

LA-UR-21-32269

Accepted Manuscript

Noninvasive acoustic time-of-flight measurements in heated, hermetically-sealed high explosives using a convolutional neural network

Greenhall, John James

Zerkle, David Karl

Davis, Eric Sean

Broilo, Robert M.

Pantea, Cristian

Provided by the author(s) and the Los Alamos National Laboratory (2022-08-08).

To be published in: Machine Learning with Applications

DOI to publisher's version: 10.1016/j.mlwa.2022.100391

Permalink to record:

<http://permalink.lanl.gov/object/view?what=info:lanl-repo/lareport/LA-UR-21-32269>




Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Journal Pre-proof

Noninvasive acoustic time-of-flight measurements in heated, hermetically-sealed high explosives using a convolutional neural network

John Greenhall, David Zerkle, Eric Sean Davis, Robert Broilo, Cristian Pantea



PII: S2666-8270(22)00073-1
DOI: <https://doi.org/10.1016/j.mlwa.2022.100391>
Reference: MLWA 100391

To appear in: *Machine Learning with Applications*

Received date: 10 June 2022
Revised date: 12 July 2022
Accepted date: 25 July 2022

Please cite this article as: J. Greenhall, D. Zerkle, E.S. Davis et al., Noninvasive acoustic time-of-flight measurements in heated, hermetically-sealed high explosives using a convolutional neural network. *Machine Learning with Applications* (2022), doi: <https://doi.org/10.1016/j.mlwa.2022.100391>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2022 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Revised manuscript (Clean Version)

Noninvasive acoustic time-of-flight measurements in heated, hermetically-sealed high explosives using a convolutional neural network

John Greenhall (corresponding author): jgreenhall@lanl.gov, Los Alamos National Laboratory, P.O. 1663, Los Alamos, NM, 87545, USA

David Zerkle: dzerkle@lanl.gov, Los Alamos National Laboratory, P.O. 1663, Los Alamos, NM, 87545, USA

Eric Sean Davis: esdavis@lanl.gov, Los Alamos National Laboratory, P.O. 1663, Los Alamos, NM, 87545, USA

Robert Broilo: broilo@lanl.gov, Los Alamos National Laboratory, P.O. 1663, Los Alamos, NM, 87545, USA

Cristian Pantea: pantea@lanl.gov, Los Alamos National Laboratory, P.O. 1663, Los Alamos, NM, 87545, USA

Noninvasive acoustic time-of-flight measurements in heated, hermetically-sealed high explosives using a convolutional neural network

John Greenhall^{1,a}, David Zerkle^a, Eric Davis^a, Robert Broilo^a, and Cristian Pantea^a

^aLos Alamos National Laboratory, Los Alamos, NM, 87545, USA

¹Corresponding author: John Greenhall, jgreenhall@lanl.gov

Abstract

We present a data-driven technique for measuring the time-of-flight through material sealed within a container. Time-of-flight measurement provides a noninvasive means of quantifying the sound speed profile within a material by transmitting an acoustic burst and then measuring the time required for the burst to arrive at an opposing receiver. In a hermetically-sealed cylindrical container, a portion of the acoustic energy propagates through the material as a bulk wave, while the remainder of the acoustic energy propagates around the container walls as guided waves. As a result, interference from the guided waves obscures the bulk arrival, inhibiting measurement of the sound speed. The technique uses a Convolutional Neural Network (CNN) to identify critical features in the measured waveforms and identify bulk wave arrivals. We demonstrate this time-of-flight measurement technique on high explosive-filled containers as they are heated from room temperature to detonation. This is a particularly challenging application for acoustic time-of-flight measurements as the high explosives have significant sound speed gradients as they undergo heating, and they lead to significant attenuation of the bulk wave, as opposed to the guided waves, which do not suffer significant attenuation. We characterize the performance of the CNN as a function of the high explosive temperature and as a function of the CNN hyperparameters. We then provide physical insight into the error trends.

1. Introduction

The ability to detect changes in temperature and temperature gradients within high explosives (HE) inside a closed container is critical to assessing the level of danger presented by the device and how to safely handle the device. High explosives exposed to thermal insults due to accidents or other difficult to anticipate sources may undergo thermal decomposition and ignition. This insult process may result in reaction violence ranging from benign pressure bursting of the containment vessel, through extremely violent detonation. The safety of personnel attending to a high explosive charge insulted in such ways is of critical concern.

A principal difficulty in assessing when a heated high explosive charge will ignite is not knowing the spatially and temporally varying temperature field and other thermal damage states interior to the charge. Further complications are often present in high explosive formulation consisting of energetic material particles distributed with high mass loading into a continuous matrix of trinitrotoluene (TNT). Examples include Cyclotol (RDX/TNT, 75/25), Octol (HMX/TNT, 70/30), and Pentolite (PETN/TNT, 50/50). In these examples the energetic material particles are cyclotrimethylenetrinitramine (RDX; $(\text{CH}_2\text{N}_2\text{O}_2)_3$), cyclotetramethylenetetranitramine (HMX; $(\text{CH}_2\text{N}_2\text{O}_2)_4$), and pentaerythritol tetranitrate (PETN; $\text{C}_5\text{H}_8\text{N}_4\text{O}_{12}$) and are found in the 10-100 micron size range. These high explosive mixtures are convenient because they are *melt castable*, that is, the TNT component melts at a temperature near 80°C, which allows high mass loadings of other high melting point crystalline explosives to be stirred in, and the mixture re-solidified with minimal decomposition. Difficulties arise during prediction of thermal insult response because at temperatures above 80°C the suspended particles partially or fully dissolve into the molten TNT, settle under the influence of gravity in the relatively less dense molten TNT, and possibly undergo bulk

advection due to buoyancy forces. A means for measuring temperature internal to an explosive charge undergoing such thermal changes is crucial for predicting ignition response.

