Audio Engineering Society

Convention Paper

Presented at the AES 161st Convention

2026 October 30-November 1, Nashville, TN, USA

This paper was peer-reviewed as a complete manuscript for presentation at this convention. This paper is available in the AES E-Library (http://www.aes.org/e-lib) all rights reserved. Reproduction of this paper, or any portion thereof, is not permitted without direct permission from the Journal of the Audio Engineering Society.

Gunfire location estimation with unsynchronized audiovisual recordings

Robert C. Maher

Electrical and Computer Engineering, Montana State University, Bozeman, MT 59717-3780

Correspondence should be addressed to rmaher@montana.edu

ABSTRACT

In audio forensic analysis, estimation of sound source location such as audible gunshots can be accomplished using time difference of arrival (TDOA) calculations from multiple concurrent audio recordings using principles known as multilateration. The TDOA and multilateration principles ordinarily require time synchronization of the audio recordings, and precise knowledge of the position of the multiple audio recording devices. However, it is increasingly common to have gunshot incidents recorded simultaneously by randomly located and unsynchronized ad hoc combinations of mobile phone videos, residential and business video surveillance systems, and body-worn cameras. As described in this paper, forensically meaningful location estimates may be deduced in some circumstances from these unsynchronized audiovisual recordings by identifying known landmarks in the videos for geolocation, and identifying coincident sounds in the corresponding audio recordings for relative time synchronization. Case study examples are included.

1 Introduction

The increasing deployment of home surveillance devices that record video and audio, and the increasing likelihood that bystanders may record civil and criminal incidents and other community disturbances with mobile phones, means that there is an increasing likelihood of gunshot sounds being captured by personal recording devices. Criminal investigations frequently include seeking and collecting the user-generated recordings following a shooting incident [1][2].

Theoretically, if the precise location of the recording devices at the time of the gunfire could be determined, and if the precise time synchronization of the various audio files could be obtained, it would be feasible to use the time of arrival of a gunshot sound in each recording to perform multilateration: computing an estimate of the likely position of the firearm that produced the muzzle blast sound [3][4]. However, unlike purpose-built gunshot location technology such as the commercial ShotSpotter® system from SoundThinking™, user-generated audio recordings generally do not have precise location information for the recording microphones, nor do the recordings have sufficiently accurate and precise time-of-day information to allow time synchronization at the scale of individual audio samples. Furthermore, common user-generated recordings these days are almost always in the form of perceptual audio compressed files (e.g, MP4), which utilized lossy data compression.

Despite these shortcomings, audio forensic examiners can utilize ad hoc information from the evidentiary audiovisual recordings to estimate the microphone locations, and then to determine time synchronization by identifying a sound present in each of the recordings that was produced at a known location, thereby creating a “slate” signal for relative time difference calculations [1]. The proposed procedure, and successful examples of applying this methodology from actual forensic audio cases, are described in the following sections of this paper.

This paper is organized as follows. First, the characteristics of gunshot sounds and their acoustical propagation in urban areas are reviewed, including references to the growing body of scholarly publications in the field. Next is a summary of the typical methodology for utilizing user-generated audio recordings that captured the sound of gunfire. Examples of the technique applied to actual audio forensic cases are presented next, followed by several concluding remarks.

2 Characteristics of Gunshot Sounds

Although any sound that has distinctive temporal characteristics can be suitable for unique identification and time estimation, the audio synchronization methodology described in this paper is applied primarily to gunshot sounds. Therefore, it is helpful to review the characteristics of the recorded sounds of gunfire [5][6][7][8].

Conventional firearms produce a variety of sounds when they are used, including the relatively quiet sounds of the gun's mechanical action, loading, cocking, and spent cartridge ejection, and, of course, the relatively loud sound of the muzzle blast [9].

Ammunition in a conventional firearm consists of a sealed cylindrical cartridge containing a hollow casing filled with a gunpowder propellent, and a bullet crimped onto the open end of the casing, or in the case of a shotgun, a receptacle containing shot pellets or a slug. The triggering mechanism of the gun causes a firing pin to strike the ammunition cartridge, initiating rapid combustion of the gunpowder inside the cartridge case. This creates a hot, expanding gas mixture that propels the projectile out of the gun barrel. The *muzzle blast* refers to the abrupt report of the expanding gas emitted from the barrel as the bullet leaves the gun. The duration of the muzzle blast is typically less than 3 milliseconds, and the "bang" sound propagates at the speed of sound (e.g., c=343 m/s at 20°C) in all directions from the muzzle, with the loudest sound level in the direction the gun barrel is pointing.

