

ACOUSTIC EXPLOSION AND ROCKET SIGNATURES FROM
SURFACE AND AIRBORNE SMARTPHONES

A DISSERTATION SUBMITTED TO THE GRADUATE DIVISION OF THE UNIVERSITY
OF HAWAII AT MĀNOA IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR
THE DEGREE OF

DOCTOR OF PHILOSOPHY

IN

EARTH AND PLANETARY SCIENCE

MAY 2025

By

Sarah K. Popenhagen

Dissertation Committee:

Milton A. Garcés, Chairperson
Garrett Apuzen-Ito
Helen Janiszewski
Sloan Coats
Brian Powell

Keywords: rocket acoustics, infrasound, smartphone, monitoring, data

Acknowledgements

I would like to thank my advisor, Milton Garcés, for the opportunity to work with him and his research group and for sharing his knowledge, expertise, and advice with me over the last five years. I am also grateful to my fellow ISLA students and researchers, especially Kei Takazawa, Anthony Christe, Shirin Wyckoff, Meritxell Collet, Brian Williams, and Tyler Yoshiyama, for their support and friendship. I would like to thank my committee members, Garrett Apuzen-Ito, Helen Janiszewski, Sloan Coats, and Brian Powell, for their insight and encouragement throughout the long process of writing and defending this dissertation.

I am grateful to have had the support of the Consortium for Monitoring, Technology, and Verification, as well as that of the Consortium for Enabling Technologies and Innovation throughout my time as a graduate student. Through these consortia, I was able to connect and collaborate with numerous scientists at several national laboratories, all of whom were wonderful to work with and provided invaluable insight and advice, both scientific and professional. In particular, I would like to thank Danny Bowman, Cleat Zeiler, Luis Ocampo Giraldo, Edna Cárdenas, J’Tia Hart, and Megan McGarry for their mentorship and support.

I would also like to thank the staff members in both the Department of Earth Science and the Hawai‘i Institute of Geophysics and Planetology, including Lily Shao, Anela Nishimoto, Jason McCartney, Ann Yamauchi, James Cromwell, and Patrick Bullard. I am also eternally grateful to everyone who sacrificed their time, energy, and at times sanity working towards a brighter future through Huliāmahi, the JEDI committee, ALU, and numerous other groups. I’d especially like to thank those who made Huliāmahi what it is today, including Diamond, Chiara, Kelly, Annie, Kei, and Lucas, as well as Jordan, Sue, Mahina, Will, and everyone else involved in Huliāmahi this past year for their support, camaraderie, and refusal to succumb to the apathy.

I would like to express my heartfelt gratitude to everyone who supported me on the long road to the submission of this dissertation, including teachers and professors, friends back home and friends from undergrad, numerous extended family members, and of course all of the friends I've made here at UH. I'd especially like to thank Kei and Mayu for being my home away from home here in Hawai'i, and for becoming part of my family; my grandparents for their enduring love and support, which has been felt tangibly even from four thousand miles away; and my siblings for being my first and best friends, and for always being there both when I needed to laugh and when I wanted to cry. Finally, I would like to express my deepest appreciation to my parents for their endless support and encouragement, for the sacrifices they have made and continue to make for me and my siblings, and for raising me to value and strive for curiosity, integrity, rationality, and above all, compassion.

Abstract

Rocket launches produce complex acoustic signatures with large amounts of energy in low (<300 Hz) and infrasonic (<20 Hz) frequency ranges. These acoustic signatures carry information about their source as they propagate through the atmosphere, which, if correctly detected and classified, can aid in monitoring efforts. Many infrasonic rocket signatures have been successfully collected after propagating great distances (>1000 km), but most of these signatures are highly attenuated, yielding broad insight but little detail into the nature of the launch that produced them. In addition, very few of these data are available to the public, limiting the rate of progress in the field of rocket acoustics.

The overarching goal of the work covered in this dissertation is to develop a fast and reliable method of detecting acoustic rocket launch signatures. Results from each of the three content chapters contribute towards this goal. In Chapter 2, a surface chemical explosion signal is observed from an ascending balloon in the middle stratosphere. From this case of a stationary surface source and an ascending airborne receiver, we can gain insight into the inverse case of an ascending airborne source (a rocket) and a stationary surface receiver by invoking reciprocity. In addition, a comparison of low-cost, low-maintenance sensors (smartphones) with traditional infrasound sensors is made. In Chapter 3, a dataset of acoustic rocket signatures collected on smartphones at estimated propagation distances of 10-70 km is presented and released to the public, accompanied by preliminary analysis of the chronology and time-frequency characteristics observed in the dataset. In Chapter 4, this dataset is used in concert with two other open-access datasets to train and test machine learning models for near-real-time detection of acoustic rocket launch signatures on mobile platforms. The best performing model showed promising results, with an overall accuracy of 97% and a false positive rate of <1%. Measures to improve the model's suitability for persistent monitoring are discussed,

implemented, and evaluated, resulting in an estimated effective true positive rate of 99% and false positive rate of 0.07%.

Table of Contents

Acknowledgements	ii
Abstract	iv
List of Abbreviations	xi
Chapter 1: Introduction	1
Chapter 2: Acoustic Waves From a Distant Explosion Recorded on a Continuously Ascending Balloon in the Middle Stratosphere	3
2.1 Introduction	4
2.2 Data and Methods	6
2.3 Results	12
2.3.1 Acoustic Data	12
2.3.2 Acceleration Data.....	16
2.4 Discussion and Conclusions	16
2.5 Acknowledgements	17
2.6 Data Availability Statement	18
2.7 Supporting Information	19
Chapter 3: Acoustic Rocket Signatures Collected by Smartphones	23
3.1 Introduction	24
3.2 Materials and Methods	26
3.2.1 Data Collection	26
3.2.2 Data Alignment.....	29
3.3 Results	34
3.3.1 Space Launch System B1	35
3.3.2 Falcon Heavy.....	39
3.3.3 Delta IV Heavy	42
3.3.4 Vulcan Centaur.....	44
3.3.5 Falcon 9.....	46
3.3.6 Atlas V.....	48
3.3.7 Small-lift launch vehicles.....	52
3.4 Discussion	54
3.5 Acknowledgements	55
3.6 Data Availability Statement	56
Chapter 4: Rocket Launch Detection with Smartphone Audio and Transfer Learning	57
4.1 Introduction	58
4.1.1 Transfer Learning with YAMNet.....	59
4.2 Data and Methods	60
4.2.1 Aggregated Smartphone Timeseries of Rocket-generated Acoustics (ASTRA)	60
4.2.2 Smartphone High-explosive Audio Recordings Dataset (SHAREd).....	62
4.2.3 Environmental sound data from ESC-50	63
4.2.4 Preprocessing and splitting	63
4.2.5 Model design	64
4.3 Results	66

4.3.1 Overall performance.....	66
4.3.2 Misclassification analysis	68
4.4 Discussion and Conclusion.....	74
4.5 Data Availability	75
Chapter 5: Conclusion.....	76
5.1 Summary.....	76
5.2 Future Work	77
References	78

List of Figures & Tables

Figure 2.1 Visualization of sensor locations (black triangles) relative to the blast pad location (red star). The launch site of the balloon is marked by a blue dot, and the balloon's ascending path is shown as a solid blue line.	9
Figure 2.2 Spectrograms showing the short-time Fourier transforms of the (a) short-range and (b) long-range surface smartphone acoustic data, (c) airborne acoustic data from the condenser microphone, (d) airborne smartphone acoustic and (e) acceleration data, and (f) balloon altitude in kilometers for the high-altitude segment of the flight. Time is relative to the acoustic wave peak at the balloon stations and observed arrival times are indicated.	10
Figure 2.3 Results of propagation modeling. Solid gray lines show possible ray paths, the dashed black lines show the propagation paths from the source to the balloon and ground stations, and the station locations are marked with black triangles.	11
Figure 2.4 Waveforms of the signals from surface microphones at (a) 3.1 km and (b) 46.4 km, airborne (c) condenser and (d) smartphone microphones, and (e) the smartphone accelerometer. The modeled waveform evolutions are plotted as orange lines in panels (b)–(d). Time is relative to the acoustic wave peak at each location.	13
Figure 2.5 Spectrograms showing the continuous wavelet transforms of the acoustic signals. From top to bottom, the rows display: the smartphone microphone data from the surface at 3.1 km (a–b) and 46.4 km (c–d), the airborne condenser microphone data (e–f), and the airborne smartphone microphone data (g–h).	15
Figure 2.S1 Waveforms of the airborne (a) smartphone microphone, (b) condenser microphone, and (c) infraBSU microbarometer data.	20
Figure 2.S2 Waveforms of the airborne smartphone acceleration data along (a) the x-axis, (b) the y-axis, and (c) the z-axis.	21
Figure 2.S3 Spectrograms and waveforms of stratospheric airborne data from (a-b) the condenser microphone, (c-d) the smartphone microphone, and (e-f) the smartphone z-accelerometer, which shows a significantly higher signal to noise ratio than the acoustic data at this scale.	22
Figure 3.1 Bar plots of the distribution in the dataset of signals (a) recorded on different makes and models of smartphones, (b) originating from different types of rockets, and (c) collected in different range categories.	29
Figure 3.2 Normalized waveforms collected at three different stations during NASA's Artemis I launch, plotted relative to the start-aligned estimated time of arrival at each station. A vertical green line indicates the arrival time estimate, and labels on the left-hand y-axis indicate the estimated propagation distance of each signal.	31
Figure 3.3 A visualization of the peak alignment estimated arrival time selection process. In panel (a), the normalized, unfiltered waveform is plotted in black, the selection window is indicated by green shading, and a dashed green line marks the closer-range station's estimated arrival time, which is used to place the window. In panel (b), the windowed, bandpassed, and re-normalized waveform is shown, with the selected median-filter window indicated with darker shading. In panel (c), the result of median filtering the absolute value of the bandpassed waveform from panel (b) is shown, with the selected peak indicated by a solid green line.	33

