

Sound of Guns: Digital Forensics of Gun Audio Samples meets Artificial Intelligence

Simone Raponi[†], Isra M. Ali^{*}, Gabriele Oligeri[‡]

Division of Information and Computing Technology,
College of Science and Engineering,
Hamad Bin Khalifa University, Doha, Qatar

Email: [†]sraponi@mail.hbku.edu.qa, ^{*}isali@mail.hbku.edu.qa, [‡]goligeri@hbku.edu.qa

Abstract—Classifying a weapon based on its muzzle blast is a challenging task that has significant applications in various security and military fields. Most of the existing works rely on ad-hoc deployment of spatially diverse microphone sensors to capture multiple replicas of the same gunshot, which enables accurate detection and identification of the acoustic source. However, carefully controlled setups are difficult to obtain in scenarios such as crime scene forensics, making the aforementioned techniques inapplicable and impractical.

We introduce a novel technique that requires zero knowledge about the recording setup and is completely agnostic to the relative positions of both the microphone and shooter. Our solution can identify the category, caliber, and model of the gun, reaching over 90% accuracy on a dataset composed of 3655 samples that are extracted from YouTube videos. Our results demonstrate the effectiveness and efficiency of applying Convolutional Neural Network (CNN) in gunshot classification eliminating the need for an ad-hoc setup while significantly improving the classification performance.

I. INTRODUCTION

Gunshot analysis have received significant attention from both the military and scientific communities. Acoustic analysis of gunshots can provide useful information, such as the position of the shooter, the projectile trajectory, the caliber of the gun, and the gun model. Although acoustical evidence may significantly contribute to audio forensic reconstruction and analysis, the forensic analysis of gunshots is characterized by many challenges due to the broadcast and noisy nature of the acoustic channel.

Consider a scenario where a microphone is deployed in a close neighborhood to the shooter. The recorded audio sample can be significantly affected by the environmental surroundings, such as trees, foliage, and buildings, which attenuate and reflect the main component of the shock wave. The resulting audio sample may feature different echoes of the gunshot that are characterized by different attenuation factors as a function of their paths. This naive approach is impractical, which motivated the development of more complex ad-hoc acoustic data acquisition strategies over the last decade. To mitigate echoes and overcome the intrinsic lack of information, that the aforementioned scenario suffers from, additional microphones are deployed. The comparison

of multiple replicas of the same gunshot enables *shooter localization* and *weapon identification*. The physical characteristics of acoustic propagation can be exploited to infer the position of the shooter and the category of the gun. Multiple spatially diverse acoustic sources enable the estimation of Angle of Arrival (AoA), Time of Arrival (ToA), and Time Difference of Arrival (TDoA). The obtained recordings can be modeled by geometrical acoustics that enable the localization of the shooter. Furthermore, multiple replicas of the same acoustic source allow to filter out echoes and background noise affecting a subset of the deployed microphones, thus enabling a deep characterization for both the time and frequency domains.

Acoustic acquisition via Wireless Sensor Network (WSN) requires a specialized infrastructure overlay to enable sensor communication, data processing, and computation distribution. Solutions that rely on the spatial diversity provided by the WSN introduce several types of burdens. Firstly, each soldier has to carry a wearable device equipped with a microphone and other sensors, such as a compass, to collect meaningful information about AoA, ToA, and TDoA. Secondly, in a military scenario, the WSN should feature a jamming-resistant communication protocol and non-interfering radio channels. Both assumptions are difficult to achieve given the resource constraints of WSNs in terms of CPU, battery, and memory. In most cases, WSNs cannot afford the computational burden of multimedia processing. Therefore, the captured data should be first off-loaded to a remote server, then downloaded and distributed again. This represents a challenge from the connectivity perspective since, in many cases, military WSNs are unattended or provided with a discontinued link to the control center.

In this work, we do not rely on ad-hoc acquisition setups, but we exploit publicly available audio recordings of gunshots, considering their temporal and spectral representations. Spectral analysis of sound has been adopted in many contexts to detect and identify recurrent patterns. In particular, the combination of time-frequency decomposition of audio samples with Convolutional Neural Network (CNN)s provides promising performance in detecting recurrent patterns. The CNN is trained over several “images” constituted by a three-

dimensional representation of time, frequency, and amplitude. The result is a robust solution that can “recognize” the same sound by cross-matching similar images.

Contribution. We propose an inexpensive solution that is able to detect and identify gunshots without resorting to any ad-hoc infrastructure. Contrary to other studies, our solution requires only an audio sample of a gunshot that can be easily obtained by any commercially available microphone. Our approach is agnostic to the microphone position with respect to the shooter, and it does not require multiple spatially different replicas of the gunshot; we consider recordings from mono-channel setups with different sample rates. We proved the effectiveness of our solution by considering 3655 samples of gunshots constituted by 30 pistols, 18 rifles, and 11 shotguns for a total of 7 different calibers. The proposed approach guarantees an accuracy higher than 90% for all of the considered cases, namely, the category, model and caliber of the gun.

Paper organization. The remainder of this paper is organized as follows. Section II summarizes recent contributions in the field of weapon classification. Section III introduces the background concepts related to frequency domain analysis, CNNs, and acoustic characteristics of gunshots. Section IV describes our dataset and Section V discusses the the dataset generation process. The neural network architecture is presented in Section VI. Section VII shows the performance of our solution. Finally, Section VIII draws some concluding remarks.

II. RELATED WORK

Firearm classification based on the acoustic evidence generated by its discharge has long been investigated, but not extensively studied in the literature. Proposed solutions vary in many aspects, including the source of acoustic data, the type of analysis applied, the type of features extracted, and the application area. Table I summarizes prior studies, that provide gunshot classification and firearm identification, according to these aspects.

The source of the data is characterized by the type, the quality, and the environmental conditions of the deployed audio recording setup, which defines the amount of information that can be leveraged for classification. Most of the gunshot recordings used in the literature are either obtained under carefully controlled conditions, where a distributed set of microphone sensors are deployed [1], [2], [3], or extracted from a conventional recording device in less controlled environments [4], [5], [6], [7], [8].

In the former case, where a Wireless Acoustic Sensor Network (WASN) is deployed, spatial information can be obtained by performing array processing and triangulation techniques. Direction of Arrival (DoA) and ToA estimation methods are applied to the obtained audio signals to determine the projectile speed and trajectory, as well as to infer the position of the shooter. Such information may also provide discriminant features, such as the bullet speed [1], that can be used to identify the firearm category. Furthermore, the

distributed nature of the recording setup provides spatial diversity, where multiple acoustic observations from different locations of the same gunshot are obtained, which can be leveraged to increase the classification accuracy. Sánchez-Hevia et al. [3] exploited this feature and proposed a multi-observation weapon classification system that leverages various classifier ensembles to enhance classic decision fusion techniques. Each node in the sensor network produces a classification decision using Least Squares Linear Discriminant Analysis (LS-LDA). The decisions are later fused using a Maximum Likelihood-based fusion rule that weights the decision of each node based on its location.

The main constraint induced by this type of analysis is the requirement of spatial information, which can only be obtained by deploying a distributed sensor network. Therefore, limiting the applicability of gunshot detection and firearm classification to a carefully controlled recording setup only. Consequently, various pattern recognition approaches were proposed that identify the firearm category in the absence of spatial information. The most used classifiers for firearm identification are Gaussian Mixtures Model (GMM) [4], [5], [6] and Hidden Markov Model (HMM) [7], [8].

Most of these approaches can be described as frame-based feature classification approaches [4], [5], [6], [8], where the time-domain acoustic signal is subdivided into a sequence of short-time windowed frames. From each frame, a set of predetermined features is extracted and used for gunshot classification. The most common extracted features are statistical measures of the spectrum and intensity of the signal, in addition to perceptual features such as Mel-Frequency Cepstrum Coefficient (MFCC) or Perceptual Linear Prediction Coefficients (PLP). Temporal features, such as energy and Zero Crossing Rate (ZCR), are also used, but only in conjunction with spectral or perceptual features.

Morton et al. [7] proposed an alternative classification approach that does not rely on frame-based features aiming to eliminate the dependency on performance-driven parameters, which are often optimized over a finite training set. They proposed modeling each firearm category as an HMM with AutoRegressive (AR) source densities using non-parametric Bayesian priors to allow automated model order selection. The AR defines a set of energy and spectral characteristics of the captured gunshot, while the HMM identifies the transitions of these states.