The muzzle blast acoustic disturbance encounters the surrounding surfaces (ground, walls, and other physical obstacles), and a microphone in the vicinity will record the direct sound of the gun, along with acoustic reflections from the surfaces, diffracted sound bending around obstacles, and the arrival of multiple-order reflections and reverberation. Thus, a microphone recording a gunfire incident will receive a mixture of direct, diffracted, and reflected sounds from the gunshot, plus environmental sounds and competing noise from other concurrent sound sources. This means that an audio forensic examiner will need to interpret the amplitude and timing information carefully: each gunshot sound is unique to the location and physical surroundings, type of firearm and microphone, orientation of firearm and microphone, type of audio coding, etc.

Although perhaps obvious, it is essential to recognize the effect of the speed of sound when interpreting the temporal characteristics of sounds in recordings from microphones at different locations with respect to the sound source. The relative time of arrival of a sound at a microphone located, say, 50 meters from the source and at another microphone 200 meters from the source will be $\Delta(\text{dist})/c$, so in this example, (200-50)/343 = 437 ms. As will be described below, the observed time difference in concurrent and synchronized audio recordings is the basis for the proposed multilateration estimation methodology.

3 Characteristics of Typical User-Generated Recordings for Forensic Use

Audio *user-generated recordings* (UGRs) come from audiovisual devices such as handheld mobile phones, memo recorders, and consumer home surveillance systems such as Nest Cam™, Ring™, Blink™, Wyze™, and many others. In some cases, the UGRs may come from law enforcement body-worn cameras, journalists' newsgathering recordings, or commercial security and surveillance services. These recording devices are generally designed to capture ordinary audio sources such as background soundscapes and speech conversations, so the presence of unusually loud, impulsive sounds like gunshots in a recording may be distorted by possible clipping, uncertainty due to competing noise, and interpretation challenges due to acoustic reflections, diffraction, and reverberation [10][11].

Most UGR audio is encoded and stored using a lossy audio perceptual coding (e.g., MP4), which retains human perceptual fidelity but may not give an accurate waveform desired for audio forensic analysis. In a situation in which the audio to be encoded has already been clipped or distorted by the microphone and preamplifier system, the lossy audio coder will represent the distorted audio in a manner that would give a representation for a human listener's ear, not necessarily in a from for audio forensic analysis.

With user-generated recordings, both the spatial location of the audio device and the precise time-of-day reference are generally not known with verifiable accuracy. This means that alternative methods need to be established to determine microphone location and precise timing for synchronization.

In contrast to common consumer audio recordings, it should be noted that purpose-built systems intended for gunshot detection, such as the commercial ShotSpotter system from SoundThinking, are designed with appropriate microphone sensitivity and preamplifier gain so that gunfire sounds do not clip the system, and perceptual audio coding is not used. ShotSpotter uses GPS signals at each microphone sensor to determine precise sensor location and to synchronize audio timing [12].

4 The Ad Hoc Synchronization and Localization Methodology

Spatial interpretation of UGRs is essential for efforts to determine likely sound source location. Most contemporary surveillance cameras are in a fixed position such as attached to a residential building or pole, and the audio microphone is generally part of the camera module. Thus, the microphone position of surveillance camera recordings can often be determined by verifying the camera's position from police reports or from using resources such as Google Street View™.

Handheld mobile devices often include a means of geolocating the latitude and longitude of the device, and the geolocation can often be found within the metadata of the original recorded file. In cases where the device's geolocation is missing or suspect, details visible in the video portion of the recording can be used to estimate the camera's position at the time of the incident. The ad hoc location approach is to identify objects and landmarks visible in the video scene, and then to compare the objects and landmarks from the camera's perspective of the scene to known imagery from the location, such as contemporary still photography of the scene, or imagery from Google Street View. The precision of the location estimate may vary, and this uncertainty must be considered when using the location estimates. At present, the location uncertainty varies from case to case because of the unique aspects of each UGR.

Like spatial localization, time synchronization of UGRs is also essential for efforts to determine likely sound source location. Many recording systems use a time reference of some kind to indicate the date and time that the recording began or ended, and sometimes a fiducial "timestamp" display is included in the video itself. However, the veracity of the time indications for a video file and a visible video frame timestamp are generally not known, and therefore not sufficiently reliable for synchronization of multiple UGRs. Moreover, the use of video frame counts for reliable timebase estimates can be problematic due to variable frame rate and decoding details when stepping through the video frames. The decoded audio track consisting of successive digital samples can generally give a valid relative time sequence count, at least for intervals of a few minutes or less. Nevertheless, it is important to understand the potential waveform timing challenges at the audio sample level when using perceptually encoded and decoded material.

The recommended time estimation approach is to identify the onset time of a distinctive sound that is audible in the audio track of each of the UGRs, label that onset time as a reference, and then to identify the time of occurrence of other sounds within each recording with respect to the identified reference time. Because the various UGRs came from devices at different spatial locations, the identified reference time in each recording represents the time-of-arrival of that sound at the recording position, not an absolute time of day. Under the assumption that the recording devices are not moving during the episode of interest, other sounds in each recording that have the same relative time offset with respect to the reference can be assumed to be from the same source location as the reference sound, while sounds with differing relative delay in the various recordings can be assumed to come from a different location than the reference sound.