Current temperature measurement techniques typically require invasive temperature probes such as thermocouples, thermistors, or platinum resistance (Childs et al., 2000). However, the invasive temperature measurements can be dangerous, as penetrating the container and agitating the HE could lead to detonation. Furthermore, a physical probe cannot provide information about the spatial distribution of internal temperature and will never be able to report the maximum internal temperature, which might drive thermal runaway and accidental explosion. Alternatively, noninvasive techniques such as infrared thermography are limited to measuring surface temperatures, and, thus cannot provide information about the temperature distribution within the HE.

Other noninvasive inspection techniques like x-ray radiography suffer limited contrast between HE at different temperatures, require bulky equipment, and/or require long measurement times. In contrast, acoustic techniques based on time-of-flight are highly sensitive to changes in temperature, are portable, and enable rapid measurement times on the order of milliseconds. In addition to monitoring temperature, acoustic measurements enable noninvasive measurement for in situ process control (Chillara et al., 2017; Remenieras et al., 1994; Sinha, 1998), monitoring for biohazards (Sinha et al., 1997), and chemical reaction monitoring (Shen et al., 2017). In these situations, the sound speed of the material within the container is measured and correlated with a physical property of interest such as temperature, material phase, species concentration, etc. To measure the sound speed in the HE, opposing transducers are typically attached to the container walls. One transducer acts as a transmitter and excites an acoustic burst that propagates through the container walls and the HE, before being measured by one or more receivers. The sound speed is then determined from the container geometry and the time-of-flight required for the acoustic burst to propagate from the transmitter to the receiver(s) using ray tomography (Williamson, 1991). With the estimated sound speed profile within the container, the temperature profile can be measured via an empirical model of the temperature as a function of sound speed for a given material.

However, within a closed container, a portion of the acoustic signal propagates through the HE as a “bulk wave”, while the remainder propagates around the container walls as guided waves. In some situations, where the bulk and guided wave arrivals are separated in time, the bulk wave arrivals can be identified via time-gating, wherein only time windows containing the bulk wave arrivals are considered. However, in many applications, such as ToF measurements of HE in cylindrical containers, there is significant interference between the bulk and guided waves, which precludes time-gating to separate the arrivals. Currently, there are several techniques for detecting the bulk waves through cylindrical containers in the presence of guided waves in the container walls. Chillara et al. separated the bulk and guided waves in a fluid-filled container by first measuring the guided waves and then subtracting the guided waves from the combined bulk and guided wave signal (Chillara et al., 2017). To measure the guided wave, the authors are emptying the container, leaving only a thin fluid residue on the container walls to eliminate any bulk waves. Alternatively, Greenhall et al. implemented broadband linear chirp excitations to exploit the dispersive behavior of the guided waves (Greenhall et al., 2021). The frequency-dependent sound speed of the guided waves distort the guided waves, whereas the bulk waves remain undistorted. The undistorted bulk waves can then be identified using cross-correlation. However, many applications preclude collecting baseline measurements on the empty container, and exploiting the guided wave dispersion typically requires high frequencies (>1 MHz) that attenuate significantly in HE. Thus, the work presented in this manuscript is motivated by the need for a technique to detect bulk waves in the presence of interfering guided waves that works for low-frequencies (100s of kHz) without baseline measurements. Recently, there has been significant work using Convolutional Neural Networks (CNNs) to classify acoustic signals (Adavanne et

al., 2018; Barat et al., 2017; Bianco et al., 2019; Chen et al., 2019; Liu et al., 2021a, 2021b; McBrearty et al., 2019; Perol et al., 2018; Ross, Meier, & Hauksson, 2018; Ross, Meier, Hauksson, et al., 2018). CNNs are commonly applied to image and acoustic signal processing because they work well for signals with spatial/temporal relationships, e.g. an object is identified by observing the neighboring pixels or an acoustic signal can be identified by observing the neighboring times in the recording. By considering only a small region of the signal at a time, the CNN reduces the number of weights and biases that must be learned, which significantly improves the model generalizability, thereby increasing accuracy and reducing the number of training samples required, compared to other neural networks. CNNs work by learning descriptive features within a signal by passing the signal through several layers, wherein it is convolved with template signals, activated by a nonlinear function, and down-sampled before being passed to the next layer. The templates of the signals are learned by passing training data consisting of, e.g. waveforms with known Time-of-Flight (ToF) values to the CNN and then minimizing the error between the estimated ToF and the true ToF. CNNs have been applied successfully for identifying acoustic waves in seismology (Chen et al., 2019; McBrearty et al., 2019; Perol et al., 2018; Ross, Meier, & Hauksson, 2018; Ross, Meier, Hauksson, et al., 2018), oceanic transmission (Liu et al., 2021b), air transmission (Adavanne et al., 2018), and transmission in wave guides (Barat et al., 2017). In these applications, waves typically propagate either as purely bulk waves (P- and/or S-waves) or guided waves. Conversely, within a cylindrical container, acoustic waves can propagate as bulk waves within the container contents and as guided waves through the container walls. Additionally, it is not uncommon for the bulk and guided waves arrive at similar times, which causes interference between the waves. As a result, the bulk waves are occluded by the guided waves, which have similar spectral properties and amplitudes that are comparable to or even greater than that of the bulk waves.