The aforementioned techniques may perform adequately in matched experimental conditions, however, their effectiveness could reduce significantly when capture conditions vary in challenging unstructured environments, where noise and distortion are present. Although Khan et al. [4] addressed this problem by using an exemplary embedding approach to bridge between varying recording conditions, the achieved classification accuracy is relatively low (i.e., 60-72%). The authors used a dataset of 100 gunshot samples obtained from 20 different firearm models, where each model is represented by 5 to 15 gunshot samples. The different conditions included in their experiments were simulated, namely, “Room Reverb”,

TABLE I: Prior Gunshot Classification Approaches

Name	Technique/Classifier	Features	Dataset	Result	Varying conditions	No Ad-Hoc Setup
[5]	Hierarchical GMM classification	Cepstral	50-100 gunshots of 10 gun types	90% (category) 85% (caliber)	✗	✓
[4]	Exemplar embedding using hierarchical GMM classification	Cepstral	100 gunshots of 20 gun types	95-100% (category) 60-72% (model)	✓	✗
[1]	DoA & ToA estimation	Projectile speed	194 gunshots of 4 gun types	86% (caliber)	-	✗
[7]	HMM using Non-parametric Bayesian techniques	None	~46 gunshots of 5 gun types	95.65% (model)	✗	✓
[6]	Hierarchical GMM classification	Cepstral & Temporal	230 gunshots of 5 gun types	96.29% (category)	✗	✓
[8]	HMM classification & Viterbi based decoding	Spectral & Temporal	372 gunshots of 4 handgun types	80% (model)	-	✓
[3]	LS-LDA & Maximum Likelihood decision fusion	Cepstral, Spectral & Temporal	840 gunshots of 14 gun types	94.1% (model)	✓	✗
Our solution	Convolutional Neural Networks	Spectral & Temporal	3655 gunshots of 59 gun types (pistols, rifles, and shotguns)	90% for Category, Caliber, and Model	✓	✓

“Concert Reverb”, and “Doppler Effect”, which may not match real-life environmental conditions and do not include directional variations. Furthermore, their approach assumes prior knowledge of the recording conditions which is not always possible, especially in audio forensic reconstruction analysis.

Our solution, being the only one considering varying environment conditions and not requiring an ad-hoc setup, outperforms the state of the art studies in terms of dataset richness, including the number of gunshots samples and range of weapon models, reaching 90% accuracy.

III. BACKGROUND

A. Spectrogram

A spectrogram is one of the most widely adopted visual representations of the frequencies spectrum of a signal over time. Being defined as an intensity plot of the Short-Time Fourier Transform (STFT) magnitude, a spectrogram is usually portrayed as a bi-dimensional graph, where one axis (usually the x-axis) represents time and the other axis (usually the y-axis) represents frequencies. An example of spectrogram is depicted in Fig. 1. Each intersection between time and frequency is assigned a color that refers to the Power Spectral Density (PSD) of that specific frequency at that particular time, which is considered a third dimension of the graph. To compute the spectrogram of a signal y , the signal is divided into shorter fixed-length segments y_1, \dots, y_n , and the Fourier transform is applied separately to each segment. The spectrogram describes the changes of the signal frequencies spectrum as a function of time. This implies that, if the time is discrete, the data to be transformed may be partitioned into overlapping frames. The STFT is applied to each of the frames and the result, consisting of both phase and magnitude for each

intersection between time and frequency, is stored in a matrix, as showed in Equation 2.

$$STFT\{y_n\}(m, \omega) = \sum_{n=-\infty}^{\infty} y_n \omega[n - m] e^{-j\omega n} \quad (1)$$

$$spectrogram\{y_n\}(m, \omega) = |STFT\{y_n\}(m, \omega)|^2 \quad (2)$$

The result consists of a bi-dimensional matrix that maps the audio frequencies to the time-localized points [9].

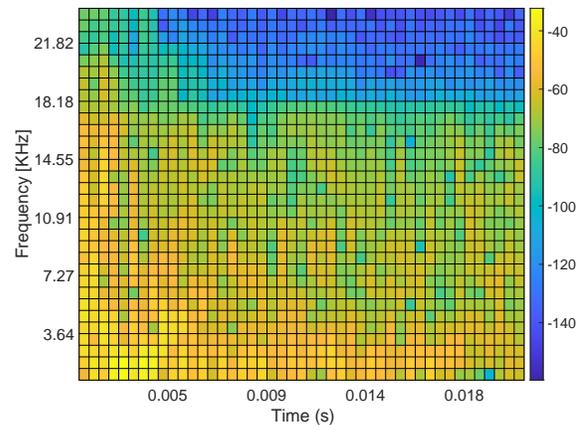


Fig. 1: Example of a gunshot spectrogram. The x-axis represents the time expressed in seconds, while the y-axis represents the frequencies expressed in kHz. The color represents the PSD at the given time-frequency.

The visual representation of audio traces through spectrograms have been extensively leveraged in the literature in

the context of audio classification [10], sound event classification [11], emotion recognition [12], human activity recognition [9], cross-modality feature learning [13], and gunshot classification [14].

B. Convolutional Neural Network

A CNN belongs to the class of deep neural networks that have one or more convolutional layers (i.e., layers that perform convolution operations). A convolution is a linear operation that consists of a slide of a parametric-sized filter over the input representation (usually a visual image). The application of the same filter to different overlapping filter-sized portions of the input generates a feature map. There are several types of filters, also known as operators. Each filter tries to identify a specific feature within the input representation. For example, the Sobel, the Prewitt, and the Canny operators highlight edges, the Harris and the Shi and Tomasi operators highlight corners, etc. One of the most powerful features of CNNs, that is also the reason behind their wide adoption, is the ability to automatically apply an extensive number of filters to the input representation in parallel, thus highlighting specific features in every part of the input image simultaneously.

CNNs can be seen as regularized versions, that discourage learning complex models, of multilayer perceptrons. While in multilayer perceptrons several fully connected layers are used—a layer is fully connected if all the neurons it is composed of are connected to all the neurons of the next layer, CNNs exploit a hierarchical structure that allows building complex patterns by using small and simple patterns.

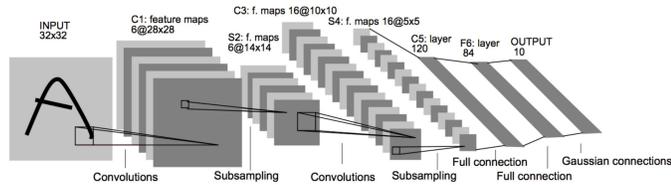


Fig. 2: Example of a CNN. LeNet-5 [15] is able to identify handwritten digits for zip code recognition in the postal service

Fig. 2 depicts a typical architecture of a CNN. The dimension of the input image (in this case representing a handwritten digit), keeps decreasing while going deeper in the neural network, while the number of filters, thus the features the architecture desires to highlight, increases. A CNN usually has three types of layers: (i) convolutional layers, to perform the convolution operations to the input, (ii) pooling layers, to discretize the input and reduce the number of learnable parameters; and (iii) fully connected layers, that are essentially feed-forward neural networks, usually placed at the end of the architecture. The goal of the fully connected layers is to hold the high-level features found during the convolutions and try to learn non-linear combinations of these features before assigning the input image a label. Details about these layers contextualized in our model are provided in Section VI-A.

One of the fundamental decisions to be taken when designing a CNN, or generically a neural network, concerns the

representation of the input data. Several input representations are available in the literature, each bringing its advantages and drawbacks. Although for visual images the choice is straightforward, for audio samples numerous alternatives are possible, including MFCC, raw digitized sample stream, machine discovered features, and hand-crafted features. Even if the best input representation to adopt is strongly dependant on the problem to solve, several studies in the literature show that feeding CNN with spectrograms is effective in many fields, including musical onset detection [16], human detection and activity classification [17], music classification [18], and other interesting activities [19].

C. Guns and Gunshots

Gunshots are the result of multiple acoustic events, namely, the *muzzle blast* created by the explosion inside the barrel and the *ballistic shockwave* that is generated by the supersonic projectile. These phenomena are the results of many characteristics and variables that eventually sum up and generate the acoustic blast, which include the firearm type, model, barrel length, ammunition type, powder quantity, weight and shape of the projectile, and possibly others. The aim of this work is to estimate at what extent it is possible to use a gunshot as a unique fingerprint that uniquely identifies one or more of the aforementioned variables. Figure 3 summarizes the most important characteristics affecting the acoustic blast generated by a gun.