5 Examples From Audio Forensic Cases

It is helpful to consider practical examples of the ad hoc synchronization and localization methodology using actual case studies. Two incidents involving asynchronous user-generated recordings are considered next.

5.1 The Trump Assassination Attempt (13 July 2024)

On Saturday, July 13, 2024, a shooting incident took place at an outdoor rally held by then presidential candidate Donald J. Trump at the Butler Farm Show grounds in Butler, Pennsylvania. Two individuals were killed and two others were critically injured. Within a few hours of the shooting, audio forensic examination of sound recorded by the lectern

microphone identified ten gunshot sounds, with the initial shots coming from approximately 135 meters away from the podium. All ten gunshot sounds occurred within 16 seconds. [13]

In the Trump rally incident, more than a dozen user-generated recordings capturing the gunfire were eventually collected by professional journalists reporting from the scene. The authorities later identified a single perpetrator who shot an AR-15 rifle from the roof of a building 135 meters north of the podium—confirming the rapid audio forensic interpretation of audio from the lectern microphone—but among the investigative questions was whether multiple shooters could have been involved.

The approach to answering the question about multiple shooters was: (a) determining a set of the available UGRs that captured all ten gunshot sounds from locations spread across the outdoor venue, (b) identifying the onset time of the muzzle blast sound of the initial gunshot report in each UGR, and (c) determining the relative time of each successive muzzle blast sound with respect to the initial onset time in each recording. Some of the UGRs were in locations that were near the path of the supersonic bullets, so care had to be taken to distinguish between the muzzle blast sounds and ballistic shock wave sounds (and reflections) due to the supersonic projectiles [6]. As noted in the previous section, the absolute time of day of the gunshots was not observable, but using the time-of-arrival of the sound of the first gunshot gave a relative time reference for each of the recordings.

To illustrate the procedure, the approximate position of five of the available mobile phone recordings (*I*, *II*, *III*, *IV*, *V*) are shown in Figure 1. The corresponding audio waveforms aligned at the sound of the first muzzle blast are shown in Figure 2.

The first important observation from Figure 2 is that the relative time-of-arrival of the muzzle blast sounds for shots A-H are coincident: the time elapsed between the first shot to the arrival of each subsequent shot is the same at each of the five different microphone locations. This is only possible if the position of the gun that produced the muzzle blast sounds and the position of each microphone did not change during the eight successive shots. Had a gun from a different relative position been fired, then time-of-arrival of that shot would differ among the various microphones, since the path lengths would differ. Thus, it is possible to conclude that the first eight shots came from a single location, and were all fired within about six seconds.