In this work, we present a noninvasive technique for measuring bulk wave ToF through HE in a sealed container using pairs of opposing piezoelectric transducers. The objective of this manuscript is to demonstrate the use of a CNN to identify bulk wave arrivals through the HE in the presence of interfering guided waves propagating around the container walls. Additionally, we present a data preprocessing technique based on cross-correlation and time-windowing to reduce the network complexity and reduce network bias. We implement the technique on experimental data collected from sealed containers of pentolite as they are heated to the point of detonation. We discuss the algorithm error, compared to hand-picked “true” ToF measurements, and the effects of various CNN hyperparameters.

2. ToF estimation

Acoustic ToF measurements were conducted within HE-filled containers (Al-6061, inner diameter 140 mm, thickness 6 mm) using 16 piezoelectric disk transducers (PZT-J, diam. 7 mm, thickness 0.5 mm) evenly spaced around the container circumference. Over the course of several hours, the container was heated from below. As it was heated, we repeatedly transmitted acoustic bursts from one transducer, while measuring the acoustic response through the HE and container using the remaining 15 transducers. We would then shift the excitation to the next transmitter, which became the transmitter, while the remaining 15 transducers served as receivers. We measured acoustic bursts this way using each possible combination of 1-transmitter and 15-receivers, resulting in 240 measured waveforms. We then repeated this process, measuring 240 waveforms every 100 s until the heating led to detonation of the HE. Figure 1(a) shows the acoustic measurements within the HE-filled container transmitted from a single transmitter/receiver (Tx/Rx) pair, separated by angle $\theta = 180^\circ$, over the duration of the heating experiment. Figure 1(b) shows the corresponding temperatures measured at the outer wall (blue) and in the HE at the center of the container (orange) as the container is heated from below. We observe a significant increase in the ToF of the bulk wave arrivals propagating through the HE due to the decrease in the HE sound speed at increased

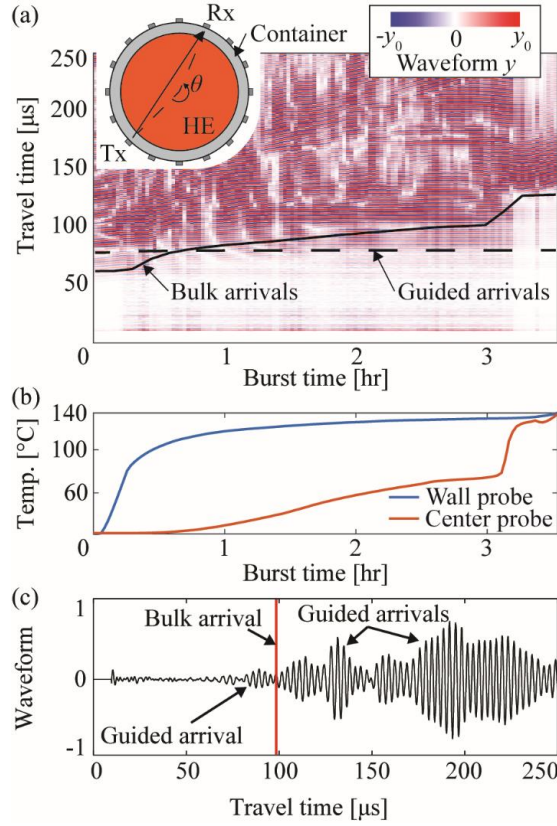


Figure 1: (a) Example acoustic measurements in an HE-filled container between a single transmitter/receiver pair ($\theta = 180^\circ$) with (b) the wall and center temperatures as the container is heated. (c) A single waveform (i.e. a single vertical line from (a)) illustrates the interference between the bulk and guided waves.

temperatures. Alternatively, we observe that the ToF of the guided waves propagating through the container walls do not change significantly due to the small (approximately 5 %) change in sound speed of the Al container over the measured temperature range. As a result, identifying the bulk wave can be relatively simple when presented with all acoustic measurements over the duration of the heating experiment. However, distinguishing the bulk wave from the guided waves from a single acoustic burst is considerably more challenging. This is apparent in Figure 1(c), where dispersion causes the guided waves to elongate so that they arrive before, during, and after the bulk wave arrival; and the guided waves may even have larger amplitudes and obscure the bulk waves. Thus, we implement a data-driven technique based on a CNN to identify the bulk ToF from a single waveform.

Figure 2 shows an example of the excitation signal (green) and the resulting measured waveform (blue) with the ToF probability (red) of a bulk wave. We preprocess the waveform by computing the cross-correlation envelope (CCE) (black in Figure 2) between the measured signal $y(t)$ and the excitation signal $y_{ex}(t)$ using

$$CCE(t) = |H(y(t) \star y_{ex}(t))|, \quad (1)$$




Figure 2: Example waveform and corresponding preprocessed signal (CCE) and Time-of-Flight (ToF) probability for CNN training.

where $H(\cdot)$ is the Hilbert transform and \star indicates cross-correlation. In contrast with the measured waveform, the CCE excludes the high-frequency content, while retaining the amplitude information. As a result, we can down-sample the CCE signal so that a longer time-span can be represented by a smaller number of data points. We note that the CCE signal may result in lower resolution when estimating the ToF.