Figure 3.4 Normalized waveforms collected at three different stations during NASA’s Artemis I launch, plotted relative to the peak-aligned estimated time of arrival at each station. A vertical green line indicates the arrival time estimate, and labels on the lefthand y-axis indicate the estimated propagation distance of each signal.	34
Figure 3.5 Normalized, start-aligned waveforms of all signals collected from the Artemis I launch. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.	37
Figure 3.6 Continuous wavelet transform (CWT) power of start-aligned signals collected from the Artemis I launch at 24.4 km (a), 37.6 km (b), and 56.9 km (c) from the launch pad.	38
Figure 3.7 CWT power of peak-aligned signals collected from the Artemis I launch at 24.4 km (a, b), 37.6 km (c, d), and 56.9 km (e, f). All panels show the CWT power of the normalized signals, with the left-hand panels showing frequency on the y-axis and time on the x-axis, and the right-hand panels showing time on the y-axis and frequency on the x-axis.	39
Figure 3.8 Normalized, start-aligned waveforms of all signals collected from launch FH-003 of SpaceX’s Falcon Heavy. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.	40
Figure 3.9 CWT power of start-aligned signals collected from launch FH-003 at 23.0 km (a), 37.1 km (b), and 52.2 km (c) from the launch pad.	41
Figure 3.10 CWT power of peak-aligned signals collected from launch FH-003 at 23.0 km (a, b), 37.1 km (c, d), and 52.2 km (e, f). All panels show the CWT power of the normalized signals, with the left-hand panels showing frequency on the y-axis and time on the x-axis, and the right-hand panels showing time on the y-axis and frequency on the x-axis.	42
Figure 3.11 Normalized, start-aligned waveforms of all signals collected from Delta IV Heavy launch D-389. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.	43
Figure 3.12 CWT power of start-aligned signals collected from launch D-389 at 29.1 km (a) and 43.4 km (b) from the launch pad.	44
Figure 3.13 Normalized, start-aligned waveforms of all signals collected from Vulcan Centaur launch Cert-1 V-001. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.	45
Figure 3.14 CWT power of start-aligned signals collected from Vulcan Centaur launch Cert-1 V-001 at 26.6 km (a), 34.2 km (b), and 49.1 km (c) from the launch pad.	46
Figure 3.15 Normalized, start-aligned waveforms of all signals collected from SpaceX Falcon 9 launch SL-G607. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.	47
Figure 3.16 CWT power of start-aligned signals collected from launch SL-G607 at 20.0 km (a), 32.2 km (b), and 51.5 km (c) from the launch pad.	48
Figure 3.17 Normalized, start aligned waveforms of all signals collected from Atlas V launch AV-093. Labels on the lefthand y-axis indicate the estimated propagation distance of each signal.	49
Figure 3.18 CWT power of start-aligned signals collected from launch AV-093 at 20.5 km (a), 34.1 km (b), and 51.9 km (c) from the launch pad.	50

Figure 3.19 Normalized, start-aligned waveforms of all signals collected from Atlas V launch AV-104. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.	51
Figure 3.20 CWT power of start-aligned signals collected from launch AV-104 at 21.6 km (a), 34.2 km (b), and 53.9 km (c) from the launch pad.	52
Figure 3.21 Normalized, start-aligned waveforms of signals collected from Terran 1. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.	53
Figure 3.22 CWT power of start-aligned signals collected from Terran-1 at 17.1 km (a) and 26.9 km (b) from the launch pad.	53
Figure 4.1 Bar plots of the distribution of signals in the high-confidence subset of ASTRA (a) recorded on different makes and models of smartphones, (b) originating from different types of rockets, and (c) collected in different range categories.	62
Figure 4.2 Pie charts showing the distributions of YAMNet’s predicted classes for (a) rocket launch audio from ASTRA, (b) ‘explosion’ (50%) and ‘silence’ (50%) audio from SHAReD, (c) pre-event audio from ASTRA, and (d) environmental audio recordings from the ESC-50 dataset. All data were upsampled to 16 kHz from 800 Hz sampling rate.	65
Figure 4.3 The confusion matrix of the model on all 15 test sets. Each quadrant shows the relevant rate as a percentage, as well as the total number of samples in the category in parentheses.	67
Figure 4.4 The confusion matrix of the best performing split. Each quadrant shows the relevant rate as a percentage, as well as the total number of samples in the category in parentheses.	68
Table 4.1 All events with ‘rocket’ samples misclassified as ‘noise’ by the best split model.	69
Table 4.2 All events with ‘noise’ samples misclassified as ‘rocket’ by the best split model.	70
Figure 4.5 Plot showing the best split model’s unfiltered performance on all signals from the SpaceX Falcon 9 launch, Transporter-4. Green lines indicate ‘rocket’ classification scores >0.6.	71
Figure 4.6 Plot showing the best split model’s median-filtered performance on all signals from the SpaceX Falcon 9 launch, Transporter-4. Green lines indicate ‘rocket’ classification scores >0.6.	72
Figure 4.7 A horizontal bar plot representing the best split model’s false positive classifications of ESC-50 data, plotted according to class labels included with ESC-50. Light colored bars represent the false positive rates before thresholding and median filtering, and dark colored bars represent the same rates after thresholding and median filtering are applied. Results for individual classes are shown in blue and labeled with black text, while the overall results are shown in orange and labeled with white text.	73

List of Abbreviations

ASTRA	Aggregated Smartphone Timeseries of Rocket-generated Acoustics
ESC	Environmental Sound Classification
ETI	Consortium for Enabling Technologies and Innovation
CWT	Continuous Wavelet Transform
HIGP	Hawai‘i Institute of Geophysics and Planetology
IMS	International Monitoring System
INL	Idaho National Laboratory
LSECE	Large Surface Explosion Coupling Experiment
MTV	Consortium for Monitoring, Technology, and Verification
NASA	National Aeronautics and Space Administration
NNSA	National Nuclear Security Administration
NNSS	Nevada National Security Site
SHAReD	Smartphone High-explosive Audio Recordings Dataset
SLS-B1	Space Launch System Block 1
SNL	Sandia National Laboratories
STFT	Short Time Fourier Transform
TNT	Trinitrotoluene
ULA	United Launch Alliance
YAMNet	Yet Another Mobile Network

Chapter 1: Introduction

Along with audible sound, rocket launches generate infrasound: acoustic waves with frequencies below 20 Hz, the lower frequency limit of human hearing. As those waves propagate, they lose information to atmospheric attenuation, the effects of which increase with increasing frequency. This dependence on frequency results in a wave's higher frequency content being disproportionately lost as the wave propagates, eventually leaving only infrasound, which can remain detectable after traveling thousands of kilometers for signals from sufficiently energetic sources. Despite the earliest examples of recorded acoustic rocket signatures dating back to the early days of spaceflight (Balachandran & Donn, 1971; D. Cotten & Donn, 1971; D. E. Cotten et al., 1971; Donn et al., 1968; Kaschak et al., 1970), the relative rarity of rocket launches has resulted in our understanding of these signatures lagging behind that of better-studied infrasound sources such as earthquakes and explosions.

In recent decades, the dramatic increase in annual rocket launches has accelerated progress in this field due to the resulting increase in the volume of rocket launch acoustic data available and number of studies conducted (Blom et al., 2016; Evers et al., 2018; Mclaughlin et al., 2000; Olson, 2012; Pilger, Hupe, et al., 2021; Pilger & Hupe, 2024; Smith et al., 2018; Tenney et al., 2003). Due to the sparsity of acoustic collection stations, however, most of the currently available data were collected at great distances (>1000 km) from the source and are highly attenuated and dispersed. Thus, to understand the chronology of rocket launch signatures, more near field data are required.

The overarching goal of this dissertation is to develop a fast and accurate method of detecting the acoustic signatures generated by rocket launch sequences, including ignition, launch, and trajectory signals. To achieve this goal, several requirements must be met. Firstly, to achieve rapid (<5 minute) detection, we must limit propagation distances to less than 100 km. In application, this level of density in a network would likely be cost-prohibitive using traditional sensors, and thus relatively low-cost,

low-maintenance sensors must be employed. In Chapter 2, recordings of a chemical explosion signal collected by both traditional infrasound sensors and smartphone microphones are compared, serving as proof-of-concept for using smartphones to collect infrasound signals. Through reciprocity, Chapter 2 also provides insight into the expected changes in characteristics of an infrasound signal generated by an elevated source and collected on the ground, as the signal being examined originated on the surface and was collected by a balloon in the stratosphere.

Secondly, a great deal more relatively near-field (<100 km propagation) acoustic rocket signatures than are currently available are necessary, both to improve our understanding of the characteristics of such signatures as well as to train and test machine learning models. Chapter 3 addresses this lack of publicly available rocket acoustic data by releasing an aligned and annotated collection of 1089 smartphone microphone recordings of rocket launches from the Kennedy Space Center and Cape Canaveral Space Force Station. Preliminary time-frequency analysis of the signatures is also presented, and features of the signatures are identified.

Lastly, in Chapter 4, the rocket acoustic signatures discussed in the previous chapter are used to train and test a machine learning model designed for near-real-time rocket launch detection. The performance of the model is evaluated, and its suitability for real-world applications is assessed. In Chapter 5, the results and conclusions of the previous chapters are summarized, and possible directions of future work are discussed.