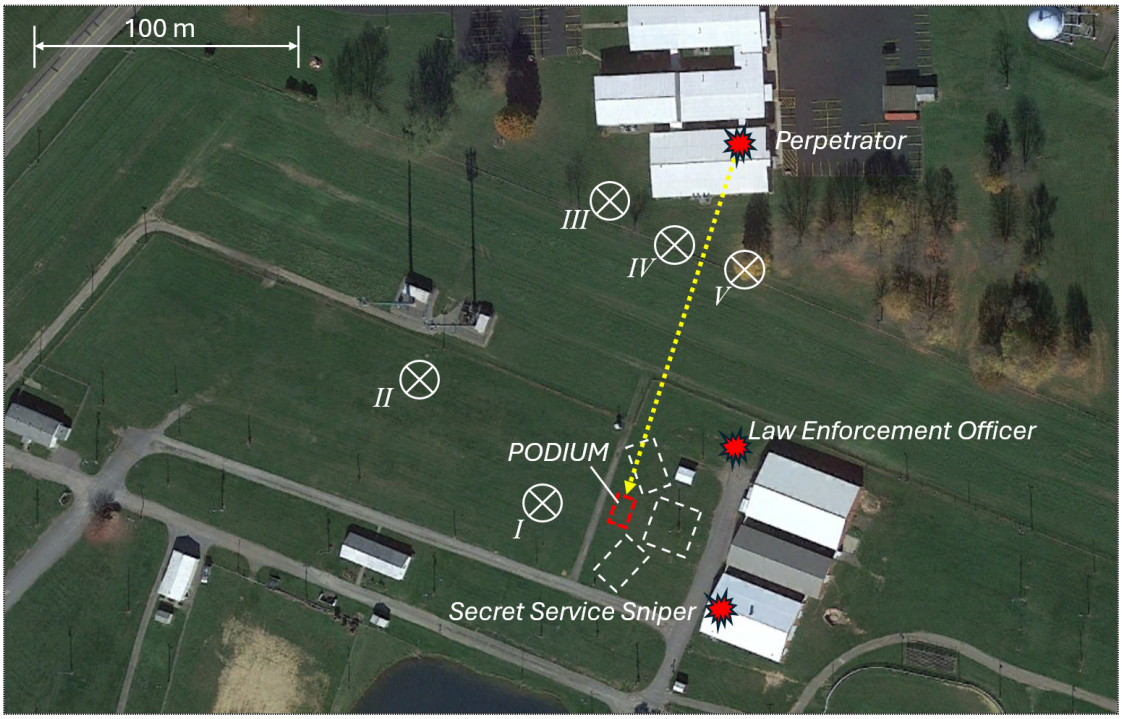


Figure 1. Aerial view of location of Trump assassination attempt, showing the estimated position of five user-generated recordings (*I*-*V*), and the reported location of the audible gunshots.

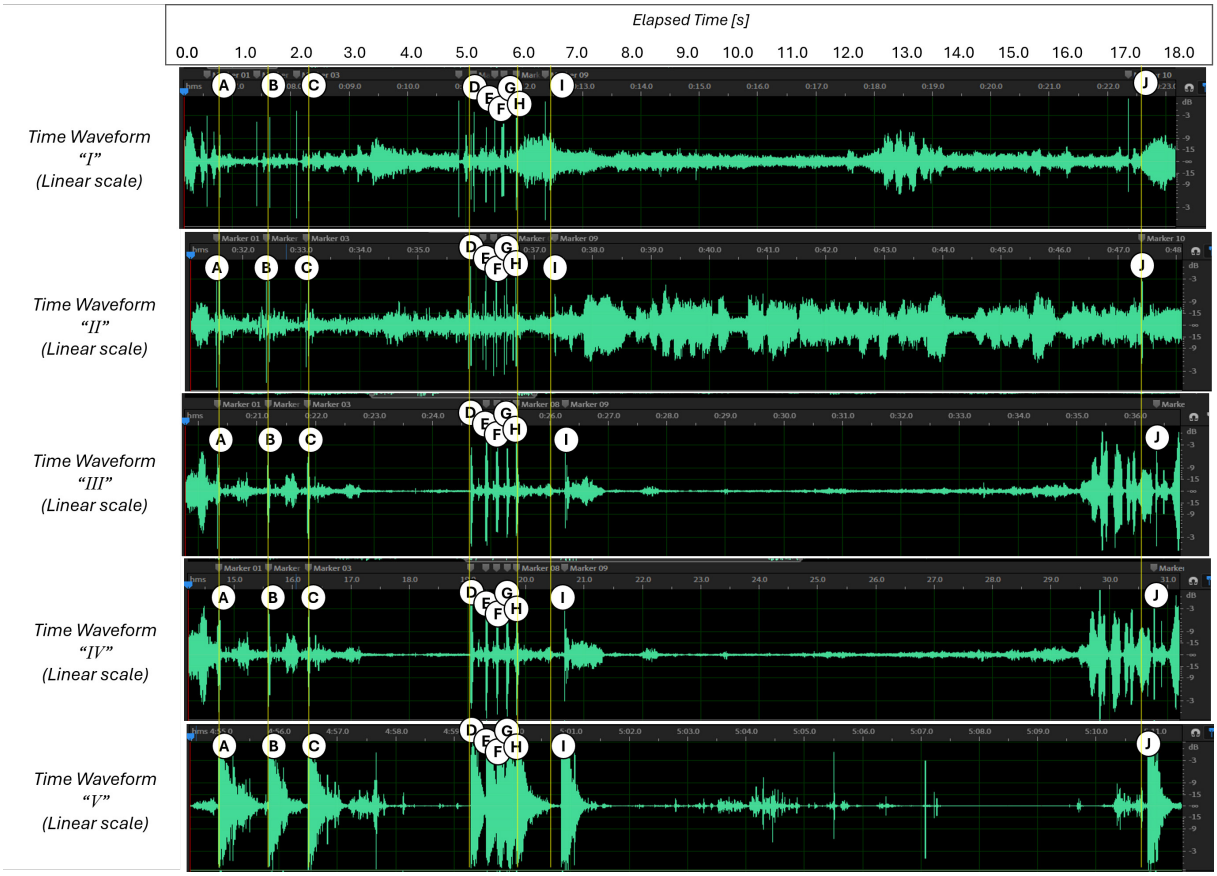


Figure 2: Five user-generated recordings of the ten audible gunshots from the Trump assassination attempt.

The second observation from Figure 2 is that, unlike shots **A**-**H**, the relative time-of-arrival of shot **I** and shot **J** differ among the five recordings. Since the microphones did not move appreciably during the 16 seconds between shot **A** and shot **J**, the differing time-of-arrival of **I** and **J** must be attributable to gunshots from different locations than the shots **A**-**H**. The *relative* timing results for the five UGRs *I*-*V* are summarized in Table 1.