To reduce the CNN complexity, we aim to arrange a network that takes in a CCE vector and outputs a scalar label. To achieve this, we break our CCE time-series into multiple windows, centered at t_q with span Δt and outputs a scalar label indicating whether the time window is centered on a bulk wave ToF, which we define as $P(t_q) > .5$, where $P(t_q)$ is the ToF probability at time t_q . Here, the window span Δt is treated as a hyperparameter, which is selected empirically.

Figure 3 shows a schematic of a classification CNN, which utilizes a series of layers that transform and reshape the input signal x to extract useful features (Goodfellow et al., 2016). Typically, CNNs consist of multiple sequences of convolution “Conv” (blue), Activation (orange), and Pooling layers (green) that are grouped together, followed by a single Fully-connected layer and Softmax layer. Additionally, we indicate the input signals to each layer $x^{(l)}$ and the output signal y in gray. The Conv convolves an input signal with multiple convolution filters to produce filtered signals with multiple channels (one per convolution filter) that are then passed into the next layer. The idea behind the convolutional filters is to extract meaningful

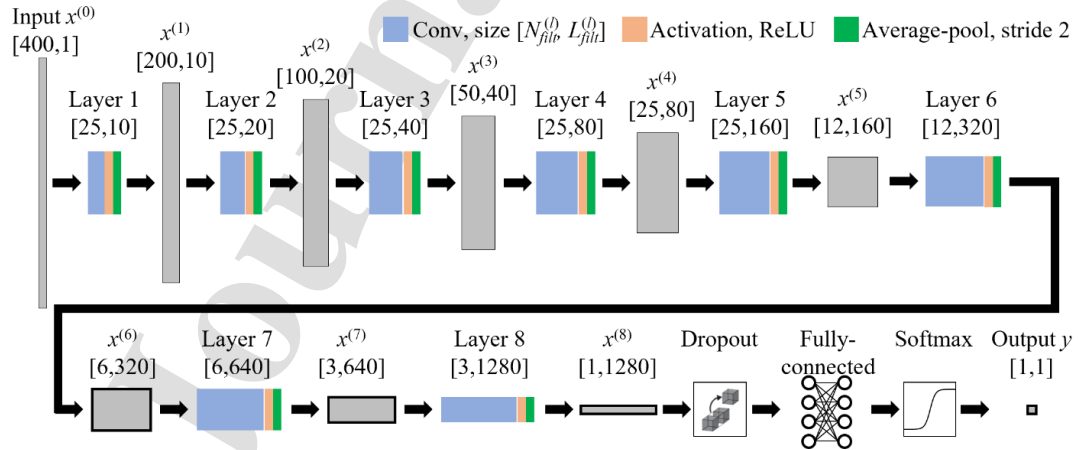


Figure 3: Schematic of the Convolutional Neural Network workflow, indicating the change in dimension of the signal $x^{(l)}$ between each layer.

patterns from the signal that are indicative of the signal label, e.g. whether or not there is a bulk wave arrival at the signal center. For the l^{th} convolution layer, the filtered signal is given

$$x_i^{(l)}(t) = \sum_j \left(x_j^{(l-1)}(t) * w_{ij}^{(l)}(t) \right) + b_i^l, \quad (2)$$

where $x_i^{(l)}(t)$ is the i^{th} channel of the filtered signal, $x_j^{(l-1)}(t)$ is the j^{th} channel of the signal from the previous layer, $w_{ij}^{(l)}(t)$ and b_i^l are the learned weights and biases of the l^{th} convolutional layer, and $*$ is the convolution operator. Our CNN utilizes 8 Conv layers, where the l^{th} Conv layer consists of $N_{filt}^{(l)}$ convolution filters, each of length $L_{filt}^{(l)}$. Here, the number of filters increases according to $N_{filt}^{(l)} = 2N_{filt}^{(l-1)}$, where the number of filters in the first layer is $N_{filt}^{(1)}=10$. We select a fixed filter length $L_{filt}^{(l)}=25$ if the input signal length $\text{len}(x^{(l)}) \geq 25$, otherwise we set $L_{filt}^{(l)} = \text{len}(x^{(l)})$. Here, the hyperparameters including the number of layers, the number of filters in the first layer $N_{filt}^{(1)}$, and the upper bound on $L_{filt}^{(l)}$ by sweeping through different permutations of hyperparameters and then selecting the combination with the lowest expected error (See Sec. 4 for details).

After each convolution layer, activation layers introduce nonlinearities that saturate or eliminate features that do not improve the CNN accuracy. As a result, the nonlinear activation functions enable the CNN to learn more complex models, as compared to purely linear networks (Hao et al., 2020). To insert nonlinearities into our CNN, we utilize the Relative Linear Unit (ReLU) activation, defined as

$$\text{ReLU}(x) = \begin{cases} x, & x > 0 \\ 0, & x \leq 0 \end{cases} \quad (3)$$