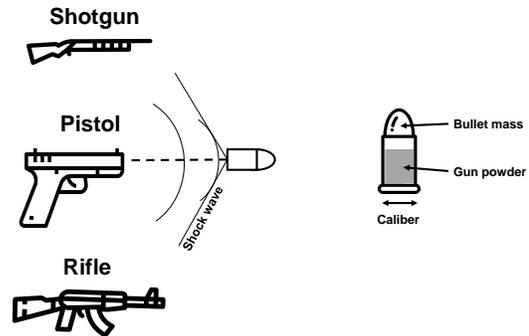


Fig. 3: Variables taken into account in our analysis: category of firearm, caliber, and gun model.

Our observation is that different configurations of the aforementioned parameters may lead to unique gunshot patterns that can be detected by analyzing the frequency-time decomposition of the gunshot blast. In the next sections, we demonstrate how Convolutional Neural Networks (CNNs) can be effectively used to detect these patterns, thus uniquely identifying the category of gun, the caliber, and finally, the model of the gun.

IV. DATASET DESCRIPTION

Table II shows the dataset considered in this work. We collected the samples from several YouTube videos, such as *C4Defense*, *hickok45*, *EmanuelRJSniper*, *mixup98*, *OneGear*,

and *ReloaderJoe*. Our choice of guns takes into account two main aspects: the *Category of Guns* and the *Caliber*.

Category of guns. We considered 30 different *pistols*, 18 *rifles*, and 11 *shotguns*. As for pistols, we considered 22 *revolvers* and 10 *semiautomatic*.

Caliber. We took into account the most popular calibers in U.S. and world-wide [20], [21], such as 9mm and .45acp for automatic pistols, .44M and .357M for revolvers, 7.62x39 and 5.56NATO for rifles, and 12 gauge caliber for shotguns.

A. Muzzle blast: preliminary considerations

When a gun is fired, there are two distinct acoustic phenomena, the *muzzle blast* and the *ballistic shockwave* [22]. The latter is generated by the bullet that compresses the air in front of itself creating a sonic boom that propagates with a shape of a cone where the vertex is the bullet itself. Conversely, the muzzle blast is a high energy acoustic signal originated by the gun’s muzzle with a spherical wavefront, propagating at the speed of sound, and with center the muzzle of the gun. The ballistic shockwave is a very important source of information to locate a sniper in an open field [23], [24]. However, to achieve that, the ballistic shockwave has to be sampled from different locations requiring an array of microphones. The ballistic shockwave cannot be observed for subsonic projectiles such as those used in shotguns and pistols.

Given the aforementioned considerations, we focus on the muzzle blast and the echoes associated with it. In the following, we discuss and highlight three critical parameters that have to be carefully set in order to maximize the detection performance of a neural network: (i) *muzzle blast duration*, (ii) *number of frequency bins*, and (iii) the *number of time slots*.

Figure 4 shows the acoustic signal amplitude recorded from a Beretta PX4 Storm, 9mm. The muzzle blast lasts for a few milliseconds (up to 5ms in the figure), depending on the model of gun and caliber. We also observe some echo effects (Echo 1, Echo 2, and Echo 3) at 10ms, 22ms, and 63ms due to reflections of the sound from obstacles around the shooter. We highlight that this is consistent with previous findings from other studies [22], while the *muzzle blast duration* will be a critical parameter from the analysis carried out in this work.

Figure 5 shows the PSD as a function of time and frequency (spectrogram) associated with the muzzle blast in Fig. 4. We consider both the bi-dimensional and the three-dimensional representation of the spectrogram. We observe that the muzzle blast (time less than 5ms) takes all the frequency components between 0 and 24KHz with a significant power spanning between -30dB (lower frequencies) and -80dB (higher frequencies). As soon as the blast finishes, the echoes take the frequencies less than 18KHz with a decreasing power between -40dB and -60dB. The aforementioned spectrogram components constitute the input for the training process of our neural network.

We identify two more critical parameters affecting our algorithm performance: the *number of frequency bins* and the *number of time slots*. For our analysis, we adopted the

TABLE II: Dataset: Gun Model, Caliber, and number of extracted samples.

Gun Model	Caliber	Category	N. Samples	ID
Glock 45	9mm	Pistol	154	24
Beretta 98 FS	9mm	Pistol	29	2
Beretta 92 FS	9mm	Pistol	361	11
Beretta PX4 Storm	9mm	Pistol	164	12
Glock 21	.45acp	Pistol	145	22
M&P Shield	.45acp	Pistol	77	27
Colt 1911	.45acp	Pistol	81	34
Walther PPQ	.45acp	Pistol	115	57
Glock 30S	.45acp	Pistol	103	23
S&W 629 4-inch	.44M	Pistol	97	5
S&W 629 TrailBoss	.44M	Pistol	62	6
S&W 629 Performance Center	.44M	Pistol	42	7
S&W 629-8	.44M	Pistol	48	47
S&W 69	.44M	Pistol	35	49
S&W 69 2.75-inch	.44M	Pistol	35	50
Charter Arms .44 Special Bulldog	.44M	Pistol	40	15
S&W Model 29 Dirty Harry	.44M	Pistol	42	21
S&W Model 29 4-inch	.44M	Pistol	51	32
Ruger Redhawk Big Game Hunt	.44M	Pistol	23	42
Ruger Super Black Hawk	.44M	Pistol	25	44
Ruger Red Hawk 8-shots	.357M	Pistol	48	41
S&W 357Magnum	.357M	Pistol	73	4
Chiappa Rhino	.357M	Pistol	47	16
Coonan 1911	.357M	Pistol	56	17
Ruger GP100 Match Champion	.357M	Pistol	36	38
Ruger SP101	.357M	Pistol	45	43
S&W Model 19 3-inch	.357M	Pistol	42	45
S&W Model 27	.357M	Pistol	42	46
S&W Model 66	.357M	Pistol	54	48
Dan Wesson Revolver	.357M	Pistol	30	19
CZ Bren 2 MS	7.62x39	Rifle	49	1
PWS MK107	7.62x39	Rifle	73	3
CZ 527	7.62x39	Rifle	38	13
Century Arms C39 AK-47	7.62x39	Rifle	37	14
Maadi AK47	7.62x39	Rifle	60	28
Micro Draco AK47 Pistol	7.62x39	Rifle	44	29
N-PAP AK	7.62x39	Rifle	46	33
Ruger American Ranch Rifle	7.62x39	Rifle	38	37
Ruger Mini 30	7.62x39	Rifle	51	40
SKS	7.62x39	Rifle	26	59
Daniel Defense M4 A1 SOCOM	5.56 NATO	Rifle	68	20
Ruger AR	5.56 NATO	Rifle	78	36
Ruger Mini-14	5.56 NATO	Rifle	60	39
Ruger AR556 MPR	5.56 NATO	Rifle	30	35
SIG 556 Classic SWAT Model	5.56 NATO	Rifle	72	51
Springfield Armory Saint	5.56 NATO	Rifle	68	54
Tactical Edge Warfighter	5.56 NATO	Rifle	49	56
M&P 15 Sport II	5.56 NATO	Rifle	84	26
Benelli M2 SBS	12	Shotgun	44	8
Benelli M4	12	Shotgun	45	9
Benelli Nova	12	Shotgun	36	10
DP-12	12	Shotgun	74	18
Kel-Tec SG12	12	Shotgun	52	25
Winchester Model 12	12	Shotgun	46	31
SRM 1216	12	Shotgun	34	52
Serbu Super Shorty	12	Shotgun	33	53
Standard Manufacturing SKO Shorty	12	Shotgun	47	55
Winchester SXP Defender	12	Shotgun	36	58
Winchester Model 12 SlugFest	12	Shotgun	35	30

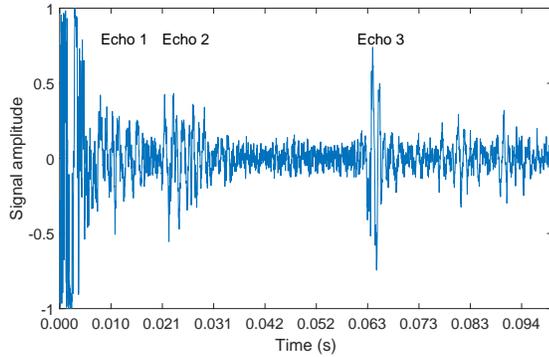


Fig. 4: Acoustic signal amplitude of a muzzle blast for a Beretta PX4 Storm, 9mm.

spectrogram function of MATLAB-R2019b considering as input the acoustic sample (0.1 seconds from the beginning of the blast), a window of size $w = 44$ to divide the signal into segments and performing windowing according to the Hann function, $no = \lfloor w/2 \rfloor$ as the number of overlapping samples between adjacent segments, $fl = 65$ as the FFT length, and $fs = 48000$ as the number of samples per seconds acquired by the microphone. Assuming the previous parameters, the *spectrogram* function returns the PSD at $\lfloor fl/2 + 1 \rfloor$ frequencies and $\lfloor \frac{length(x)-no}{w-no} \rfloor$ time bins, where x is the vector of the acoustic samples. For instance, in the previous example, the frequency range (0 to 24KHz) has been divided into 65 bins, while the time has been divided into 46 slots.