Chapter 2: Acoustic Waves From a Distant Explosion Recorded on a Continuously Ascending Balloon in the Middle Stratosphere

Published as: Popenhagen, S.K., Bowman, D.C., Zeiler, C., and Garcés, M.A. (2023). Acoustic Waves From a Distant Explosion Recorded on a Continuously Ascending Balloon in the Middle Stratosphere. *Geophysical Research Letters*, 50, 20.

Abstract:

A helium-filled mylar balloon carrying a smartphone and infrasound sensors ascended to a stratospheric height of 35 km over the surface detonation of a chemical explosive, with a total acoustic propagation distance of 127 km. The smartphone was configured to collect multi-modal data at high rates from internal sensors. Analysis of the data shows successful collection of the explosion signal by both the smartphone's microphone and its accelerometers, the first from an ascending balloon. Comparison of the acoustic signal with that collected by other infrasound sensors, both airborne and ground-based, provides insight into the possibilities and limitations of collecting acoustic data from the stratosphere.

2.1 Introduction

Large-scale events, such as earthquakes, volcanic eruptions, and explosions, produce low-frequency acoustic waves called infrasound, which carry information about their source as well as the medium through which they travel. Since acoustic attenuation generally increases with increasing frequency, infrasound waves can propagate over much greater distances than audible sound. Ground-based monitoring of infrasound has been widely used for decades, successfully detecting signals from numerous sources including earthquakes, tsunamis, volcanoes, bolides (Garces & Le Pichon, 2009) and explosions (Pilger, Hupe, et al., 2021). Ground-based infrasound monitoring is also used (alongside other detection methods) by the global International Monitoring System of the Comprehensive Nuclear-Test-Ban Treaty (Vergoz et al., 2022).

Airborne infrasound, on the other hand, has not been as extensively studied or utilized. Although balloon-based sensors were used in early efforts to detect infrasound signals in the 1950s (Peebles, 1997; Weaver & McAndrew, 1995) and 1960s (Wescott, 1964), relatively few studies were conducted until the 2010s saw a resurgence of research in the area. In the years since, airborne infrasound sensors have been used to successfully capture signals from various sources, including lightning (Lamb et al., 2018), earthquakes (Brissaud et al., 2021), chemical explosions (Bowman & Albert, 2018; Bowman & Krishnamoorthy, 2021; Silber et al., 2023; Young et al., 2018), and volcanic eruptions (Podglajen et al., 2022).

We present the first observations of surface blast signals recorded by sensors in an ascending stratospheric balloon. The signals were collected by multiple sensors onboard the balloon, which was 35.2 km above the geoid and at a total propagation distance of 126.8 km from a 1,000 kg TNT-equivalent surface chemical explosion. Explosion signals have previously been collected from balloons in the upper stratosphere (Young et al., 2018), but they are few in number. Thus, careful consideration

was made of the expected properties of an acoustic signal from the surface collected in the stratosphere. The balloon's ascent must also be considered, as similar observations are usually made from passively drifting balloons during periods of neutral buoyancy (Bowman et al., 2022; Bowman & Albert, 2018; Bowman & Krishnamoorthy, 2021; Young et al., 2018). Finally, we show that the vertical accelerometer on the smartphone is able to distinguish the acoustic signal from wind noise recorded during the ascent.

We first consider atmospheric attenuation. The degree of attenuation of an acoustic wave propagating through the atmosphere increases with the frequency of the wave, resulting in low attenuation for infrasonic signals. However, it also depends on atmospheric conditions such as temperature and humidity, which can be considered approximately constant when horizontal changes in position are small ($<1,000$ km) and altitude is held constant. As attenuation tends to increase with increasing altitude (Sutherland & Bass, 2004), we expect to see a greater loss of energy in airborne acoustic data than in data from surface stations at comparable range.

Second, ambient infrasonic noise varies substantially with altitude. Noise levels are high in the troposphere (<10 km above the geoid), with natural and anthropogenic noise persistently present in near-surface acoustic data (Christie & Campus, 2010). In the stratosphere (≈ 10 – 50 km above the geoid), relatively little of this noise remains (Bowman & Lees, 2015). Unlike in the troposphere, where colder, denser air overlying warmer, thinner air leads to turbulent overturning and other meteorological and topographical effects, the stratosphere has a more stable stratification, less ambient noise, and more importantly, a low-velocity zone that efficiently traps sound in an elevated waveguide (Garcés et al., 1998).

Lastly, we consider possible sources of additional noise in the stratosphere. Noise due to cable vibrations (strumming) is expected, and other possible relative movements of the balloon system

should also be considered, primarily the ability of the payload box to spin (Garcés et al., 2022). Acoustic signals produced by aircraft may also be present (Wescott, 1964). The primary noise concern for an ascending balloon, however, is additional wind noise due to the payload box traveling in the balloon's wake (Barat et al., 1984). Research has shown that wind noise dominates during a balloon's ascent (Krishnamoorthy et al., 2020), thus previous efforts to collect explosion signals have focused on free-floating balloons during periods of neutral buoyancy (Bowman & Albert, 2018; Bowman & Krishnamoorthy, 2021; Young et al., 2018). Studies of stratospheric noise, however, have shown that wind noise during ascent through the stratosphere is drastically reduced compared to the troposphere (Bowman & Lees, 2015). Thus, we show that high signal-to-noise ratios are achievable in the stratosphere even during ascent.

The ability to record sound waves as an airborne platform ascends or descends has important impacts for terrestrial and planetary acoustics. Balloons can “stationkeep” in the vicinity of a target region by rising and falling to take advantage of different wind speeds (Bellemare et al., 2020), but it is not clear whether such maneuvers would prevent concurrent infrasound monitoring. Airborne sensing has also been proposed as a means of recording acoustic waves from seismic activity on Venus (Brissaud et al., 2021), but the previous balloon missions in that planet's atmosphere encountered substantial vertical motion (Sagdeev et al., 1986). Our results suggest that continuous airborne infrasound monitoring of sources of interest on Earth and geophysical activity on Venus (e.g., Brissaud et al. (2021), Rossi et al. (2023)) will be less challenging than previously thought.

2.2 Data and Methods

A high-altitude balloon was deployed near a $1,000 \text{ kg} \pm 10 \text{ kg}$ TNT-equivalent surface chemical explosives test at Nevada National Security Site on 27 October 2020, at 6:37 a.m. local time during the Large Surface Explosion Coupling Experiment (LSECE) (Silber et al., 2023; Wermer et al.,

2021). The balloon's payload consisted of acoustic sensors and a SPOT TRACE asset tracker contained in a foam carton. The sensors consisted of a Samsung S10 smartphone using the RedVox application (Garcés et al., 2022) to collect geophysical data from the phone's internal sensors, including the microphone and three-dimensional accelerometers. In addition, two InfraBSU microbarometers (Marcillo et al., 2012) and a prototype condenser microphone digitized on a DiGOS DATA-CUBE logger at 400 Hz were included. The frequency response of the condenser microphone is unknown. The payload did not have wind noise mitigation screens. The balloon was launched at 5:16 a.m. local time and achieved a stratospheric height of 35.2 km before the signal of interest was detected. The flight system, the various smartphone sensors recording data, and the application's data collection protocols are detailed in Garcés et al. (2022).

A surface network of smartphones running the RedVox app was also collecting data during the experiment. The individual phones were located at various distances from the explosion site, including a station at a horizontal distance of 3.1 km, and a station at a horizontal distance of 46.4 km. Despite several stations deployed at ranges comparable to the total distance (126.8 km) of the balloon from the explosion site at the time of the signal, the station at 46.4 km was the furthest with a confirmed signal arrival. The locations of the surface and airborne stations relative to the explosion site are shown in Figure 2.1, along with the balloon's launch site and flight path. In Figure 2.2, time-frequency representations of the recordings from the smartphones on the surface, two of the airborne acoustic sensors, and the vertical channel of the airborne accelerometer are shown for the high-altitude segment of the flight, along with the balloon's altitude for reference. While the background noise levels of surface stations are mostly constant over the time period, this is not true for the airborne sensors. Around 2,200 s before the airborne acoustic wave peak, we see a clear decrease in the background noise levels, indicating the balloon entered the stratosphere. The acoustic explosion signal is not

obvious on this time scale, but we can see the airborne accelerometer signal clearly in Figure 2.2e. Finally, about 200 s after the acoustic wave peak, the noise levels increase dramatically as the balloon bursts and the payload begins its descent. For a more in-depth view of noise levels in the stratosphere, see Figure 2.S3 in Supporting Information S1.

Ray tracing and waveform evolution modeling were calculated using the open-source *infraGA/GeoAc* software and the “eigenray” methods described in Blom & Waxler (2017). The signal recorded at 3.1 km was used as the input for the evolution modeling. Temperature and wind data were extracted from Ground-to-Space atmospheric specifications (Drob, 2019; Drob et al., 2003) provided by the National Center for Physical Acoustics of the University of Mississippi. A variety of possible ray paths and the calculated eigenrays between the source and the sensors' locations are shown in Figure 2.3. Due to an idiosyncrasy in the software, it was necessary to perform the eigenray search for the path to the balloon in reverse, by setting the balloon as the source and the detonation site as the receiver. To correct for this, the winds were reversed in the atmospheric profile when calculating this eigenray. Due to reciprocity, we don't expect this to have any effect on the modeled propagation time.




Figure 2.1 Visualization of sensor locations (black triangles) relative to the blast pad location (red star). The launch site of the balloon is marked by a blue dot, and the balloon's ascending path is shown as a solid blue line.