Additionally, between convolution layers, we use mean-pooling to reduce the length of the filtered signals. After passing a measured signal through multiple sequences of convolution, activation, and pooling layers, we obtain an output consisting of multiple channels, each containing a shortened feature vector (or scalar). To achieve a scalar label, we flatten the output into a single vector, and then pass it to a fully-connected layer that performs multiplication and summation, resulting in a scalar output feature. Finally, we pass the output feature through a softmax layer, which serves as an activation layer to normalize the output into probability distributions, i.e. it represents the probability of a bulk arrival occurring at the center of the time-windowed signal. Finally, dropout layers have been demonstrated to reduce overfitting in neural networks by neglecting a fraction of features that are randomly selected each training iteration (Srivastava et al., 2014). We test the effect of the dropout layer by adjusting the fraction of dropout samples and measuring the effect on the test error (see Sec. 4) and identify that a dropout rate of 25% achieves the lowest error. To learn the CNN weights and biases, we minimize the cross-entropy loss

$$L = -\frac{1}{N} \sum_n [y_n \log \hat{y}_n + (1 - y_n) \log(1 - \hat{y}_n)], \quad (4)$$

where y_n and \hat{y}_n are the true and estimated outputs of the CNN (the probability of a bulk arrival occurring at the center of the time-windowed input signal) for training sample $n = 1, 2, 3, \dots, N$. We perform the optimization using the Adaptive Moment Estimation (ADAM) solver, which typically demonstrates improved accuracy and learning time, compared to the more commonly used stochastic gradient descent (Kingma & Ba, 2017).

We train the CNN in epochs, where one epoch comprises training over all training samples once, and repeat epochs until the validation accuracy fails to increase by 1%. We observe that, for most hyperparameters,

the mean training error reaches approximately 90% after several hundred training samples and oscillates around that value for the duration of training. We also observe that the validation error reaches approximately 90% after the first epoch and only increases slightly (<1%) after the second epoch. As a result, the training sessions terminate after the second epoch.

Once we have trained the CNN, we then estimate the bulk ToF for a given waveform as follows. We first compute the CCE and break it into overlapping time-windowed samples, centered at times $0 \leq t_q \leq t_f$. We pass the samples through the trained CNN to estimate the bulk wave ToF probability $P(t_q)$ at each time t_q . Finally, we estimate the ToF as the time t_q with highest ToF probability

$$t_{arr} = \underset{t_q}{\operatorname{argmax}} P(t_q). \quad (5)$$

3. Experimental setup and data collection

Figure 4 shows an aluminum (Al) cylindrical container filled with HE, lined with 16 piezoelectric transducers around the circumference, and thermocouple probes measuring the temperatures at the container wall and at the center of the HE. We cast pentolite 50/50 within a sealed Al cylinder that is equipped with electric heaters, 16 acoustic transducers attached via high-temperature epoxy around the container circumference, and two thermocouples at approximately the same height as the acoustic transducers (approximately 50 mm from the container's interior upper/lower surface), one located at the outer surface and one located in the HE at the center of the container. We control the current to the heaters to heat the container to a soak temperature (138°C) from below, which rapidly heats the container walls before the heat propagates into the HE. During the heating process, we perform thermocouple measurements and acoustic measurements every minute. To measure the acoustic ToF, we use a single transducer as a transmitter to generate an acoustic Gaussian burst with center frequency 350 kHz, and standard deviation 150 kHz. We then use a Data Acquisition (NI PXIe-8880 embedded controller with 3x NI PXIe-5105 digitizers) to simultaneously measure the acoustic burst from the remaining transducers at a sampling rate of 2.0 MHz. We repeat this transmission/measurement process, where in each transmission we use a different transducer as the transmitter and measure the burst with the remaining transducers. Here, we use a multiplexer (Keithley 7001 Switch System with a 7011-S multiplexer card) to shift the excitation to each transducer, one at a time. We observe that for transmitter/receiver pairs with angular separation $\theta < 135^\circ$ or $\theta > 225^\circ$, the guided waves dominate, and the bulk waves cannot be identified reliably for the training samples that are labeled manually. As a result, we only consider measurements for transmitter/receiver separations of $135^\circ < \theta < 225^\circ$, i.e. we record measurements from five receivers per




Figure 4: Experimental setup used to measure acoustic arrivals and temperatures of HE within a sealed container.

transmitter. Additionally, we observe that the position and epoxy layer used to attach the hand-placed transducers affects the resulting measurements, and that heating the cylinder further changes the epoxy layer resulting in a reduction in the bulk wave amplitude and additional ringing of the transmitted burst. To ensure robustness of the CNN to these changes, we repeat the experiment on three different HE-filled containers that were heated to soak temperatures of either 132°C or 138°C.

At each burst time in the experiment, transmit bursts from 16 transducers, each measured by five receivers, resulting in 80 measurements, where 40 of those are unique (the burst from transducer n to transducer m is approximately equivalent to that from transducer m to transducer n). We collect measurements from 110 unique transmitter/receiver (Tx/Rx) pairs. For each Tx/Rx pair, we recorded between 125 and 140 acoustic bursts, depending on the time required to heat the HE to the point of detonation. In total, 14,795 acoustic bursts are collected. To train, validate, and test the CNN, we group acoustic measurements based on the Tx/Rx pair. These results in validation and testing samples from Tx/Rx combinations that the CNN has not experienced during training, which more closely represents real-world implementation case, where the CNN likely has not been trained on the Tx/Rx pair in the configuration being used to measure the acoustic ToF in HE.