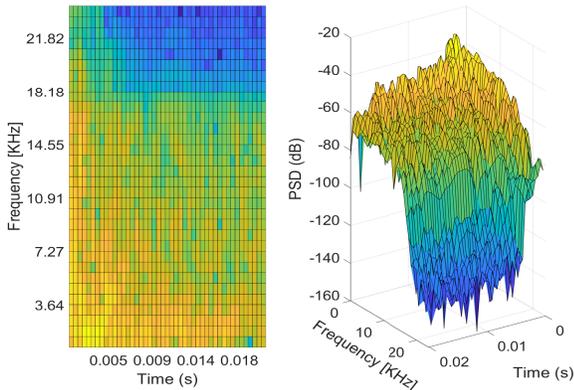


Fig. 5: Spectrogram of a muzzle blast: bi-dimensional and three-dimensional PSD of a muzzle blast (Beretta PX4 Storm, 9mm) as a function of time and frequency.

B. Quality of the audio samples

In the following, we provide a quantitative analysis of the quality of the collected audio samples. As a quality metric, we consider the Signal-to-Noise Ratio (SNR) computed on each muzzle blast from the actual starting of the blast for a period of 400ms. For each audio sample, we consider a pre-defined reference noise pattern constituted by random samples of amplitude 0.1, i.e., one-tenth of the maximum signal amplitude taken by the microphone. The previous sound pressure is equivalent to a classical background noise that can be sampled from an outdoor environment characterized by a gentle wind. Figure 6 shows the probability distribution function associated with the SNR computed as described before. The overall audio quality is very high since the muzzle blast is +20dB higher than the reference noise pattern. We observe that even the echoes can be easily identified from the noise reference.

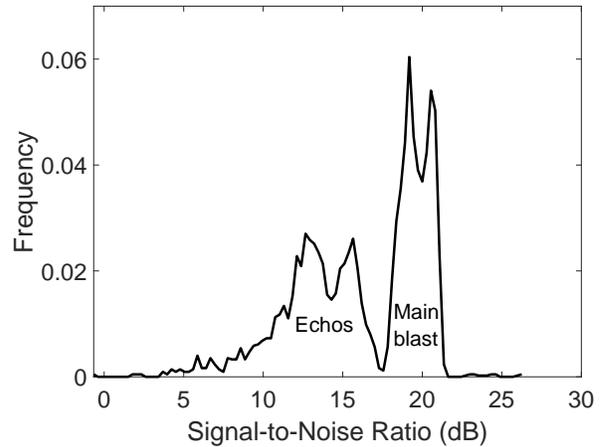


Fig. 6: Sound quality analysis: SNR of main blasts and their associated echos assuming a reference random noise of amplitude one tenth of the microphone saturation threshold.

V. DATASET GENERATION

We generated a dataset of 3655 samples extracted from videos found on YouTube. Each of the collected audio samples has a sample rate of either 48000 or 44100 samples per second. Generating a dataset of gunshots extracted from YouTube videos involves the following steps:

- **Audio extraction.** We performed the audio extraction (MP3 format) from the selected videos using *youtubedl* [25] and *ffmpeg* [26] tools.
- **Abrupt change detection.** A preliminary filtering is performed by identifying abrupt changes in the audio signal.
- **Gunshot detection.** Gunshots are detected among blasts by relying on a Support Vector Machine (SVM) learning algorithm.

In the following, we describe the procedure of automatically extracting gunshots from an audio trace focusing on *Blast detection* and *Gunshot detection*.

A. Identification of abrupt changes in an audio trace

To detect *abrupt changes* in an audio trace, we computed the variance over a sliding window of 5ms, equivalent to either 220 or 240 samples depending on the quality of the audio trace, i.e., 44100 or 48000 samples per second, respectively. Subsequently, we searched for the peaks adopting windows of size 0.3 seconds and a minimum peak prominence of 0.3. Fig. 7 shows the three computation stages from the sound pressure to the blast sequences that are passing by the moving window averaging. This figure refers to two sound chunks extracted from an audio trace, where the first part (i.e., $0 \leq t \leq 5.5$ seconds) is a sequence of gunshots, while the second part (i.e., $t > 5$ seconds) is mainly constituted by voice. We stress that the main aim of this part is to detect abrupt changes in the sound pressure, while subsequently we will show how gunshots are identified.

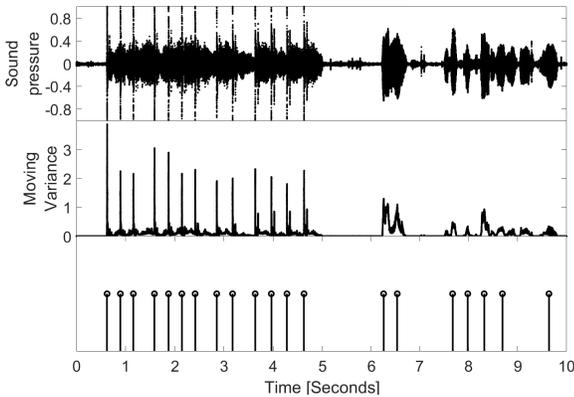


Fig. 7: Detection of abrupt changes in audio traces: from sound pressure to abrupt change detection by computing moving variance and peak detection.

B. Gunshot detection

Gunshot detection is performed via a *human-assisted supervised learning* approach. The intention is to have a growing training set of actual gunshots that is supervised by the user. The user checks for both false positives and false negatives by listening to the newly generated samples in the training set. Figure 8 shows the training, validation, and testing procedures. We assume that the training set is populated with an initial dataset of actual gunshots that have been manually selected. In our case, we started from an initial dataset of 10 gunshot samples only. At each cycle, a new model is trained with the current training set (Step 1 in Fig. 8). Subsequently (Step 2 in Fig. 8), new samples are selected from the list that is

generated by the procedure presented in Section V-A. Finally, the generated samples are tested with the current training set. The output is assessed by the supervisor (Step 3 in Fig. 8), and the verified samples are added to the training set (Step 4 in Fig. 8).

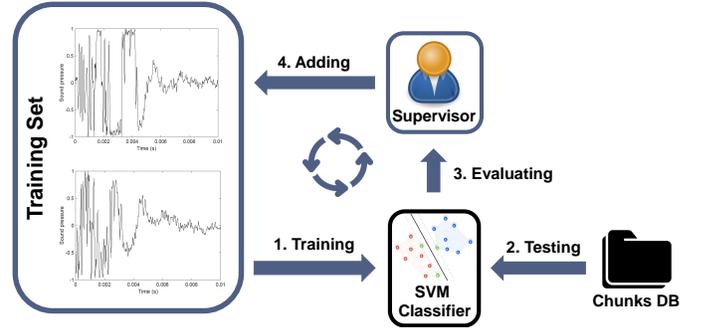


Fig. 8: Gunshot detection via human-assisted supervised learning. The SVM classifier is trained by verified samples of gunshots. When new gunshot samples are tested, the output of the classifier is verified (by the user) and added to the Training Set.

Classification performance. To assess the quality of the classification procedure, we considered 6 additional videos (V1, ..., V6) downloaded from YouTube, which are not included in the training set. For each video, we detected the abrupt changes according to the procedure presented in Section V-A, and we executed the gunshot detection procedure presented in Fig. 8. As for the Training Set, we considered the one we generated from the samples found in Table II. Figure 9 shows the frequency of the similarity indexes provided by the SVM classifier for the Shot and No-Shot audio samples with red crosses and green circles, respectively. The similarity indexes were categorized into bin width of 10, where each cross/circle aggregates adjacent similarity indexes. Figure 9 represents the decision after one iteration of the procedure presented in Fig. 8. The decision Shot vs No-Shot is taken as a function of the threshold Thr , which has been empirically set to zero. We observed that 96% of the No-Shot samples feature a similarity index of -189.5 , while the remaining 4% are spread between -178.9 and -0.55 . There are no samples from the No-Shot class with a similarity index that is greater than 0. As for the Shot class, the samples are distributed between 0.41 and 275, with frequencies between 1% and 11%. Even in this case, we highlight that there are no samples from the Shot class with a similarity index that is less than 0.