Figure 2.2 Spectrograms showing the short-time Fourier transforms of the (a) short-range and (b) long-range surface smartphone acoustic data, (c) airborne acoustic data from the condenser microphone, (d) airborne smartphone acoustic and (e) acceleration data, and (f) balloon altitude in kilometers for the high-altitude segment of the flight. Time is relative to the acoustic wave peak at the balloon stations and observed arrival times are indicated.




Figure 2.3 Results of propagation modeling. Solid gray lines show possible ray paths, the dashed black lines show the propagation paths from the source to the balloon and ground stations, and the station locations are marked with black triangles.

Time-frequency analysis of the signals was performed by computing the continuous wavelet transform (CWT) of eleven-second-long windows centered on the absolute maximum (in bits) of the acoustic signal. The CWT transforms the one-dimensional waveform into two-dimensional coefficients representing scale and position by maximizing the similarity between the signal and scaled and shifted versions of a specified wavelet function. Like the short-time Fourier transform (STFT), the CWT is a windowed transform. Unlike the STFT, however, the size of the window varies with frequency to maximize the resolution of both time and frequency. This improves frequency resolution at low frequencies and time localization at high frequencies compared to the STFT. Choice of wavelet

varies between uses. The CWT used in this work is constructed according to the standardized constant-Q variation of the Gabor atom detailed in Garcés (2020), with band order $N = 6$. The time-frequency representations in this work are shown in bits, a binary logarithmic unit more natural to digital systems than the traditional decibel (Garcés, 2020, 2023).

2.3 Results

2.3.1 Acoustic Data

The observed waveforms from the surface and airborne sensors are shown in Figure 2.4, along with the results of infraGA's weakly non-linear waveform evolution modeling along the calculated eigenray. All the waveforms are unfiltered, and all the acoustic waveforms are normalized. Looking at the results from the two airborne audio sensors, we see that both detected a signal at 13:43:46 UTC, at which time the location data shows a height of 35.2 km and a total distance from the explosion of 126.8 km. This timing is consistent with propagation modeling results within 5 s, corresponding to a true propagation path about 2 km longer than the modeled one, or a true speed of sound slightly slower than that estimated by the modeling software.




Figure 2.4 Waveforms of the signals from surface microphones at (a) 3.1 km and (b) 46.4 km, airborne (c) condenser and (d) smartphone microphones, and (e) the smartphone accelerometer. The modeled waveform evolutions are plotted as orange lines in panels (b)–(d). Time is relative to the acoustic wave peak at each location.

The two airborne audio signals differ. The first difference is due to the sensitivities of the microphones. The condenser microphone was able to collect the signal in significantly more detail compared to the smartphone microphone, which was approaching its sensitivity limit. This can be seen in the waveform itself, where individual bits of information can be observed visually. There is

also a difference in phase between the two signals, likely due to the phase change of the smartphone in the 2–10 Hz passband (Asmar et al., 2019). The smartphone frequency response at the temperatures and pressures in the stratosphere is unknown and beyond the scope of this paper. However, theory suggests that the infraBSU microbarometer (see Figure 2.S1 in Supporting Information S1) frequency response should be stable at high altitudes in the passband of interest (see Equation 4 in Marcillo et al. (2012)). The phase of the infraBSU microbarometer signal matches that of the condenser microphone signal, indicating that the phase of the condenser microphone signal is correct.

For this to be true, the classic “blast wave” signal recorded on the ground must have experienced a phase shift during its journey into the middle stratosphere. This is supported by infraGA's weakly linear waveform evolution modeling, as the shape of the resulting waveform (seen in Figures 2.4c and 2.4d) closely resembles the condenser microphone signal. As the wave followed a direct path from source to receiver, this distortion was unexpected (Bowman & Krishnamoorthy, 2021), but explainable. We can attribute this phase shift to Hilbert transforms resulting from the acoustic wave being trapped near turning points in the tropopause temporarily before escaping into the stratosphere. The nearly horizontal path in the calculated eigenray between 20 and 50 km (see Figure 2.3) and the modeled waveform evolutions (see Figure 2.4) mentioned above both support this conclusion.

For further confirmation of the validity of the airborne detection, time-frequency analysis was performed on all the acoustic signals. The resulting multiresolution spectrograms are shown in Figure 2.5. From the CWT of the signal collected on the ground at a horizontal distance of 3.1 km from the explosion shows (Figures 2.5a and 2.5b), we can identify two high energy regions in the frequency domain. The first, near 5 Hz, is clearly present in all the acoustic signals. The presence of the second, near 10 Hz, is less obvious in the signals collected at greater range, particularly in the

signals from the airborne sensors. This is consistent with the expected loss of energy to attenuation during propagation, as lower frequency energy is less effected, and attenuation is greater at higher altitudes. Additionally, we see that despite the phase difference between the two airborne acoustic signals, their energy distributions are nearly identical in the frequency domain (Takazawa, Popenhagen, Ocampo Giraldo, Cárdenas, et al., 2024). Thus, time-frequency analysis of the acoustic signals provides additional evidence of successful airborne detection.




Figure 2.5 Spectrograms showing the continuous wavelet transforms of the acoustic signals. From top to bottom, the rows display: the smartphone microphone data from the surface at 3.1 km (a–b) and 46.4 km (c–d), the airborne condenser microphone data (e–f), and the airborne smartphone microphone data (g–h).

2.3.2 Acceleration Data

The airborne acceleration data in the z-direction is shown in Figure 2.4e. A sudden increase in energy at the moment of the signal's arrival is clearly seen in the second half of the shown window. Damped oscillation continues for approximately 5 s post-arrival, after which the amplitude returns to pre-arrival levels. This behavior is most obvious in the z-direction, but an increase in energy was also seen in the x- and y-directions (see Figure 2.S2 in Supporting Information S1). We see in Figure 2.2e that this sudden increase in energy was not observed at any other point throughout the stratospheric flight, providing further evidence of the acoustic wave arrival.

2.4 Discussion and Conclusions

Using data collected by multiple sensors, we successfully observed the arrival of the acoustic wave generated by an explosion from a stratospheric height of 35.2 km and a total range of 126.8 km from the site of the explosion. The balloon carrying the sensors was still ascending when the signal was collected, as evidenced by location data from both the smartphone and a Balloon Ascent Technologies HAB Boulder Balloon Cut-Down Device included in the balloon's payload, and supported by the smartphone's barometer and accelerometers. Propagation modeling predicted the arrival at this location within a few seconds of the observation.

Comparison of the airborne audio data and data collected on the ground at 46.4 km and at 3.1 km revealed disproportionate loss of energy above 10 Hz frequency as the acoustic wave propagated through the atmosphere. This confirms the expected increase in attenuation of acoustic waves with frequency. The effect of the attenuation was significantly greater in the airborne data than in the ground-based data, which was also expected, as attenuation generally increases with altitude, and the propagation path to the balloon was significantly longer. It may also suggest other nonlinear effects enhanced by stratospheric conditions.

Nevertheless, the energy of the signal at frequencies below 9 Hz, particularly the peak near 5 Hz seen in the short-range surface data, was clear in the airborne audio data collected by both the smartphone and the condenser microphone, while the surface stations at comparable ranges were unable to detect the signal at all. Given the decrease in signal-to-noise ratio from the surface station at 3.1 km to that at 46.4 km, it is not surprising that the signal was not detected at surface ranges comparable to the propagation distance of the balloon. These results demonstrate that infrasound signals can not only be successfully collected from an ascending balloon at sufficiently high altitudes, but can be successfully collected in the stratosphere at propagation distances much greater than the detectable range of the signal on the ground. The ability to collect from a balloon while it's ascending could potentially lengthen the signal collection window of future balloon deployments by allowing for collection during periods of both neutral buoyancy and ascent.

Additionally, analysis of acceleration data at the time of the arrival suggests that accelerometers included with the other sensors carried by a balloon can provide evidence of an acoustic wave arrival. As wave propagation modeling is only possible if the location of the source is known, using acceleration data to discriminate between noise and true acoustic wave arrivals would be a valuable tool for infrasound monitoring applications both terrestrial and extraterrestrial.

2.5 Acknowledgements

The Large Surface Explosion Coupling Experiment (LSECE) would not have been possible without the support of many people from several organizations. The authors wish to express our gratitude to the LSECE working group and multi-institutional and interdisciplinary group of scientists and engineers for the field experiment. This work was supported by the Department of Energy National Nuclear Security Administration under Award Numbers DE-NA0003920 (MTV) and DE-NA0003921 (ETI). Work was also done by Mission Support and Test Services, LLC, under Contract

No. DE-NA0003624 with the U.S. Department of Energy and the NNSA Office of Defense Nuclear Nonproliferation (DOE/NV/03624–1538). This report was prepared as an account of work sponsored by agencies of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. The United States Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. We thank the National Center for Physical Acoustics for providing Ground 2 Space profiles, and the NOAA-SORD team at the Nevada National Security Site for assistance during the balloon launch. We are grateful to Fransiska Dannemann Dugick for reading through the manuscript draft and offering comments.

2.6 Data Availability Statement

The LSECE data set is archived through the International Federation of Digital Seismograph Networks at <https://doi.org/10.7914/1p2e-0167> (IRIS, 2023; Walter et al., 2020). Data recorded using RedVox can be accessed through a report on the RedVox website (RedVox Inc., 2023). Figures

in this work were made with Matplotlib version 3.5.1 (Caswell et al., 2021; Hunter, 2007), available under the Matplotlib license at <https://matplotlib.org/>. Propagation modeling was performed with infraGA (Blom, 2014; Blom & Waxler, 2017). Time-frequency analysis was performed using libquantum version 1.4.1 (RedVox Inc., 2022).