4. Learning results and discussion

To train and test the CNN, we randomly select measurements for 22 Tx/Rx pairs (approximately 20%) to exclude from the training data. 11 Tx/Rx measurement sets are set aside for validation and the remaining 11 Tx/Rx measurement sets are set aside for testing. Figure 5 shows the results of the CNN ToF estimation

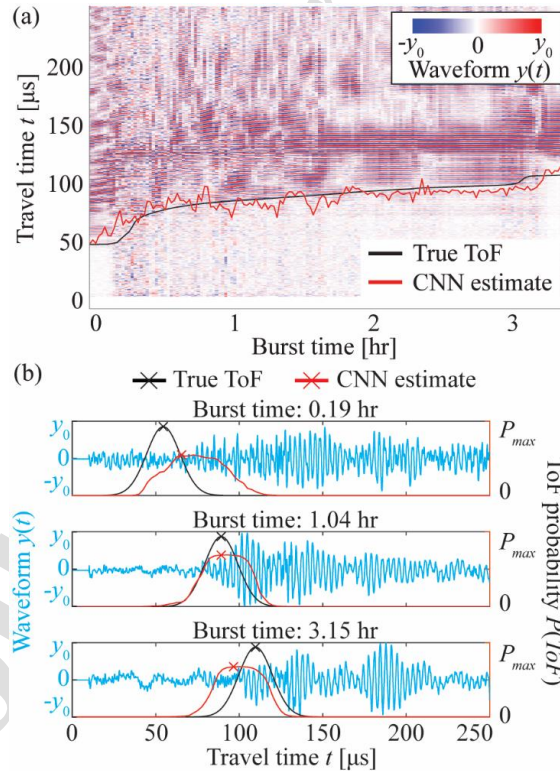


Figure 5: Comparison of true and CNN-estimate ToF. (a) All waveforms for one transducer/receiver pairs in the test set, superimposed by the true and estimated ToF. (b) Select example waveforms superimposed by the true and estimated ToF probabilities $P(ToF)$.

algorithm for two Tx/Rx pairs from the test sets. Figure 5(a) shows the measured waveforms (colormap) for all burst times T_b , which we normalize to be a fraction of the total experiment time. Here we normalize the amplitude of each waveform by the maximum waveform magnitude for that burst time. On top of the waveforms, we superimpose the “true” ToF (black) and the estimated ToF from the CNN (red). Figure 5(b) shows selected waveforms (blue) pulled from Figure 5(a) and (b), respectively, superimposed by the true (black) and estimated ToF probability lines (red). From Figure 5(a) and (b), we observe good agreement between the true and estimated ToF for the majority burst times. In the regions with larger error, we observe two a reduction in the waveform amplitude near the true ToF, compared to the noise level. From Figure 5(b), we observe that the estimated ToF probabilities from the CNN deviate from the Gaussian peaks of the “true” ToF probability, resulting in a shortened, broadened arrival curve. Here, the height and width of the estimated ToF probabilities is due to the width of the prescribed “true” ToF probabilities, because all training windows with ToF probabilities $>0.5 \times P_{max}$ are assigned training labels “1”. As a result, there is a range of ToF that are considered equally correct, as indicated by the broadened peaks of the estimated ToF probability curves.

To better understand the conditions resulting in large ToF error, we measure the error in the ToF as a function of burst time, normalized by the time required for detonation. Figure 6(a) shows the mean and standard deviation of the ToF errors at each burst time for the test sets. Figure 6(b) shows the thermocouple temperatures measured at the container wall (solid line) and at the center of the HE (dashed line) as a function of burst time. We observe that the largest errors occur early in the heating process (approximately 5% into the experiment) and right before detonation (approximately 90% through the experiment), as indicated by red vertical lines. From Figure 6(b), we observe that the periods of high ToF error correspond to temperatures of approximately 80°C at the wall (left line) and at the HE center (right line), respectively. The temperature of 80°C is critical because it is the melt temperature of the TNT binder. Initially, when the wall first reaches 80°C, it is likely that the expansion of the liquid TNT pushes the solid HE away from the wall, resulting in poor acoustic coupling, and, thus reducing the Signal-to-Noise Ratio (SNR). At approximately 90% of the total experiment time, when the center temperature reaches 80°C, the solid HE

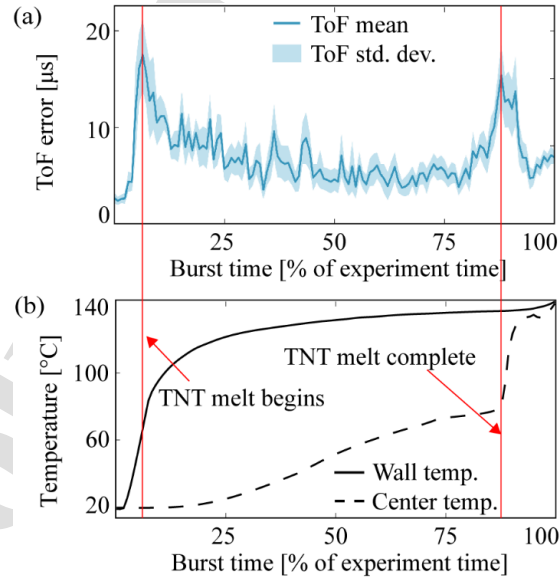


Figure 6: (a) Mean and standard deviation in test set ToF error at each burst time of the experiment. (b) Probe temperatures measured at the container wall and the HE center throughout the experiment, with red lines indicating the initiation and completion of the TNT melt.

becomes small enough that a portion of the acoustic burst propagates around the solid HE and a portion propagates through the solid region, both arriving at roughly the same times. As a result, there is destructive interference between the waves propagating through and around the solid region, which inhibits identifying either bulk wave ToF.