Without overall synchronization of the UGRs, the proposed ad hoc method is to use the initial muzzle blast sound as the reference for each UGR, then determine the elapsed time from the first blast to each subsequent blast. As shown in Figure 2 and in Table 1, the timing comparison for all five UGRs gives comparable time span variability within 20 milliseconds of each other for shots **A**-**H**, while the difference is 356 ms for shot **I** and 589 ms for shot **J**, again affirming that the first eight audible shots came from the same location.

	Microphone					
Shot:	***I***	***II***	***III***	***IV***	***V***	***Span***
A	0.000	0.000	0.000	0.000	0.000	0.000
B	0.853	0.859	0.859	0.863	0.861	0.010
C	1.526	1.538	1.536	1.528	1.529	0.012
D	4.311	4.329	4.325	4.312	4.320	0.018
E	4.564	4.584	4.581	4.568	4.571	0.020
F	4.752	4.768	4.766	4.748	4.752	0.020
G	4.924	4.945	4.942	4.922	4.930	0.023
H	5.084	5.104	5.101	5.084	5.085	0.020
I	5.580	5.932	5.936	5.755	5.934	**0.356**
J	15.574	16.066	16.163	15.818	16.134	**0.589**

Table 1: Time in seconds of muzzle blasts A-J with respect to time of first audible shot A.

For shots **I** and **J**, the time span among the UGRs becomes several hundred milliseconds, indicating that the source of those shots differed from the source of shots **A**-**H**, and that shot **I** came from a different

location than shot **J**. With shot **I** determined by official reports to have come from a law enforcement officer on the ground north east of the podium, and shot **J** attributed to a Secret Service sniper posted on the roof of a building southeast of the podium, the relative timing shown in Table 1 determines geometrically that recording *I* was relatively closer to the law enforcement and secret service positions and farther from the perpetrator (shortest time gaps), and that recording *III* was relatively closer to the perpetrator and farther from the law enforcement and Secret Service rifles (longest time gaps). These results show the usefulness of the relative time synchronization of the user-generated recordings to confirm the officially reported sequence of events at the incident in Butler, Pennsylvania.

5.2 The Terrorist Shootings at Bondi Beach in Australia (14 December 2025)

On Sunday, December 14, 2025, in Sydney, Australia, a terrorist attack involving two perpetrators with multiple firearms took place in the Bondi Beach area of the city during an outdoor celebration of Hanukkah attended by about 1,000 Jewish residents. 16 individuals were killed and more than 40 were injured. Investigative journalist reports determined that the perpetrators fired approximately 82 shots in 6 minutes, and law enforcement officers fired an additional 26 shots [14].

The two attackers used three firearms: one Beretta BRX1 straight-pull rifle, and two Stoeger M3K 12-gauge semiautomatic shotguns [15]. The law enforcement officers are believed to have used 9-mm Glock handguns.

Journalists from the Australian Broadcasting Corporation were able to review more than 130 user-generated videos showing some portion of the attack, and 84 of the videos could be time synchronized using the audible gunshot sounds [16]. Several of the videos clearly showed visible puffs of smoke or debris emanating from the perpetrators' guns, which allowed localization of the attacker's firearms for nearly every shot as the attack progressed. However, there was very little visible evidence of the law enforcement gunshots in the videos due to the concealed location of the officers and the concentration of those filming the incident upon the perpetrators.

A question arose during the investigation about what time in the 6-minute sequence of the terrorists' gunfire the first law enforcement gunshot could be discerned, and the likely location of the officer who made the first response shot. The reporters from Australian Broadcasting pointed to a distinctive gunshot sound approximately 172 seconds after the first audible gunshot from the attackers. Reviewing a user video that showed nearly the entire attack, and checking several other UGRs from the time interval of interest, the distinctive gunshot could not be attributed to shots by either of the two perpetrators. Thus, the investigating reporters were interested to estimate the location of that *unknown* gunshot, assuming that it was likely the first shot from one of the law enforcement officers on the scene.

The terrorist attack took place in the Archer Park area of Bondi Beach. The two attackers began the shooting from an elevated pedestrian bridge west of the Hanukkah gathering, and one of the attackers also proceeded to walk west on a sidewalk toward the victims in Archer Park, shooting along the way. At the moment of interest, the shooter and microphone positions are summarized in Figure 3. The figure shows the location of the two perpetrators, and the estimated location of four user-generated recordings covering the time interval of interest.



Figure 3: Map of Archer Park portion of Bondi Beach Park, showing positions of perpetrator #1, perpetrator #2, and the location of four mobile phones (*I-IV*) that were recording approximately 172 seconds after the first shots.