2.7 Supporting Information

This supplementary document contains additional figures supporting the results. Figure 2.S1 shows the waveforms of the acoustic signals collected by the airborne smartphone and condenser microphone (included in article), and the signal collected by the infraBSU microbarometer. In Figure 2.S2, the airborne acceleration signal is shown in all three dimensions (only z-acceleration included in article). Figure 2.S3 shows the spectrograms and waveforms of data from the airborne condenser microphone, smartphone microphone, and the vertical channel of the smartphone accelerometer for the stratospheric segment of the balloon's flight.

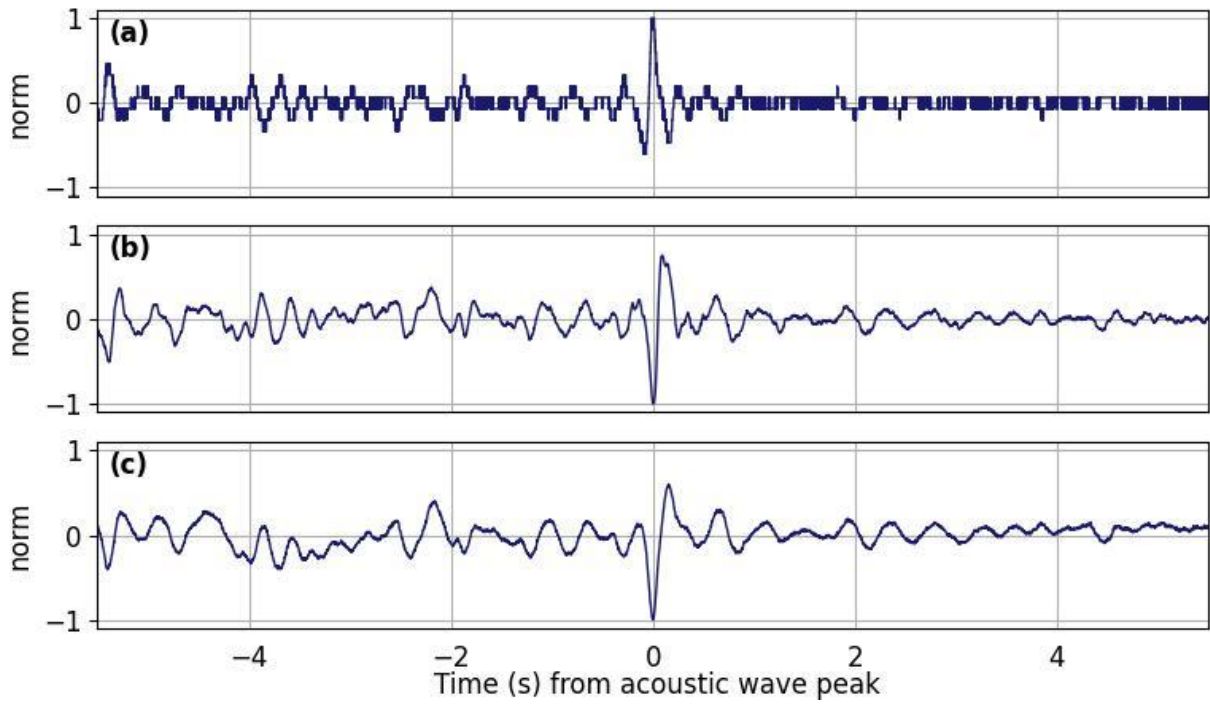


Figure 2.S1 Waveforms of the airborne (a) smartphone microphone, (b) condenser microphone, and (c) infraBSU microbarometer data.

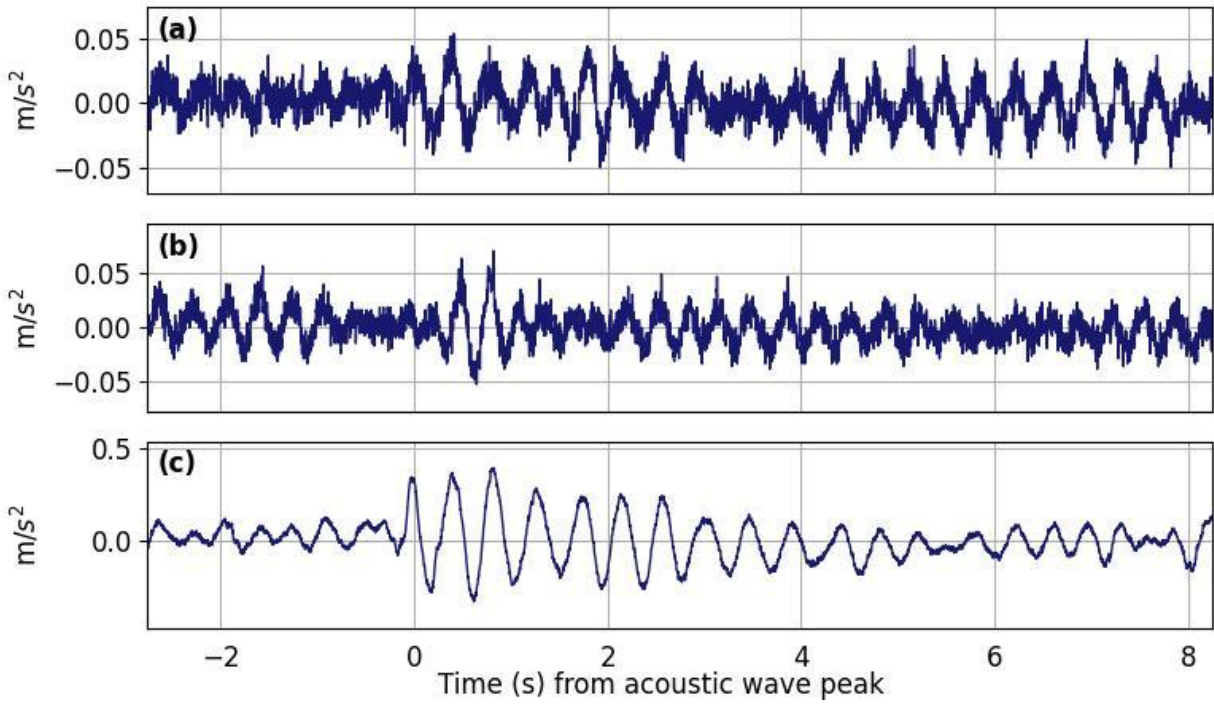


Figure 2.S2 Waveforms of the airborne smartphone acceleration data along (a) the x-axis, (b) the y-axis, and (c) the z-axis.




Figure 2.S3 Spectrograms and waveforms of stratospheric airborne data from (a-b) the condenser microphone, (c-d) the smartphone microphone, and (e-f) the smartphone z-accelerometer, which shows a significantly higher signal to noise ratio than the acoustic data at this scale.

Chapter 3: Acoustic Rocket Signatures Collected by Smartphones

Published as: Popenhagen, S. K. and Garcés, M.A. (2025) Acoustic Rocket Signatures Collected by Smartphones. *Signals*, 6, 5.

Abstract:

Rockets generate complex acoustic signatures that can be detected over a thousand kilometers from their source. While many far-field acoustic rocket signatures have been collected and released to the public, very few signatures collected at distances less than 100 km are available. This work presents a curated and annotated dataset of acoustic signatures of 243 rocket launches collected by a network of smartphones stationed at distances between 10 and 70 km from the launch sites, resulting in 1089 individual recordings. Due to the frequency dependence of atmospheric attenuation and the relatively short propagation distances, higher-frequency features not preserved in most publicly available data are observed. The signals are time-aligned to allow for different segments of the signal (ignition, launch, trajectory, chronology) to be more easily examined and compared. Initial analysis of the features of these rocket launch stages is performed, observed features are compared to those found in the existing literature, and comparisons between signals from launches of different rocket types are made. The dataset is annotated and made available to the public to aid future analysis of the characteristics and source mechanisms of rocket acoustics as well as applications such as rocket detection and classification models.

3.1 Introduction

Acoustic waves traveling through the atmosphere carry information about both the source that generated them and the medium through which they travel. As waves propagate, they lose some of that information to atmospheric attenuation. Atmospheric attenuation, however, does not affect all waves equally. Lower-frequency sound is less attenuated and can remain detectable over much longer distances than higher-frequency signals. Since both natural and anthropogenic large-scale events can generate significant amounts of low frequency energy, this phenomenon is particularly useful for detecting and monitoring such events. Sound waves with frequencies below 20 Hz (the lower frequency limit of human hearing) are called infrasound. A wide variety of events have been successfully detected using infrasound, including tsunamis, volcanoes, bolides, earthquakes (Schwardt et al., 2022), lightning (Lamb et al., 2018), explosions (Pilger, Gaebler, et al., 2021), and rocket launches (Pilger, Hupe, et al., 2021).

Many of these detections were made using infrasound data collected by the International Monitoring System (IMS) of the Comprehensive Nuclear-Test-Ban Treaty (Hupe et al., 2022). The IMS is a global network of stations that collects data from a number of different types of sensors, including infrasound microphones. While the IMS has proven invaluable for global detection of large-scale events, the size of the network necessitates sparseness in order to limit cost, resulting in most signals being detected only after propagating great distances. This trend towards far-field data is not unique to the IMS; many of the infrasound signals collected are collected at large propagation distances. Due to the low attenuation rate of infrasound in the atmosphere, this is not generally a problem for detectability given a sufficiently energetic low-frequency signal, but it severely limits the expediency at which detections can be made due to the relatively slow (compared to light) speed of sound.