Away from the initial and final transition regions, we see good agreement between the theoretical ToF and the ToF predicted by the CNN. This results in an overall mean ToF error of 4.7 μs , which is approximately 5% of the theoretical ToF, and would result in approximately 5% error in the measured sound speed in the HE based on error propagation, neglecting errors in the path length. To understand how the algorithm is affected by the various hyperparameters in the time-windowing preprocessing and the CNN, we perform a parameter sweep, over the length of the time-window, the number of CNN layers (one CNN layer consists of 1 \times Conv, 1 \times ReLU, and 1 \times Mean-Pooling Layer), the number of filters per layer, and the length of filters in each layer. To simplify the hyperparameter characterization, we define the number of Conv filters for the first Conv layer, and then increase the number of filters by a factor of two at each subsequent layer. Additionally, we create Conv filters to have the same length in each layer. For each combination of hyperparameter values, we use K-fold cross validation, where we randomly separate the data into five groups, such that all measurements from a given Tx/Rx pair are placed in the same group. We then iterate through the five groups of samples, selecting one group as the validation set and combining the remaining four groups to form the training set. For each permutation of training/validation sets, we train the algorithm, and calculate the ToF error for the validation set. To quantify the effect of each hyperparameter combination, we calculate the average validation error across all five combinations of training/validation sets. After the CNN has been tested with all combinations of hyperparameters, we evaluate the effect of each hyperparameter by computing the correlation coefficient between the hyperparameter value and the resulting mean ToF error across each combination of training/validation sets. Table 1 shows the ranges of hyperparameters used, and the correlation coefficient between the hyperparameter value and the mean test set ToF error. Here, negative correlation coefficients indicate that increasing hyperparameter values result in decreasing ToF error. We observe the largest negative correlation attributed to the time-window length, likely because the features contain more context, e.g. whether a peak is the first peak or it follow a series of large peaks. The second most impactful hyperparameter was the number of CNN layers. In contrast, the number of Conv filters, the Conv filter lengths, and the Dropout fraction had insignificant effects on the ToF error.

Table 1: List of hyperparameter values and the correlation between hyperparameter value and the mean ToF error for the test sets.

Hyperparameter	Min value	Max value	Correlation coefficient
Time-window span, Δt [μs]	100	250	-0.89
Number of CNN layers	2	8	-0.44
Number of Conv filters	5	25	-0.07
Conv filter lengths	5	50	-0.06
Dropout fraction	0.0	0.2	-0.05

In addition to the CNN structure described in Sec. 2, there are alternative CNN architectures such as autoencoders, where an “encoder” uses a series of convolutional layers to reduce the time-dimension and increase the channel-dimension (same formulation as Sec. 2), and then a “decoder” uses a series of “transposed convolution” layers to grow the signal in the time-dimension and reduce the channel-dimension (Qian et al., 2018). However, including a deconvolution layer increases the number of weights and biases

that must be learned, which requires larger training data sets to prevent overfitting. Additionally, “dilated convolution” layers can be used to obtain multi-scale context without a reduction in resolution (Yu & Koltun, 2016). This has proven useful in semantic segmentation, but it is not typically intended for scalar classification problems like the one in this work.

5. Conclusion

We present a data-driven technique for identifying acoustic bulk wave ToF within closed-containers, where guided waves propagating through the container walls frequently interfere and obscure the bulk waves. We demonstrate this technique on cylindrical containers filled with a high explosive (HE) as it is heated from a solid to a liquid and eventually detonates. In contrast with other techniques for detecting bulk waves in closed-containers with interfering guided waves, our technique does not require baseline measurements in the empty container and they work for low acoustic frequencies that are necessary to overcome acoustic attenuation in HE. We implement a Convolutional Neural Network (CNN) to distinguish the bulk waves from the interfering guided waves, in contrast with other CNN-based acoustic ToF techniques that have been demonstrated on applications with either purely bulk waves or guided waves. We train, validate, and test our algorithm on experimental acoustic measurements of an Al cylinder filled with Pentolite that is gradually heated until detonation. We observe good agreement between the true ToF and those estimated by the CNN, with average errors on the order of 5% for the test measurements. Thus, this technique enables measuring acoustic time-of-flight within closed containers for a wide range of applications such as industrial process control, in-situ process monitoring, characterizing phase-change energy storage devices, and monitoring threat levels from HE devices, where current noninvasive and invasive techniques are not feasible.

Acknowledgements

This material was prepared by Los Alamos National Laboratory, which is operated by Triad National Security, LLC for the United States Department of Energy under contract DE-AC52-06NA25396.

References

- Adavanne, S., Politis, A., & Virtanen, T. (2018). Direction of Arrival Estimation for Multiple Sound Sources Using Convolutional Recurrent Neural Network. *2018 26th European Signal Processing Conference (EUSIPCO)*, 1462–1466. <https://doi.org/10.23919/EUSIPCO.2018.8553182>
- Barat, V., Kostenko, P., Bardakov, V., & Terentyev, D. (2017). *Acoustic Signals Recognition by Convolutional Neural Network* (No. 12). *12*(12), 9.
- Bianco, M. J., Gerstoft, P., Traer, J., Ozanich, E., Roch, M. A., Gannot, S., & Deledalle, C.-A. (2019). Machine learning in acoustics: Theory and applications. *The Journal of the Acoustical Society of America*, *146*(5), 3590–3628. <https://doi.org/10.1121/1.5133944>