Reviewing the sequence of audible gunshots in the available UGRs, the 4-second time segment of interest has a pattern of four shot sounds: the *unknown* shot believed to be from law enforcement, a shot by perpetrator #2 (audible, and visible in videos), then two shots in succession by perpetrator #1 (audible, and visible in video). The shots occurred at three different locations, and therefore the absolute arrival time of the gunshot sounds at the four different unsynchronized microphone positions differs. Note that the last two shots of the four-shot sequence were made by perpetrator #1, so the time interval between those two shots is expected to be the same in each recording at all of the observation positions.

The manual ad hoc time alignments of the four microphone recordings to match the arrival time of perpetrator #1's two shots are shown in Figure 4.

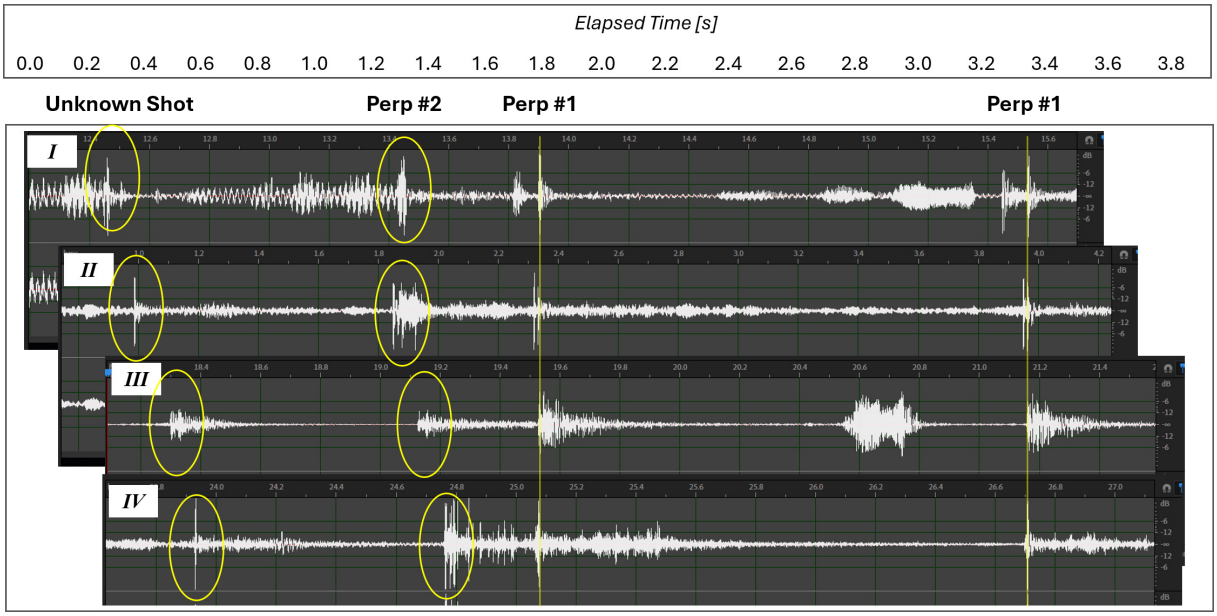


Figure 4: Relative time alignment of the four recordings (*I*-*IV*) using the two shots by perpetrator #1.

With the time alignment corresponding to the perpetrator #1 shots, it is noted that the *relative* time of arrival of the sound of the shot by perpetrator #2 at microphone *II* was slightly before microphone *I*, then arrived at microphone *III* and finally microphone *IV* in sequence. This confirms the geometry, since referring to Figure 3, the distance from perpetrator #2 to microphone *II* is shortest, then microphone *I*, then *III*, and then *IV*.

The time comparison of Figure 4 also provides useful information about the time of arrival of the sound of the *unknown* gunshot: the arrival occurs sequentially at microphones *I*, *II*, *III*, and *IV*, indicating that the unknown shot location must have been in the vicinity of microphone *I*, so that as the sound propagated away outward from that general location, it would encounter the other microphones in sequence.

Given the location estimates for the two perpetrators and the location estimates of the four microphones, it is possible to calculate the theoretical time-of-flight for the gunshot sounds based upon the distance and the speed of sound. The map-based distance estimates for the two perpetrators and the four microphones are summarized in Table 1 and Table 2.

Microphone	Distance from perp #1 [m]	Propagation time estimate [s]
I	71.10	0.207
II	58.10	0.169
III	70.62	0.206
IV	62.48	0.182

Table 2: Map-based estimated distance from perpetrator #1 to each microphone.

Microphone	Distance from perp #2 [m]	Propagation time estimate [s]
I	30.45	0.089
II	19.72	0.057
III	58.28	0.170
IV	88.95	0.259

Table 3: Map-based estimated distance from perpetrator #2 to each microphone.