To precisely assess the effectiveness of our solution, we manually checked all of the classified samples, namely, Shot vs No-Shot. Table III shows the result of our analysis. For each video, we report the number of detected abrupt changes (N), the threshold used by the SVM classifier ($Thr = 0$), True Positive (TP), False Positive (FP), True Negative (TN), False Negative (FN), the actual number of gunshots ($Actual$), and the overall accuracy of the detection algorithm. As previously stated, during our evaluation, we considered only one iteration

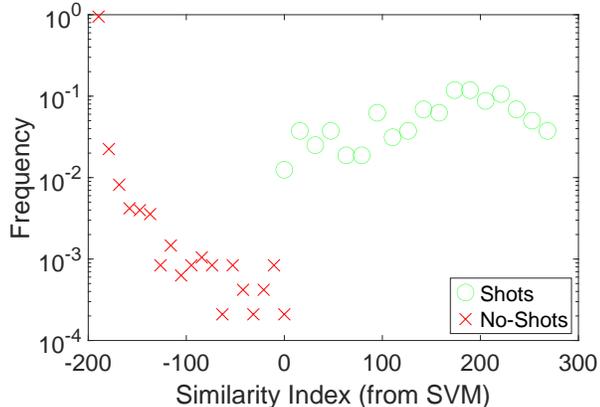


Fig. 9: Gunshot detection performance: frequency distribution of Shot and No-Shot samples as a function of their similarity index.

TABLE III: Shot detection performance considering 6 videos

	N	Thr	TP	FP	TN	FN	Actual	Accuracy
V1	643	0	24	2	617	0	24	0.99
V2	385	0	4	0	365	16	20	0.96
V3	79	0	5	0	74	0	5	1
V4	32	0	5	0	27	0	5	1
V5	1860	0	88	0	1772	0	88	1
V6	1932	0	31	0	1897	4	35	0.99

as depicted in Fig. 8. We would like to highlight that the proposed algorithm achieves the main purpose of generating a dataset of gunshot samples (i) in a fast and efficient way and (ii) with the minimum amount of false positives. The output of this phase will be the training set to be used by the CNN.

At this stage, we aim at minimizing the number of FP , which might bias the subsequent training process. We also aim at maximizing the process efficiency of creating a large dataset of gunshot samples. Therefore, the task of the supervisor resorts mainly to listening to a very few samples ($TP + FP$) despite the dataset N , in order to remove the FP , which are overall very few: only 2 out of 4931 samples. Conversely, we observe that our approach might lose some good samples ($FN = 16 + 4$). However, these samples do not affect the performance of our solution hence we consider them not important.

The above procedure has been applied to each audio sample found in Table II to generate a dataset of actual gunshot samples, that is, one dataset for each gun model.

VI. OVERALL ARCHITECTURE

Figure 10 depicts the overall architecture of our CNN consisting of five layers with weights: the first four are convolutional layers, while the last one is a fully connected layer. The output of the fully connected layer is fed to a 7-way softmax, that outputs the probability distribution over

the 7 class labels. The details of our architecture, including information about the layers and their learnable parameters, are reported in Appendix A.

Considering the dimension of the starting image and the need to give importance also to peripheral pixels, every convolution of our architecture makes use of padding to avoid losing information. By adding additional pixels to the border, every convolution layer outputs an image with the same number of pixels of the one fed into that layer. Furthermore, in our CNN architecture, we make use of a stride of 1 during convolutions, and a stride of 2 during the Max Pooling application. The stride is a critical hyperparameter in the context of CNN, as it allows to specify the number of cells by which filters (e.g., convolution filters, pooling filters) slides over the image. If the stride is equal to 2, the filter starts from the top left corner and moves over the image with jumps of 2 units at a time. By considering square filters (i.e., $f \times f$) and square initial images (i.e., $n \times n$), after having specified the dimension of the filters f , being them convolutional filters or pooling filters, the stride parameter s , the dimensions of the initial images n , and the padding p , it is possible to calculate the dimension of the square output image of a layer as:

$$\left\lfloor \frac{n + 2p - f}{s} + 1 \right\rfloor \quad (3)$$

Our choice to keep a unit stride during the convolutions and a stride equals to 2 during the pooling is guided by the intention of not losing information during convolution phases, while exploiting the pooling technique to summarize the features, thus reducing the input dimensionality.

The first convolutional layer filters the $36 \times 99 \times 1$ spectrogram image with 40 kernels of size $3 \times 3 \times 1$, without any stride and with 'same' padding. The second convolutional layer takes as input the normalized (40 channels) and pooled (3×3 max pooling, stride = 2) output of the first convolutional layer and filters it with 80 kernels of size $3 \times 3 \times 40$. The third convolutional layer takes as input the normalized (80 channels) and pooled (3×3 max pooling, stride = 2) output of the second convolutional layer and filters it with 160 kernels of size $3 \times 3 \times 80$. The fourth convolutional layer takes as input the normalized (160 channels) and pooled (3×3 max pooling, stride = 2) output of the third convolutional layer and filters it with additional 160 kernels of size $3 \times 3 \times 160$. The normalized (160 channels) and pooled (1×13 max pooling, stride = 2) output of the fourth convolutional layer is fed to a 7-neuron fully connected layer that, in turn, outputs the result to a 7-way softmax, that produces a distribution over the 7 class labels.

A. CNN Details

Activation Function. Our neural network relies on the Rectified Linear Units (ReLU) activation function [27] after each convolution. The ReLU activation function, whose operation is simplified in equation 4, outputs the maximum value between zero and the input value.

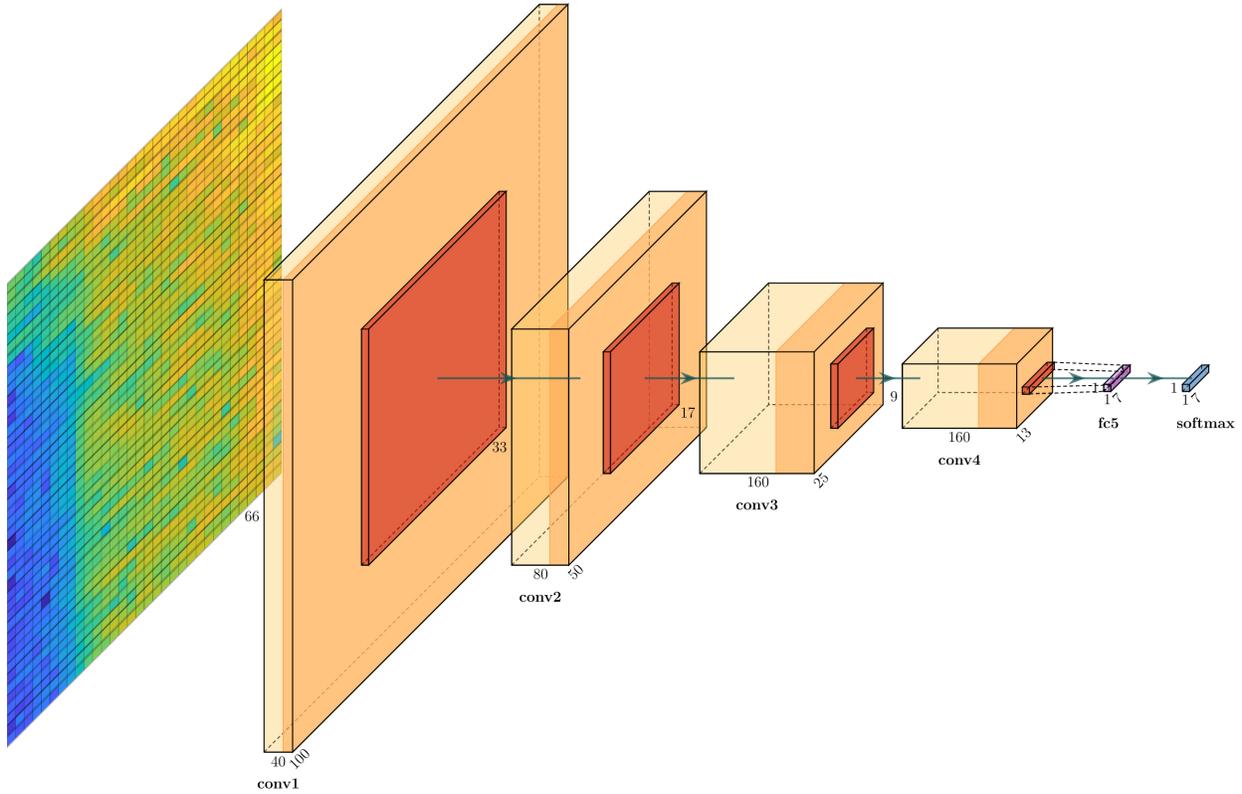


Fig. 10: Structure and details of our Convolutional Neural Network.

$$f(x) = \begin{cases} x & \text{if } x \geq 0, \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

Although the literature uses other variants (e.g., Tanh, Soft-Sign, Sigmoid), several studies show that ReLU outperforms the competitors in terms of performance [28], [29].