Acoustic signals generated by rocket launches have consistently been detected over the decades since humanity's first successful launches (Balachandran & Donn, 1971; Blom et al., 2016; D. Cotten & Donn, 1971; D. E. Cotten et al., 1971; Donn et al., 1968; Evers et al., 2018; Kaschak et al., 1970; Pilger, Hupe, et al., 2021; Pilger & Hupe, 2024; Tenney et al., 2003). Starting in the late 1950s, infrasound waves from rocket launches in Florida were detected by infrasound sensors at multiple stations in the eastern United States at ranges between 291 and 1594 km (D. Cotten & Donn, 1971; Donn et al., 1968; Kaschak et al., 1970). In 1971, Cotton and Donn reported the observation of infrasound signals believed to be shock waves generated by the exhaust plumes of Apollo rockets as they flew over Bermuda at 188 km altitude (Balachandran & Donn, 1971; D. Cotten & Donn, 1971). In the decades following these early observations, the annual number of rocket launches has increased dramatically, as has our ability to detect infrasound signals. Numerous rocket launch signals have been detected by IMS infrasound stations since the network's establishment, usually at great distances. Such signals have been analyzed in multiple studies (Blom et al., 2016; Evers et al., 2018; Mclaughlin et al., 2000; Pilger, Hupe, et al., 2021; Pilger & Hupe, 2024), and thus the characteristics of far-field rocket signatures are generally well understood.

Far fewer studies, however, have been conducted on acoustic rocket launch signals collected at shorter distances (<100 km). The existing literature includes a 2016 study that analyzed acoustic signals collected 7–100 km from the launch site of a four-stage sounding rocket (Blom et al., 2016), and results of a few studies analyzing acoustic signals collected at similar distances can be found in conference proceedings (Olson, 2012; Smith et al., 2018; Tenney et al., 2003), but few data are currently publicly available. This relative dearth of nearer-field data limits advancement in the field of rocket acoustics for two related reasons. First, due to the atmospheric propagation effects previously discussed, much of the original signal's content is lost as it propagates, including features seen in data

from previous studies (Blom et al., 2016; Mclaughlin et al., 2000; Olson, 2012; Tenney et al., 2003), limiting what can be learned about the source. Second, distortion of the signal during propagation makes relationships between characteristics and their source mechanisms more difficult to interpret in signals collected at longer ranges. Thus, we believe there is much to learn from analysis of nearer-field rocket launch acoustics.

In order to conduct such an analysis, however, it is first necessary to curate a dataset of such signals. To this end, we established a semi-permanent network of low-cost, attritable sensors at a 10–70 km range from launch sites at the Kennedy Space Center and Cape Canaveral Space Force Station in Merritt Island and Cape Canaveral, FL, USA. From the first launch recorded on 24 May 2019 until the last on 20 August 2024, the network recorded 1089 acoustic signals from 243 rocket launches. These recordings have been aggregated into a curated, labeled dataset, hereafter referred to as ASTRA (Aggregated Smartphone Timeseries of Rocket-generated Acoustics) for convenience (Popenhagen, 2024). ASTRA is an open-access dataset created to help increase our understanding of the nature and features of rocket acoustics through traditional analysis, as well as for building, training, and testing near-real-time rocket detection and identification models using machine learning methods. We present an overview of ASTRA, the network on which it was collected, the alignment and labeling procedures used, and a summary of the observed time–frequency characteristics.

3.2 Materials and Methods

3.2.1 Data Collection

The data used in this study were collected by a network of smartphones located 10–70 km from the launch sites. Smartphones are used in this study for several reasons, first of which is their low cost compared to traditional infrasound sensors, which allowed for a much denser network. In addition, smartphones are attritable, user-friendly, and able to run machine learning models in the

field, which may prove valuable for near-real-time detection and identification studies in the future (Kong et al., 2020; Takazawa, Popenhagen, Ocampo Giraldo, Hix, et al., 2024). Smartphones are also commercially off-the-shelf (COTS) devices, thus stations can be replaced quickly and easily by simply buying a new smartphone on location, which could be valuable for monitoring applications.

There are important limitations to keep in mind when working with smartphones, however. Firstly, ground-truth information about a smartphone's on-board microphone is usually unavailable, as the information is considered proprietary. In addition, we know through observation that the model of microphone used can differ between smartphone models and even individual devices of the same make and model (Slad & Merchant, 2021). The frequency response of smartphones can and has been studied (Asmar et al., 2019; Brown & Evans, 2011; Takazawa, Popenhagen, Ocampo Giraldo, Cárdenas, et al., 2024); however, it is impossible to know the exact frequency response of a specific smartphone without calibration due to the microphone model being unknown. Requiring each smartphone to be individually calibrated would negate many of the previously discussed strengths of using a smartphone network. Thus, we instead leave all the phones uncalibrated, display normalized amplitudes to avoid assigning meaning to uncalibrated values, and focus on the signal-to-noise ratio and information in the time–frequency domain, in which important characteristics of acoustic signals (such as shape) have been shown to remain stable between traditional infrasound microphones and smartphones (Popenhagen et al., 2023; Takazawa, Popenhagen, Ocampo Giraldo, Cárdenas, et al., 2024) over the relatively unstable shape and amplitude of the signal in the time domain. However, possible inconsistencies due to the lack of calibration must still be kept in mind during analysis. For those interested, examples of how signals collected by smartphone microphones compare to those collected by traditional sensors can be found in previous studies (Asmar et al., 2019; Brown & Evans, 2011; Popenhagen et al., 2023; Slad & Merchant, 2021; Takazawa, Popenhagen, Ocampo Giraldo,

Cárdenas, et al., 2024), and the raw amplitudes of the waveforms as well as the makes, models, and unique identification numbers of the smartphones are included in the dataset.

Each of the phones collected data from multiple internal sensors through the RedVox Android application, the specifications of which are detailed in Garcés et al. (2022). While data from multiple sensors can be collected using the RedVox application, the only sensor relevant to this work is the microphone. The smartphone microphones collected acoustic data at a sampling rate of 800 Hz and uploaded the data in near-real time, after which they were downloaded from the cloud using the Redvox SDK and processed in Python using open-source packages. The makes and models of the smartphones represented in ASTRA are shown in Figure 3.1 along with the distributions of rocket types and propagation ranges.

The data were largely recorded by Samsung Galaxy models (see Figure 3.1a), as the network was originally deployed with Galaxy S8 phones, which were replaced in time with S10s, and finally with S20s. The other models represented in ASTRA originate from instances where a single station was replaced or when other smartphones were temporarily added to the network. As the network's day-to-day operation was performed by citizen scientists, not every phone in the network was turned on and recording for every launch, resulting in some individual launches being recorded by only a subset of the phones.

The vast majority of the signals (949 out of 1089) are of SpaceX Falcon 9 launches (see Figure 3.1b), as Falcon 9s were the most commonly launched rocket during the collection period. ULA Atlas V and SpaceX Falcon Heavy launches were also prevalent during the collection period, resulting in 77 and 40 collected signals, respectively. All other rocket types represented in ASTRA have less than 10 collected signals, as most were launched only once in the collection period.

The range distribution is more even than the distributions previously discussed, with 390, 475, and 224 signals in the 10–30 km, 30–50 km, and 50–70 km categories, respectively (see Figure 3.1c). However, the distribution is skewed slightly towards the 30–50 km range category, and the 50–70 km range category is underrepresented. These distributions are considered when analyzing trends in the dataset.

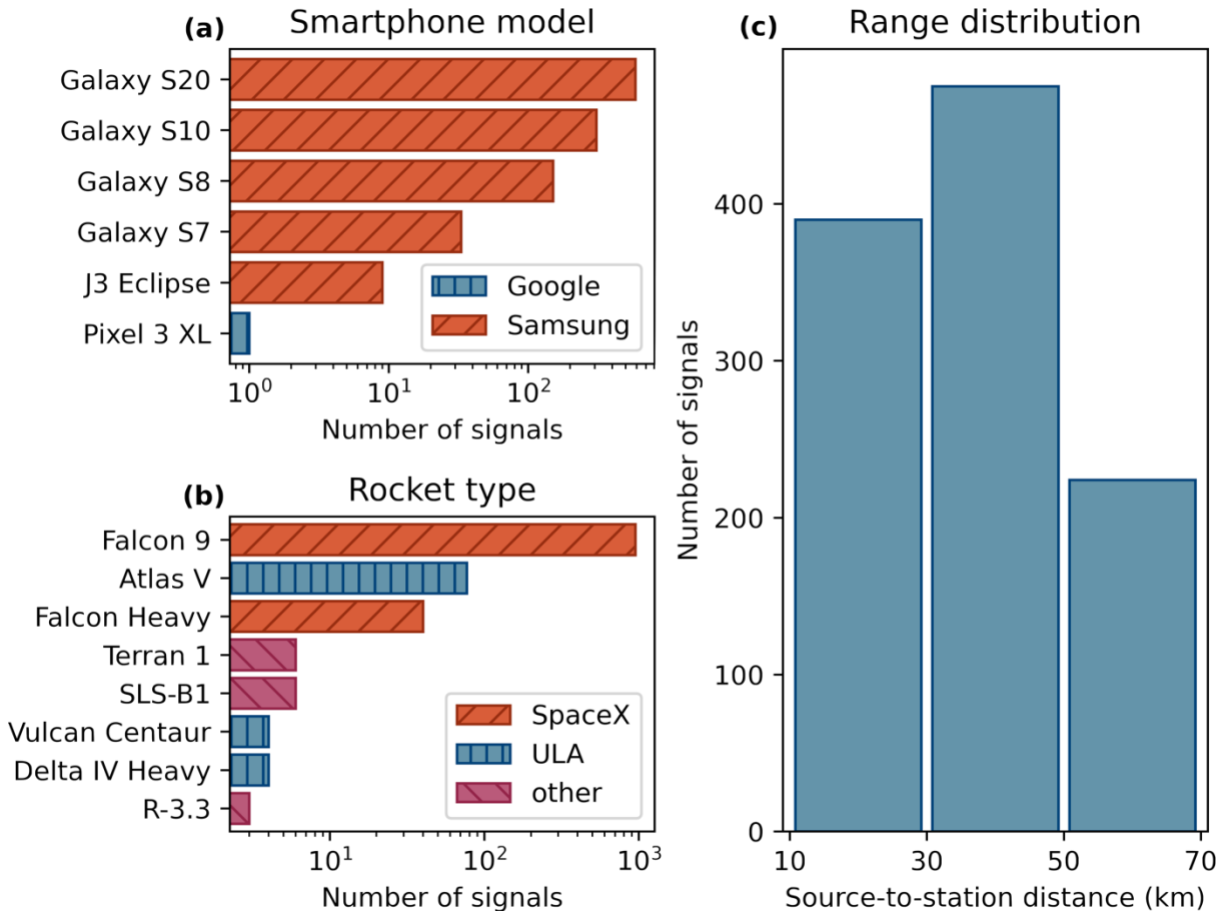


Figure 3.1 Bar plots of the distribution in the dataset of signals (a) recorded on different makes and models of smartphones, (b) originating from different types of rockets, and (c) collected in different range categories.

