation, information such as when and where things happened are crucial. Imagine a person firing a gun in a room with a surveillance system installed. The shooter sees the cameras and immediately destroys them. Now, only the speakers are left and they are not threatening to the criminal. The catch is that these speakers have built-in microphones recording continuously. With the information in these recorded tracks, clues about the event can be extracted. In this paper, a step by step guide through the methods needed to detect and position the gunshots automatically will be explained. Firstly, we need to detect our key event, a gunshot sound.

Detection

The gun has fired, windows shattered, and people are screaming. For a human being, it is easy to distinguish all these sounds, but how would a computer tell the difference? The answer is *artificial neural networks*. These are machine learning models trying to mimic the brain of a human. In recent times, they have been successful in everything from recognising objects to maintaining a conversation. Just like the brain, an artificial network consists of millions of neurons connected to each other and just like us humans, the models need to experience in order to learn.

We teach our *artificial neural network* by presenting 600 soundtracks to it, some of them containing gunshots. After training with these sounds several times, the best model has an accuracy of 98-100%.

Positioning

Once all gunshots in the recorded tracks have been pinpointed, the positioning of them can begin. As the microphones are connected to a joint network, all audio files are asynchronous. This complicates calculations as a sound source positioning system requires that every device start recording simultaneously. By modifying classic methods for a sound source positioning task, we have succeeded to solve the synchronisation problem in addition to positioning gunshots. This means, one can follow the path of our gunman during the situation.

The final system localises a single sharp sound with an error of 50 cm. In a case with multiple gunshots present, it has the ability to calibrate and solve the positioning task with a decreased error of 25 cm instead. This accuracy is adequate for larger environments such as classrooms, department stores, airports, and markets.

The detection and positioning can be done in less than 2 seconds making our system a viable choice for real-time monitoring. After a manual start of the recording, everything from the recorded audio files to calculated positions of interesting sound sources are done automatic. Everything is easy to use and maintain. Hopefully, the shooter can soon be caught.