Regularization. Our neural network relies on Dropout [30] regularization to reduce the likelihood of overfitting. The Dropout regularization technique allows to randomly cut out units (together with their connections) from the neural network during the training phase with a given probability. This discourages neurons to rely on the presence of particular other neurons and forces them to find more robust features with different ones [29], thus reducing the probability of learning the training set by heart.

Normalization. Training neural network without normalization brings to the internal covariate shift phenomena, where the distribution of each layer's input change during training, thus requiring a more sophisticated tuning of the parameters. To mitigate this issue we add Batch Normalization layers after each convolution. The Batch Normalization technique [31] performs the normalization for each training mini-batch, allowing the usage of higher learning rates and reducing the need for a cherry-picking tuning of the parameters. As summarized in the equation 7, Batch Normalization normalizes the output of an activation layer by subtracting the mean and dividing by

the standard deviation of the batch.

Given a mini-batch $\beta = x_1, \dots, x_m$:

$$\mu_\beta \leftarrow \frac{1}{m} \sum_{i=1}^m x_i \quad (5)$$

$$\sigma_\beta^2 \leftarrow \frac{1}{m} \sum_{i=1}^m (x_i - \mu_\beta)^2 \quad (6)$$

$$\hat{x} \leftarrow \frac{x_i - \mu_\beta}{\sqrt{\sigma_\beta^2 + \epsilon}} \quad (7)$$

Although the Batch Normalization technique brings a slight regularization effect to the neural network, in some cases eliminating the need for Dropout [31], we find that the combined use of the Batch Normalization and Dropout aids generalization [29].

Discretization. The application of discretization techniques to an input representation consists of reducing its dimensionality to evaluate the features within the obtained, summarized sub-regions. This process allows to mitigate the overfitting of the training set and to reduce the number of parameters to be learned for the training, thus reducing the overall computational cost. To attain these benefits, in our architecture we use a Max Pooling sample-based discretization process layer after each activation layer. Max Pooling applies a max filter to non-overlapping sub-regions of the input feature map, whose dimension is dictated by the dimension of the filter. When

Max Pooling is applied, the passage of the moving filter onto a sub-region produces, as output, a value, consisting of the maximum value of that sub-region.

Output. As for the output layer, our neural network architecture relies on the commonly used softmax function. The softmax function, taking as input a vector of real numbers, produces a probability distribution proportional to the exponential of the input numbers. In detail, the input real numbers are mapped in a (0,1) interval that sums up to zero, thus allowing to treat the output provided by the softmax function as probabilities.

In general, given a vector of real numbers $v = (v_1, \dots, v_K) \in \mathbb{R}^K$, the standard unit softmax function $\sigma : \mathbb{R}^K \rightarrow \mathbb{R}^K$ is defined by:

$$\sigma(v)_i = \frac{e^{z_i}}{\sum_{j=1}^K e^{z_j}} \quad (8)$$

B. Learning Details

TABLE IV: Training Options of our network.

Option	Value
Optimizer	Adam
InitialLearnRate	$x * 10^{-4}, x \in [1, 3]$
MaxEpochs	50
MiniBatchSize	8
Shuffle	Every Epoch
Plots	Training Progress
Validation Data	Random 20% of the data
Validation Frequency	$\lfloor \text{training_set} \rfloor / \text{miniBatchSize}$

Table IV summarizes the training options of our network, that are detailed in the following.

Optimizer. An optimizer is defined as an algorithm (or a method) used to tune the parameters of a neural network with the goal of reducing the loss function. In our architecture, we rely on the Adaptive Moment Estimation (Adam) optimizer [32], an extensively adopted optimizer that inherits the advantages of both Root Mean Square Propagation (RMSProp) and Stochastic Gradient Descent (SGD) with momentum (i.e., SGD where each gradient update is a linear combination of the previous gradient updates) optimizers. From RMSProp it inherits the squared gradients to scale the learning rate, while from SGD with momentum it inherits the concept of the moving average of the gradients. An empirical analysis conducted in [32] shows that Adam outperforms the other optimizers, thus working better in practice. As recommended in the original paper (whose algorithm is reported below with our parameters), in our implementation we set to 0.9 the gradient decay factor β_1 , to 0.999 the squared gradient decay factor β_2 , and to 10^{-8} the denominator offset (to avoid divisions by zero), respectively. However, although the original paper recommends using an initial learning rate of 10^{-3} , we empirically found (relying on the grid search hyperparameter tuning technique) that setting this value to $x * 10^{-4}, x \in [1, 3]$ provides better results.

Algorithm 1: Adam Optimizer [32]

Require: α : Stepsize

Require: $f(\theta)$: Stochastic objective function

Require: θ_0 : Initial parameter vector

$m_0 \leftarrow 0$ (Initialization of the 1st moment vector)

$v_0 \leftarrow 0$ (Initialization of the 2nd moment vector)

$\beta_1 \leftarrow 0.9$ (Initialization of the gradient decay factor)

$\beta_2 \leftarrow 0.999$ (Initialization of the squared gradient decay factor)

$t \leftarrow 0$ (Initialization of the timestep)

$\epsilon \leftarrow 10^{-8}$ (Initialization of the denominator offset)

$\alpha \leftarrow x \cdot 10^{-4}, x \in [1, 3]$ (Initialization of the learning rate)

while θ is not converged **do**

$t = t + 1$ (Increment the timestep)

$g_t \leftarrow \nabla_{\theta} f_t(\theta_{t-1})$ (Gradients w.r.t. stochastic objective at timestep t)

$m_t \leftarrow \beta_1 \cdot m_{t-1} + (1 - \beta_1) \cdot g_t$ (Update biased first moment estimate)

$v_t \leftarrow \beta_2 \cdot v_{t-1} + (1 - \beta_2) \cdot g_t^2$ (Update biased second raw moment estimate)

$\hat{m}_t \leftarrow \frac{m_t}{1 - \beta_1^t}$ (Compute bias-corrected first moment estimate)

$\hat{v}_t \leftarrow \frac{v_t}{1 - \beta_2^t}$ (Compute bias-corrected second raw moment estimate)

$\theta_t \leftarrow \theta_{t-1} - \alpha \cdot \frac{\hat{m}_t}{\sqrt{\hat{v}_t + \epsilon}}$ (Update Parameters)

end

return θ_t (Parameters)

Number of Epochs. An epoch is defined as a single pass through the training set, i.e., 1 forward pass and 1 backward pass for all the training samples. We empirically set as 50 the max number of epochs, since each of the subsequent epochs does not bring any benefit to our model learning.

Mini-Batch Size. Using mini-batch that consists of processing small subsets of training samples in every iteration, instead of processing them all together. The choice of mini-batch size (e.g., the number of training samples to process) does not affect the performance of the model in terms of accuracy, but affects the resource required during the training process. A larger mini-batch size requires more memory and takes more time per epoch, but allows to better optimize the vectorization (i.e., the linear transformation of a matrix into a column vector), while a smaller mini-batch size requires less memory but loses the speed-up given by vectorization. In our model, we set the mini-batch size to 8, to better optimize the resources of our server.

Shuffle. The “shuffle” option allows shuffling the order of which training samples are fed to the model, with the goal of reducing variance, thus reducing overfitting. Shuffling the training samples becomes crucial in case mini-batches are used, due to the need to avoid having batches containing highly correlated samples that would slow down (or, in many cases, compromise) the performance of the model. In our model, we

shuffle the training data before each training epoch, as well as the validation data before each validation.

Plot. The “plot” option in Matlab provides several information to be taken into account during the training process. Information include, but are not limited to, the mini-batch training loss and accuracy, the smoothed training loss and accuracy (i.e., the result of the application of a smoothing algorithm to the training accuracy), the validation loss and accuracy, hardware resources, etc.

Validation Data. The validation data, also known as validation set, refers to a subset of samples separated from the training set, that the model will rely on to evaluate the effectiveness of its training. In our case, by following the 80/20 rule, the validation set is represented by 20% of the whole dataset.

Validation Frequency. The validation frequency represents the number of iterations between evaluations of validation metrics. We empirically set this value to $\lfloor \frac{|training_set|}{miniBatchSize} \rfloor$.