3.2.2 Data Alignment

For each rocket launch, the reported time of the launch and the great circle distances between the source and each active station are used to compute estimates of the signal’s arrival time at each station using an estimated speed of sound. There are two likely sources of error in these estimates that

we will address. Firstly, the reported launch time is often accurate only to the minute. Secondly, we assumed direct great circle paths from the source to each phone, which, while a reasonable estimate for short propagation distances, may not be an accurate representation of acoustic wave propagation. We expect that error in arrival times arising from simplified propagation paths will increase with distance from the source, while bias error due to imprecisely or inaccurately reported launch times will be independent of distance.

For each signal in ASTRA, the effective speed of sound is estimated according to

$$c_{\text{eff}} = c + v_{\text{parallel}}, \quad (3.1)$$

using the mean temperature T in degrees Celsius and mean parallel (or negative if antiparallel) wind along the great circle path between the source and the station. Temperature and wind values were obtained from Copernicus Climate Change Service’s ERA5 (Hersbach et al., 2018), and parallel wind is calculated by projecting the wind vector onto a vector representing the estimated speed of sound in the direction of propagation. The magnitude of the parallel or antiparallel wind is then added or subtracted, respectively, from the magnitude of the estimated speed of sound, giving the effective speed of sound along the assumed path. The estimated travel time is simply the assumed propagation distance divided by the effective speed of sound.

Once the effective sound speed is estimated, the dataset is aligned in two ways, each with a different estimate of the arrival time of the signal. Both estimates of the arrival times are preserved in ASTRA as there are use cases where each estimate may prove more useful than the other. These estimates, their strengths, and their potential weaknesses are outlined in the following paragraphs.

The first arrival time estimate is made by simply adding the estimated travel time to the reported launch time. We call data aligned this way “start-aligned”, and the associated arrival time estimate is labeled as such in ASTRA. The start-aligned estimation is independent of the waveforms

themselves and thus unaffected by noise in the data, but it is sensitive to bias error in the reported launch time and assumed propagation path. An example of start-aligned data is shown in Figure 3.2.




Figure 3.2 Normalized waveforms collected at three different stations during NASA’s Artemis I launch, plotted relative to the start-aligned estimated time of arrival at each station. A vertical green line indicates the arrival time estimate, and labels on the left-hand y-axis indicate the estimated propagation distance of each signal.

To calculate the second arrival time estimate, a window from which to select a peak is determined. The window duration is equal to 180 s plus 25% of the estimated propagation time to account for the uncertainty increasing with propagation distance as previously discussed. For each individual launch, estimates are made starting with the closest station to the launch site, for which the first timestamp of the selection window is set to the start alignment estimate of the arrival time at the same station, minus a 30 s buffer to account for potential error in the reported launch time. For successive stations, the next-closest station’s new estimate is used instead of the start-alignment

estimate, and an additional buffer term inversely proportional to the difference between the two propagation distances is subtracted. This is represented mathematically in Equation (3.2).

$$, \tag{3.2}$$

The windowed waveform is then bandpassed. Previous studies (Balachandran & Donn, 1971; Blom et al., 2016; Evers et al., 2018; Kaschak et al., 1970; Pilger, Hupe, et al., 2021; Tenney et al., 2003) have shown that characteristic acoustic frequencies of rocket launch signals are higher than 0.5 Hz, and preliminary observation of the data showed a decrease in the signal-to-noise ratio below approximately 50 Hz, thus the frequency limits of the passband are set to 0.5 Hz and 50 Hz. As the main phase of the rocket launch signature persists over a timescale on the order of 10 s to 100 s, a rolling median filter is applied to the absolute value of the windowed and bandpassed waveform with a kernel duration of 15 s. The peak of the median-filtered time series is then determined, the corresponding epoch time of which is the second arrival time estimate. We call data aligned this way “peak-aligned”, and the associated arrival time estimate is labeled as such in ASTRA. Aligning the dataset this way is useful for observing characteristics of the signature that are more stable over the propagation path, such as the frequency of peak energy. Unlike start alignment, the accuracy of peak alignment is partially dependent on the signal-to-noise ratio of the waveform, but it is less reliant on the accuracy of the travel time estimation and reported launch time. The peak selection process is visualized in Figure 3.3, and an example of peak-aligned data is shown in Figure 3.4.

Which type of alignment should be used depends on which phase of the rocket launch sequence is of interest. For example, if analyzing the ignition signal, start alignment is more useful, as propagation effects cause the signal to become more emergent as it travels and peak-aligning the signals will misalign the ignition signals (see Figure 3.4). In contrast, if the goal is to analyze the characteristics of the signal generated during the rocket’s ascent, peak alignment may be more useful.




Figure 3.3 A visualization of the peak alignment estimated arrival time selection process. In panel (a), the normalized, unfiltered waveform is plotted in black, the selection window is indicated by green shading, and a dashed green line marks the closer-range station's estimated arrival time, which is used to place the window. In panel (b), the windowed, bandpassed, and re-normalized waveform is shown, with the selected median-filter window indicated with darker shading. In panel (c), the result of median filtering the absolute value of the bandpassed waveform from panel (b) is shown, with the selected peak indicated by a solid green line.




Figure 3.4 Normalized waveforms collected at three different stations during NASA’s Artemis I launch, plotted relative to the peak-aligned estimated time of arrival at each station. A vertical green line indicates the arrival time estimate, and labels on the lefthand y-axis indicate the estimated propagation distance of each signal.

3.3 Results

Which type of alignment should be used depends on which phase of the rocket launch sequence is of interest. For example, if analyzing the ignition signal, start alignment is more useful, as propagation effects cause the signal to become more emergent as it travels and peak-aligning the signals will misalign the ignition signals (see Figure 3.4). In contrast, if the goal is to analyze the characteristics of the signal generated during the rocket’s ascent, peak alignment may be more useful.

Once the signals in the dataset are aligned, characteristics of the signature are more easily observed. Through comparison of recordings of the same launch collected at different distances, the effects of propagation on the signal are apparent. As previously mentioned, different frequency

components of an acoustic signal can travel through the atmosphere at slightly different speeds, leading to dispersed signals with more emergent onsets at greater distances from the source. We also expected to observe frequency-dependent energy loss as the propagation distance increases due to atmospheric attenuation.

In addition to analyzing the differences between signals originating from the same source but collected at different distances, signals generated by different types of rockets can also be examined for differing characteristics. ASTRA includes signals from launches of eight different rocket types with lift classes ranging from small (Terran 1, Rocket 3.3) to super heavy (Space Launch System B1) with varying fuel types and numbers of engines. The results are organized by rocket type, starting with the heaviest rockets and proceeding to the lightest.

3.3.1 Space Launch System B1

NASA's Space Launch System B1 (SLS-B1) is a super-heavy-lift launch vehicle and the most powerful operational rocket to date. Its first stage is powered by four RS-25 liquid-fueled engines and two solid rocket boosters. It is represented in ASTRA by a single launch (Artemis I), from which six stations collected acoustic signals. The waveforms of all six signals can be seen in Figure 3.5. The waveforms collected less than 30 km from the source are characterized by a spindle-like shape, high chronological symmetry about the peak amplitude, short duration of the highest amplitude segment (~60 s) compared to longer-range signals, and the presence of a small transient signal preceding the high amplitude segment. The waveform collected at a 37.6 km range shows significant elongation but maintains some similarity in shape with the shorter-range data despite early signs of peak separation. In contrast, the shapes of the waveforms collected at a greater than 50 km range are significantly different, with all three showing two chronologically distinct high-energy regions separated by approximately 150 s, both of which appear to further separate into two sub-regions each.

Multiresolution time–frequency analysis was performed on all six signals, and one signal from each of the range categories (10–30 km, 30–50 km, 50–70 km) was selected to illustrate the identified characteristics, as high similarity was observed within each range category. In signals collected between 10 and 30 km from the source, we observe a sharp increase in power across nearly the whole spectrum (see Figure 3.6), but particularly near 10, 20, 45, and 100 Hz. As the signal propagates and the contribution of higher-frequency components is reduced, the peak power band near 10 Hz becomes more dominant. This phenomenon can be seen in Figures 3.6 and 3.7, in which spectrograms of the continuous wavelet transform (CWT) power of the three selected signals are displayed. As the rocket ascends, moving away from the stationary sensor, an apparent decrease in the frequency of the signal is observed. This observed shift in frequency, an example of the Doppler effect, can be seen most clearly in Figures 3.6c and 3.7.