- Chen, Y., Zhang, G., Bai, M., Zu, S., Guan, Z., & Zhang, M. (2019). Automatic Waveform Classification and Arrival Picking Based on Convolutional Neural Network. *Earth and Space Science*, 6(7), 1244–1261. <https://doi.org/10.1029/2018EA000466>
- Childs, P. R. N., Greenwood, J. R., & Long, C. A. (2000). Review of temperature measurement. *Review of Scientific Instruments*, 71(8), 2959–2978. <https://doi.org/10.1063/1.1305516>
- Chillara, V. K., Sturtevant, B. T., Pantea, C., & Sinha, D. N. (2017). Ultrasonic sensing for noninvasive characterization of oil-water-gas flow in a pipe. *AIP Conference Proceedings*, 1806, 090014.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
<https://www.deeplearningbook.org/contents/convnets.html>
- Greenhall, J., Hakoda, C., Davis, E. S., Chillara, V. K., & Pantea, C. (2021). Noninvasive Acoustic Measurements in Cylindrical Shell Containers. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 68(6), 2251–2258. <https://doi.org/10.1109/TUFFC.2021.3054716>
- Hao, W., Yizhou, W., Yaqin, L., & Zhili, S. (2020). The Role of Activation Function in CNN. *2020 2nd International Conference on Information Technology and Computer Application (ITCA)*, 429–432. <https://doi.org/10.1109/ITCA52113.2020.00096>
- Kingma, D. P., & Ba, J. (2017). Adam: A Method for Stochastic Optimization. *ArXiv:1412.6980 [Cs]*.
<http://arxiv.org/abs/1412.6980>
- Liu, Y., Chen, H., & Wang, B. (2021a). DOA estimation based on CNN for underwater acoustic array. *Applied Acoustics*, 172, 107594. <https://doi.org/10.1016/j.apacoust.2020.107594>
- Liu, Y., Chen, H., & Wang, B. (2021b). DOA estimation based on CNN for underwater acoustic array. *Applied Acoustics*, 172, 107594. <https://doi.org/10.1016/j.apacoust.2020.107594>
- McBrearty, I. W., Delorey, A. A., & Johnson, P. A. (2019). Pairwise Association of Seismic Arrivals with Convolutional Neural Networks. *Seismological Research Letters*, 90(2A), 503–509.
<https://doi.org/10.1785/0220180326>
- Perol, T., Gharbi, M., & Denolle, M. (2018). Convolutional neural network for earthquake detection and location. *Science Advances*, 4(2), e1700578. <https://doi.org/10.1126/sciadv.1700578>

- Qian, F., Yin, M., Liu, X.-Y., Wang, Y.-J., Lu, C., & Hu, G.-M. (2018). Unsupervised seismic facies analysis via deep convolutional autoencoders. *GEOPHYSICS*, 83(3), A39–A43.
<https://doi.org/10.1190/geo2017-0524.1>
- Remenieras, Cervenka, & Alais. (1994). Non intrusive measurements of the acoustic pressure and velocity fluctuations of fluids flowing in pipes. *Proceedings of IEEE Ultrasonics Symposium ULTSYM-94*, 1323–1326 vol.3. <https://doi.org/10.1109/ULTSYM.1994.401836>
- Ross, Z. E., Meier, M.-A., & Hauksson, E. (2018). P Wave Arrival Picking and First-Motion Polarity Determination With Deep Learning. *Journal of Geophysical Research: Solid Earth*, 123(6), 5120–5129. <https://doi.org/10.1029/2017JB015251>
- Ross, Z. E., Meier, M.-A., Hauksson, E., & Heaton, T. H. (2018). Generalized Seismic Phase Detection with Deep Learning. *Bulletin of the Seismological Society of America*, 108(5A), 2894–2901.
<https://doi.org/10.1785/0120180080>
- Shen, C.-C., Weng, M.-Y., Sheu, J.-K., Yao, Y.-T., & Sun, C.-K. (2017). In Situ Monitoring of Chemical Reactions at a Solid-Water Interface by Femtosecond Acoustics. *The Journal of Physical Chemistry Letters*, 8(21), 5430–5437. <https://doi.org/10.1021/acs.jpcclett.7b02384>
- Sinha, D. N. (1998). *Noninvasive identification of fluids by swept-frequency acoustic interferometry* (United States Patent No. US5767407A).
- Sinha, D. N., Springer, K., Han, W., Lizon, D., & Kogan, S. (1997). *Applications of swept-frequency acoustic interferometer for nonintrusive detection and identification of chemical warfare compounds* (LA-UR-97-3113; CONF-970962-). Los Alamos National Lab., NM (United States).
<https://doi.org/10.2172/555542>
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15, 1929–1958.
- Williamson, P. R. (1991). A guide to the limits of resolution imposed by scattering in ray tomography. *GEOPHYSICS*, 56(2), 202–207. <https://doi.org/10.1190/1.1443032>

Yu, F., & Koltun, V. (2016). *Multi-Scale Context Aggregation by Dilated Convolutions*

(arXiv:1511.07122). arXiv. <http://arxiv.org/abs/1511.07122>

Journal Pre-proof

- Acoustic Time-of-Flight critical to noninvasive inspection of sealed high explosive containers
- Interference between waves in high explosive and in container walls inhibits measurements
- Convolutional Neural Networks are demonstrated for distinguishing between interfering waves

Journal Pre-proof

John Greenhall: Methodology, Software, Validation, Data curation, Writing – Original draft preparation. **Dave Zerkle:** Supervision, Writing – Reviewing and Editing. **Eric Davis:** Software, Experimental data collection, Writing – Reviewing and Editing. **Robert Broilo:** Experimental data collection, Writing – Reviewing and Editing. **Cristian Pantea:** Conception, Supervision, Experimental data collection, Writing – Reviewing and Editing.

Journal Pre-proof

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Journal Pre