Using the map-based geometrical predictions and the observed relative time difference of arrival of the *unknown* shot sound, a multilateration calculation can be performed that represents the locus of points for which a shot would match the time difference between the relative arrival at pairs of microphones. The loci represent mathematical hyperbolas. as shown in Figure 5.

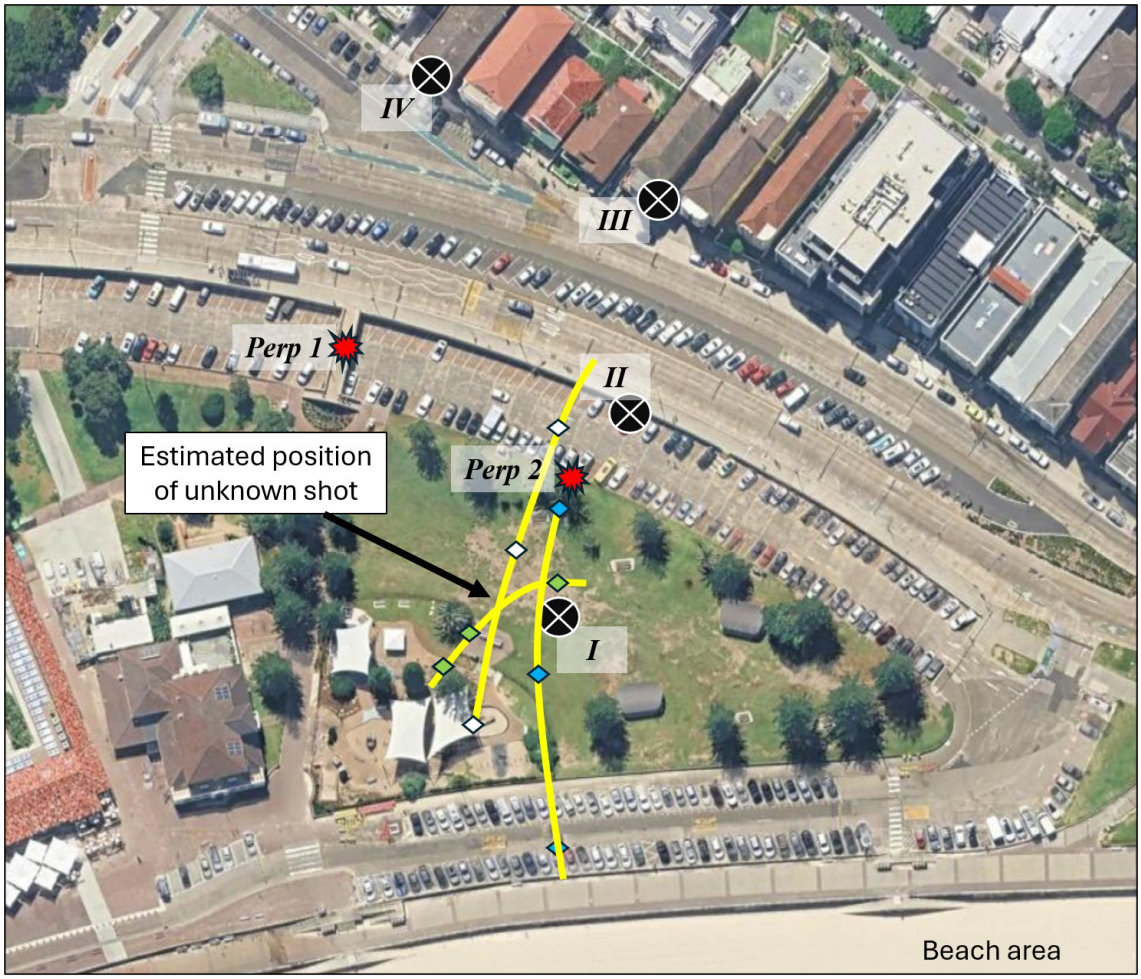


Figure 5: Indication of multilateration hyperbolas for microphone time of arrival differences using microphones *I*, *II*, *III*, and *IV*, showing estimated position for unknown shot.

The estimated hyperbolas are not perfectly aligned, which is typical of the ad hoc methodology with uncertainties about the speed of sound, signal timing, microphone and firearm positions. Nonetheless, the position estimate for the *unknown* gunshot sound shown in Figure 5 matches the location witnesses reported seeing a law enforcement officer at the time of the response, and is at least indicative that the first law enforcement shot came from the grassy park area near perpetrator #2, not from the location west of perpetrator #1 where later law enforcement shots were positively identified in the various videos. Further reporting and results of a formal investigation may provide additional information to determine if the acoustical estimate is consistent with other reports and physical evidence from the terrorist incident.