VII. PERFORMANCE

A. Category of Gun Identification

In this section, we consider the neural network previously introduced to infer the *Category of gun*. We reconsider Table II and we divide the dataset into three classes, namely, *Pistols*, *Rifles*, and *Shotguns*, according to the gun models in the dataset. Figure 11 shows the confusion matrix computed as the average of 50 training and validation runs.

		Guns' Categories				
		Pistol	Rifle	Shotgun	Accuracy	Error
True Class	Pistol	412	13	6	95.6%	4.4%
	Rifle	4	169	21	87.1%	12.9%
	Shotgun	1	9	85	89.5%	10.5%
		98.8%	88.5%	75.9%		
		1.2%	11.5%	24.1%		
		Pistol	Rifle	Shotgun		
		Predicted Class				

Fig. 11: Confusion matrix associated with the classification of the Category of Guns.

The accuracy acc can be computed according to Eq. 9.

$$acc = \frac{1}{N} \sum_{i=1}^{N_C} x_{ii} \quad (9)$$

where $N = 722$ is the total number of samples, $N_C = 3$ is the number of classes, and x_{ii} is the i^{th} diagonal element of the confusion matrix, yielding to $acc \approx 0.92$. The confusion

matrix in Fig. 11 also reports summaries of columns and rows, predicted and true classes, respectively.

We observe that the classification error spans between 4.6% and 13.9% for Pistol and Rifle classes, respectively. The class Rifle (an actual gunshot from a rifle) is incorrectly classified as either Pistol (5 times) or Shotgun (22 times) in the 13.9% of the cases. The same type of analysis can be performed column-wise, where the prediction error spans between 1.4% and 25.2%. As an example, we observe that a prediction on class Shotgun is wrong in the 25.2% of the cases (7 times for Pistol and 22 times for Rifle).

Finally, we observe that while the Pistol class is likely to be correctly classified all the times, the vast majority of errors are happening between the Rifle and Shotgun classes.

B. Caliber Identification

In this section, we report the performance of our classification algorithm when considering 7 different calibers from Table II. We group the video chunks based on gun caliber, obtaining 7 different classes, namely, 12, 357M, 44M, 45acp, 556NATO, 762x39, and 9mm. Figure 12 shows the confusion matrix computed as the average of 50 training and validation runs. The overall accuracy computed according to Eq. 9 sums up to $acc \approx 0.9$. Best and worst performance are achieved by 9mm and 762x39, respectively. In particular, class 762x39 is wrongly predicted 8 times as class 556NATO. Classes 556NATO and 762x39 are intrinsically similar, since both are from class Rifle. Therefore, they are prone to be confused. Nevertheless, we observe that this phenomenon is very limited since we have 4 cases of 556NATO classified as 762x39, and 8 cases for the opposite configuration. We also observe that 556NATO and 762x39 classes experience a significant amount of misclassifications with classes 12 and 357M. Conversely, class 9mm is the most likely to be correctly classified.

		Guns' Calibers								
		12	357M	44M	45acp	556NATO	762x39	9mm	Accuracy	Error
True Class	12	84	2			7	4		86.6%	13.4%
	357M	2	77	3		4	1	1	87.5%	12.5%
	44M		1	92	3			1	94.8%	5.2%
	45acp			9	93			2	89.4%	10.6%
	556NATO	4	2			93	3		91.2%	8.8%
	762x39	5	3			8	75		82.4%	17.6%
	9mm		1	5	3			132	93.6%	6.4%
		88.4%	89.5%	84.4%	93.9%	83.0%	90.4%	97.1%		
		11.6%	10.5%	15.6%	6.1%	17.0%	9.6%	2.9%		
		12	357M	44M	45acp	556NATO	762x39	9mm		
		Predicted Class								

Fig. 12: Confusion matrix associated with the classification of the Calibers of Guns.

Highlights. By combining Fig. 11 and Fig. 12 we can draw some interesting remarks. Rifle class is misclassified for class

Shotgun 20 times (the opposite is happening 12 times) in Fig. 11, while 6+8=14 times (5+2=7 times) in Fig. 12. We think that the error is not due to a specific caliber, either the 556NATO or the 762x39, but to the feature similarities between the two classes; Shotgun and Rifle classes.

Pistol class is also misclassified as Rifle class 15 times. By looking into the details of Fig. 12, we observe that the major source of misclassifications is coming from the 357M class, classified 4 times as 556NATO and 1 time as 762x39. We observe that the 357M is the most powerful among the pistol calibers hence it is the closest to Rifle class in terms of bullet size, pressure, and barrel diameter.

Finally, we observe that our solution is particularly robust in detecting pistols. In particular, one of the most adopted worldwide caliber (9mm) is characterized by a very limited number of misclassifications (11 out of 167 total). The same considerations apply to classes 44M and 45acp.

C. Model Identification

In this section, we consider all of the gun models previously introduced in Table II with the aim of classifying each of them. The total number of classes sums up to 59, which is the number of gun models considered throughout this paper. We report the confusion matrix associated with the aforementioned classification in Appendix B. The accuracy sums up to $acc \approx 0.90$ and the maximum number of misclassifications (per model) never exceeds 2. We observe that class 38 (Ruger GP100 Match Champion) is never correctly classified. Finally, we highlight that the number of samples for the validation process is small (20% of each gun model in Table II). Nevertheless, the diagonal of the matrix in Appendix B collects the vast majority of the samples confirming the effectiveness of our model. We are confident that a larger data sample can increase the accuracy performance and effectiveness of gun model detection from gunshot sounds.

D. Testing

To validate our methodology, we tested the model against a new set of audio samples taken from videos different than the ones considered before with varying conditions, including the background noise and relative positions between the microphone and the shooter. We consider a total of 115 audio samples constituted by 13 Pistol (Beretta 92 FS), 59 Rifles (Ruger AR, Daniel Defense M4 A1 SOCOM, Maadi AK-47) and 44 Shotguns (Maverick 88, Winchester Model 300 Defender). We observe that Pistol and Rifle classification is characterized by high performance, where only 4 Rifles samples are misclassified for Pistol. As for the Shotgun class, we highlight that the two shotguns considered are not in the training set (Table II) because we did not find any valid samples from additional videos. Although the audio samples are coming from different shotgun models, our algorithm can still detect the caliber with high probability (only 8 audio samples are misclassified), which verifies the effectiveness and correctness of our algorithm. Finally, we observe that the

overall accuracy is consistent with the validation process and sums up to about 0.9.

True Class	Predicted Class			Accuracy	
	Pistol	Rifle	Shotgun		
Pistol	13			100.0%	
Rifle	4	54		93.1%	6.9%
Shotgun		8	36	81.8%	18.2%
	76.5%	87.1%	100.0%		
	23.5%	12.9%			

Fig. 13: Confusion matrix associated with guns' categories testing while adopting new audio samples.

VIII. CONCLUSION

Although scenarios requiring in-depth digital forensic of gunshots are countless, including military operations, mass-shooting, and possibly others, current solutions are far from reaching an adequate accuracy under real conditions.

In this paper, we have proposed an effective and efficient methodology to uniquely fingerprint gunshots enabling the identification of the category, caliber, and the model of the gun with an accuracy higher than 90% regardless of the capture conditions. Unlike existing solutions, our technique requires neither ad-hoc deployment of microphone networks, nor specific sample quality, and is agnostic to the microphone position with respect to the shooter. We have demonstrated that forensic analysis in the time-frequency domain of a single gunshot audio sample recorded by a commercial microphone (44100 samples per seconds) can be effectively used to infer the gun model (and other related characteristics). The proposed solution may lead to new insights and further developments in the area of weapon classification considering more samples, different noise levels, and a much larger weapon database.

ACKNOWLEDGMENT

This publication was partially supported by the following awards: NPRP12S-0125-190013 from QNRF-Qatar National Research Fund, a member of The Qatar Foundation, and the Innovation Research Project "Anti-Spoofers - GPS Spoofing Detection Exploiting Crowd-Sourced Information" awarded by the HBKU Innovation Center. The information and views set out in this publication are those of the authors and do not necessarily reflect the official opinion of the QNRF.