6 Conclusion

The sounds of gunfire from incidents around the world are increasingly captured by home surveillance systems and bystander mobile phone video, either inadvertently or deliberately. Unlike purpose-built gunshot detection systems that have precise location and timing consistency for multiple microphone sensors, non-synchronized user-generated recordings require steps to try to determine the relative location of each recording device, and the relative time synchronization of the original recordings. The methodology requires identifying a sound found in each of the available recordings that is from a known sound source and location. The known sound can then be used as a synchronization time stamp, with adjustments made to account for the time required for sound to propagate from the known source position to the estimated microphone positions.

The two recent examples demonstrate the practical application of the methodology to help answer investigative questions through audio forensic analysis. The results can assist with investigative journalist inquiries and in the courtroom.

The practical methods to localize the recording devices and to synchronize the audio recordings can be helpful to create ad hoc multilateration networks, or at least to discern the relative geometric relationship between the recording devices and the likely position of the firearm. Application of these methods must include an understanding of the strengths and limitations of audio evidence interpretation due to the typically noisy, distorted, and perceptually-encoded form of the data, and unknown details regarding the shooting scene, meteorological conditions, and other physical factors. Thus, while ideally it would be possible to develop a general synchronization procedure applicable in all cases, the unique aspects and uncertainties of user-generated recordings currently require ad hoc methods.

References

[1] R.C. Maher, "Forensic interpretation and processing of user generated audio recordings," Preprint 10419, Proc. 149th Audio Engineering Society Convention, New York, NY (2020).

[2] B.F. Miller, F.A. Robertson, and R.C. Maher, "Forensic handling of user generated audio recordings," Preprint 10515, Proc. 151st Audio Engineering Society Convention (2021).

[3] R.C. Maher and E.R. Hoerr, "Audio Forensic Gunshot Analysis and Multilateration," Preprint 10100, Proc. 145th Audio Engineering Society Convention, New York, NY (2018).

[4] R. Calhoun and S. Lamkin, "Determining the source location of gunshots from digital recordings," Express Paper 18, Proc. 153rd Audio Engineering Society Convention, New York, NY (2022).

[5] B.E. Koenig, S.M. Hoffman, H. Nakasone, and S.D. Beck, "Signal Convolution of Recorded Freefield Gunshot Sounds," J. Audio Eng. Soc., vol. 46, no. 7/8, pp. 634-653 (1998).

[6] R.C. Maher, "Acoustical Characterization of Gunshots," Proceedings IEEE SAFE 2007: Workshop on Signal Processing Applications for Public Security and Forensics, Washington, DC, pp. 109-113 (2007). https://doi.org/10.1109/IEEECONF12259.2007.4218954

[7] R.C. Maher, Principles of Forensic Audio Analysis, Springer Nature Switzerland (2018). https://doi.org/10.1007/978-3-319-99453-6

[8] D.R. Begault, S.D. Beck, and R.C. Maher, "Overview of forensic gunshot analysis techniques," elib 20475, Proceedings 2019 Audio Eng. Soc. International Conference on Audio Forensics, Porto, Portugal (2019).

[9] R.C. Maher, "Examining tell-tale sounds in forensic gunshot recordings," elib 22634, AES 8th International Conference on Audio Forensics, Denver, CO (2024).

[10] R.C. Maher and S.R. Shaw, "Gunshot recordings from digital voice recorders," Proc. Audio Engineering Society 54th Conference, Audio Forensics—Techniques, Technologies, and Practice, London, UK (2014).

[11] R.C. Maher, "Challenges of Audio Forensic Evaluation from Personal Recording Devices," Preprint 9897, Proc. 143rd Audio Engineering Society Convention, New York, NY (2017).

[12] ShotSpotter Data Sheet, SoundThinking, Inc. https://www.soundthinking.com/wp-content/uploads/2023/03/2023-xx-xx-ShotSpotter-Datasheet.pdf (2023).

[13] R.C. Maher, "Audio forensic interpretation of the Trump rally assassination attempt (July 13, 2024)," J. Audio Eng. Soc., vol. 73, no. 11, pp. 782–790 (2025). https://doi.org/10.17743/jaes.2022.0230.

[14] I. Ting, A. Palmer, T. Brettell, D. Welch, L. Sonnenschein, and K. Shatoba, "Dozens of synced videos capture humanity amid horror of Bondi attack," *Australian Broadcasting Corporation* (2026). https://www.abc.net.au/news/2026-02-28/synchronised-video-timeline-bondi-terror-attack/106282432.

[15] T. Spencer, "What type of weapons did the Bondi Beach attackers use?" *BBC* (2025). https://www.bbc.com/news/live/cq5qqq6gv6yt?post=asset%3A0155ee15-a357-43df-a934-4cd698ee9b82.

[16] I. Ting, "How sound helped create a synchronised video timeline of the Bondi terror attack," *Australian Broadcasting Corporation* (2026) https://www.abc.net.au/news/2026-02-28/how-sound-helped-sync-videos-of-the-bondi-terror-attack/106282450.