REFERENCES

- [1] J. Sallai, W. Hedgecock, P. Volgyesi, A. Nadas, G. Balogh, and A. Ledeczki, "Weapon classification and shooter localization using distributed multichannel acoustic sensors," *Journal of Systems Architecture*, vol. 57, no. 10, pp. 869–885, 2011.
- [2] U. Libal and K. Spyra, "Wavelet based shock wave and muzzle blast classification for different supersonic projectiles," *Expert Systems with Applications*, vol. 41, no. 11, pp. 5097–5104, 2014.
- [3] H. A. Sánchez-Hevia, D. Ayllón, R. Gil-Pita, M. Rosa-Zurera, H. A. Sanchez-Hevia, D. Ayllon, R. Gil-Pita, and M. Rosa-Zurera, "Maximum likelihood decision fusion for weapon classification in wireless acoustic sensor networks," *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, vol. 25, no. 6, pp. 1172–1182, 2017.
- [4] S. Khan, A. Divakaran, and H. S. Sawhney, "Weapon identification across varying acoustic conditions using an exemplar embedding approach," in *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense IX*, vol. 7666. International Society for Optics and Photonics, 2010, p. 76662B.
- [5] —, "Weapon identification using hierarchical classification of acoustic signatures," in *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense VIII*, vol. 7305. International Society for Optics and Photonics, 2009, p. 730510.
- [6] M. Djeddou and T. Touhami, "Classification and modeling of acoustic gunshot signatures," *Arabian journal for science and engineering*, vol. 38, no. 12, pp. 3399–3406, 2013.
- [7] K. D. Morton Jr, P. A. Torrione, and L. Collins, "Classification of acoustic gunshot signatures using a nonparametric bayesian signal model," in *Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense X*, vol. 8019. International Society for Optics and Photonics, 2011, p. 80190T.
- [8] E. Kiktova, M. Lojka, M. Pleva, J. Juhar, and A. Cizmar, "Gun type recognition from gunshot audio recordings," in *3rd International Workshop on Biometrics and Forensics (IWBF 2015)*. IEEE, 2015, pp. 1–6.
- [9] D. Ravi, C. Wong, B. Lo, and G. Yang, "Deep learning for human activity recognition: A resource efficient implementation on low-power devices," in *2016 IEEE 13th International Conference on Wearable and Implantable Body Sensor Networks (BSN)*, June 2016, pp. 71–76.
- [10] S. Hershey, S. Chaudhuri, D. P. W. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, M. Slaney, R. J. Weiss, and K. Wilson, "CNN architectures for large-scale audio classification," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, March 2017, pp. 131–135.
- [11] J. Dennis, H. D. Tran, and H. Li, "Spectrogram image feature for sound event classification in mismatched conditions," *IEEE Signal Processing Letters*, vol. 18, no. 2, pp. 130–133, Feb 2011.
- [12] A. Satt, S. Rozenberg, and R. Hoory, "Efficient emotion recognition from speech using deep learning on spectrograms," in *INTERSPEECH*, 2017.
- [13] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, "Multimodal deep learning," in *Proceedings of the 28th international conference on machine learning (ICML-11)*, 2011, pp. 689–696.
- [14] M. Navrátil, V. Křesálek, and P. Dostálek, "Neural network classification of gunshots using spectral characteristics," in *Proceedings of the 13th WSEAS International Conference on Automatic Control, Modelling and Simulation*, ser. ACMOS'11. Stevens Point, Wisconsin, USA: World Scientific and Engineering Academy and Society (WSEAS), 2011, pp. 262–267. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2001466.2001516>
- [15] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov 1998.
- [16] J. Schlter and S. Beck, "Improved musical onset detection with convolutional neural networks," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 6979–6983.
- [17] Y. Kim and T. Moon, "Human detection and activity classification based on micro-doppler signatures using deep convolutional neural networks," *IEEE Geoscience and Remote Sensing Letters*, vol. 13, no. 1, pp. 8–12, Jan 2016.
- [18] Y. M. Costa, L. S. Oliveira, and C. N. Silla, "An evaluation of convolutional neural networks for music classification using spectrograms," *Applied Soft Computing*, vol. 52, pp. 28 – 38, 2017. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1568494616306421>
- [19] L. L. Wyse, "Audio spectrogram representations for processing with convolutional neural networks," *ArXiv*, vol. abs/1706.09559, 2017.
- [20] "What are the most popular calibers in the us?" <http://knowledgeglue.com/what-are-the-most-popular-calibers-in-the-us>, accessed: 2019-04-06.
- [21] "Top 5 most popular calibers for handguns," <https://www.guns.com/news/2019/06/25/top-5-most-popular-calibers-for-handguns>, accessed: 2019-04-06.
- [22] J. Sallai, W. Hedgecock, P. Volgyesi, A. Nadas, G. Balogh, and A. Ledeczki, "Weapon classification and shooter localization using distributed multichannel acoustic sensors," *J. Syst. Archit.*, vol. 57, no. 10, pp. 869–885, Nov. 2011. [Online]. Available: <http://dx.doi.org/10.1016/j.sysarc.2011.04.003>
- [23] J. Sallai, A. Lédeczi, and P. Völgyesi, "Acoustic shooter localization with a minimal number of single-channel wireless sensor nodes," in *Proceedings of the 9th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '11. New York, NY, USA: ACM, 2011, pp. 96–107. [Online]. Available: <http://doi.acm.org/10.1145/2070942.2070953>
- [24] P. Volgyesi, G. Balogh, A. Nadas, C. B. Nash, and A. Ledeczki, "Shooter localization and weapon classification with soldier-wearable networked sensors," in *Proceedings of the 5th International Conference on Mobile Systems, Applications and Services*, ser. MobiSys '07. New York, NY, USA: ACM, 2007, pp. 113–126. [Online]. Available: <http://doi.acm.org/10.1145/1247660.1247676>
- [25] "youtube-dl," <https://github.com/yt-dl-org/youtube-dl>, accessed: 2019-04-06.
- [26] "Ffmpeg," <https://www.ffmpeg.org/>, accessed: 2019-04-06.
- [27] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *Proceedings of the 27th International Conference on Machine Learning*, ser. ICML'10. USA: Omnipress, 2010, pp. 807–814. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3104322.3104425>
- [28] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, Y. W. Teh and M. Titterton, Eds., vol. 9. Chia Laguna Resort, Sardinia, Italy: PMLR, 13–15 May 2010, pp. 249–256. [Online]. Available: <http://proceedings.mlr.press/v9/glorot10a.html>
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, ser. NIPS'12. USA: Curran Associates Inc., 2012, pp. 1097–1105. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2999134.2999257>
- [30] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, Jan. 2014. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2627435.2670313>
- [31] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 448–456. [Online]. Available: <http://proceedings.mlr.press/v37/ioffe15.html>
- [32] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014.

Appendix A: Making up layers of our architecture.

Name	Type	Learnables	Total Learnables
inputImage (66x100x1 with 'zerocenter' normalization)	Image input	-	0
conv1 (40 3x3x1 convs, stride: [1 1], padding: 'same')	Convolution	Weights: 3x3x1x40 Bias: 1x1x40	400
batchnorm1 (Batch Normalization, 40 channels)	Batch Normalization	Offset: 1x1x40 Scale: 1x1x40	80
relu1	ReLU	-	0
maxpool1 (3x3 Max Pooling, stride: [2 2], padding: 'same')	Max Pooling	-	0
conv2 (80 3x3x40 convs, stride: [1 1], padding: 'same')	Convolution	Weights: 3x3x40x80 Bias: 1x1x80	28880
batchnorm2 (Batch Normalization, 80 channels)	Batch Normalization	Offset: 1x1x80 Scale: 1x1x80	160
relu2	ReLU	-	0
maxpool2 (3x3 Max Pooling, stride: [2 2], padding: 'same')	Max Pooling	-	0
conv3 (160 3x3x80 convs, stride: [1 1], padding: 'same')	Convolution	Weights: 3x3x80x160 Bias: 1x1x160	115360
batchnorm3 (Batch Normalization: 160 channels)	Batch Normalization	Offset: 1x1x160 Scale: 1x1x160	320
relu3	ReLU	-	0
maxpool3 (3x3 Max Pooling, stride: [2 2], padding: 'same')	Max Pooling	-	0
conv4 (160 3x3x160 convs, stride [1 1], padding: 'same')	Convolution	Weights: 3x3x160x160 Bias: 1x1x160	230560
batchnorm4 (Batch Normalization, 160 channels)	Batch Normalization	Offset: 1x1x160 Scale: 1x1x160	320
relu4	ReLU	-	0
maxpool4 (1x13 Max Pooling, stride: [1 1], padding: [0 0 0 0])	Max Pooling	-	0
Dropout (20%)	Dropout	-	0
fc5 (7-neuron fully connected layer)	Fully Connected	Weights: 7x1440 Bias: 7x1	10087
Softmax	Softmax	-	0