Figure 3.5 Normalized, start-aligned waveforms of all signals collected from the Artemis I launch. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.




Figure 3.6 Continuous wavelet transform (CWI) power of start-aligned signals collected from the Artemis I launch at 24.4 km (a), 37.6 km (b), and 56.9 km (c) from the launch pad.

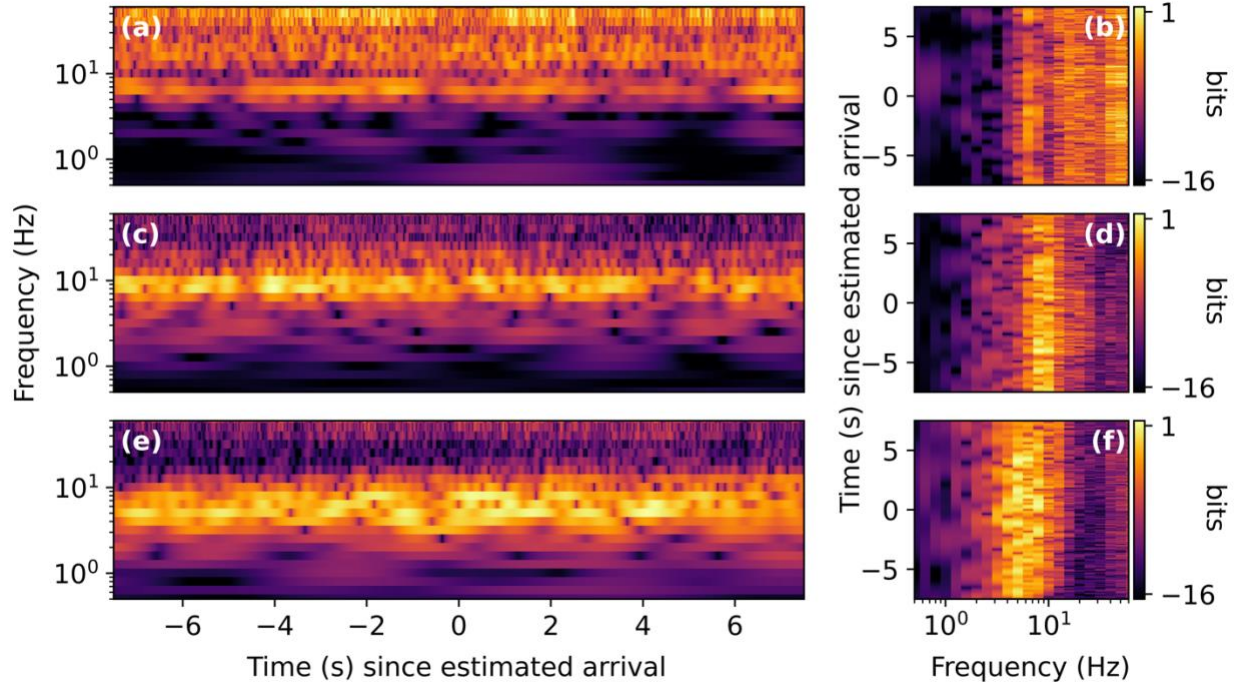


Figure 3.7 CWT power of peak-aligned signals collected from the Artemis I launch at 24.4 km (a, b), 37.6 km (c, d), and 56.9 km (e, f). All panels show the CWT power of the normalized signals, with the left-hand panels showing frequency on the y-axis and time on the x-axis, and the right-hand panels showing time on the y-axis and frequency on the x-axis.

3.3.2 Falcon Heavy

The Falcon Heavy is a super-heavy-lift launch vehicle manufactured by SpaceX. The first stage of its launch is powered by twenty-seven liquid-fueled Merlin 1D engines. It is represented in ASTRA by 40 signals originating from eight individual launches. The five waveforms collected from flight FH-003 are shown in Figure 3.8. Compared to signals collected at short range from Artemis I, the FH-003 signals at the 23.0 and 24.5 km ranges appear to have significantly less impulsive onsets and longer durations ($\sim 100\text{--}200$ s), and the leading small transients (i.e., short-duration changes in amplitude and/or frequency) seen in the Artemis I data are not present at the closest station. The peak separations seen in the Artemis I data as range increases are also observed here; however, while this

pattern was apparent only past the 30 km range in the Artemis I data, it is observed in all but the closest FH-003 signal.




Figure 3.8 Normalized, start-aligned waveforms of all signals collected from launch FH-003 of SpaceX's Falcon Heavy. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.

We once again select one signal from each range class to illustrate identified time-frequency characteristics of Falcon Heavy signals. In Figure 3.9, we see that the increase in power at the start-aligned arrival time estimate is gradual when compared with data from Artemis I, and the high-power bands near 45 and 100 Hz observed in the Artemis I data are not present. The short-range signals still contain more higher-frequency content than the longer-range signals, however, with most of the power over 10 Hz being lost to attenuation as the signal propagates. We once again note the Doppler shift observed at all stations, but it is most clearly illustrated by comparing panels (d) and (f) in Figure 3.10. The leading transient seen in the longer-range stations has a characteristic frequency near 6 Hz.




Figure 3.9 CWT power of start-aligned signals collected from launch FH-003 at 23.0 km (a), 37.1 km (b), and 52.2 km (c) from the launch pad.




Figure 3.10 CWT power of peak-aligned signals collected from launch FH-003 at 23.0 km (a, b), 37.1 km (c, d), and 52.2 km (e, f). All panels show the CWT power of the normalized signals, with the left-hand panels showing frequency on the y-axis and time on the x-axis, and the right-hand panels showing time on the y-axis and frequency on the x-axis.

3.3.3 Delta IV Heavy

The Delta IV Heavy is a retired heavy-lift launch vehicle manufactured by United Launch Alliance. Its first stage was powered by three RS-68 liquid-fueled engines. All three engines would operate at full thrust for the first 44 s of the launch, after which the central engine would reduce thrust to 55% until the two booster engines separated 242 s after launch, at which point the central engine would return to operating at full thrust. Only the final Delta IV Heavy launch (D-389) is represented in ASTRA, from which one incomplete and three complete signals were collected, the waveforms of which are shown in Figure 3.11.




Figure 3.11 Normalized, start-aligned waveforms of all signals collected from Delta IV Heavy launch D-389. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.

Without any signals in the 50–70 km range class, it is more difficult to assess the effect of propagation on the characteristics of the launch signals. However, the broad trends observed in other rocket types are seen, with the signal onset becoming more emergent and experiencing disproportionate loss of high-frequency content as propagation distance increases, and Doppler shifting of characteristic frequencies as the rocket ascends. In Figure 3.12a, we see that the frequency distribution of the closest signal is similar to that of the short-range Artemis I data, with power concentrated in bands near 10, 20, and 45 Hz. Like the Falcon Heavy signals, however, there is no leading transient. There is also significantly less power near the start-aligned estimate of arrival time than is seen in the Artemis I or Falcon Heavy signals.




Figure 3.12 CWT power of start-aligned signals collected from launch D-389 at 29.1 km (a) and 43.4 km (b) from the launch pad.

3.3.4 Vulcan Centaur

ULA’s Vulcan Centaur is a heavy-lift launch vehicle meant to replace both the Delta IV Heavy and the Atlas V. The first stage of its launch is powered by two liquid-fueled BE-4 engines and zero to six GEM 63XL solid rocket boosters (SRBs). ASTRA contains only four Vulcan Centaur signals, all of which originate from a single launch: Cert-1 V-001, which used two SRBs. The four waveforms collected from the launch are shown in Figure 3.13. The two signals collected at a less than 30 km range are spindle-shaped, with short-duration high-amplitude segments (~ 50 s) and no leading transient, as can be seen in Figure 3.14. The propagation effects and Doppler shifts are similar to those observed for other rockets. However, like the Delta IV Heavy, the Vulcan Centaur signals collected at ranges greater than 30 km show significantly less power near the start-aligned estimate of arrival time than those from Artemis I or Falcon Heavy launches.




Figure 3.13 Normalized, start-aligned waveforms of all signals collected from Vulcan Centaur launch Cert-1 V-001. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.




Figure 3.14 CWT power of start-aligned signals collected from Vulcan Centaur launch Cert-1 V-001 at 26.6 km (a), 34.2 km (b), and 49.1 km (c) from the launch pad.

3.3.5 Falcon 9

SpaceX’s Falcon 9 is a medium-lift launch vehicle with a first stage powered by nine Merlin 1D+ liquid-fuel engines. The Falcon 9 is the most represented rocket type in ASTRA by a wide margin, with 949 signals originating from 211 launches. Due to the sheer number of Falcon 9 launches in the dataset, there are numerous outliers. In general, however, the waveforms from most Falcon 9 launches are similar to those shown in Figure 3.15, which are from launch SL-G607. At a less than 30 km range, the waveforms are spindle-shaped, show high chronological symmetry about the peak amplitude, and have durations of ~ 60 s. There is no smaller transient leading the high-amplitude, spindle-shaped signal. As the propagation distance increases, the signal elongates and the onset becomes more emergent, but the strong peak separation observed in signals from the super-heavy lift class is not observed.




Figure 3.15 Normalized, start-aligned waveforms of all signals collected from SpaceX Falcon 9 launch SL-G607. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.




Figure 3.16 CWT power of start-aligned signals collected from launch SL-G607 at 20.0 km (a), 32.2 km (b), and 51.5 km (c) from the launch pad.

3.3.6 Atlas V

ULA's Atlas V is a medium-lift launch vehicle with a first stage powered by one NPO Energomash RD-180 liquid-fuel engine and up to five SRBs. It is the second most represented rocket type in the dataset, with 77 signals originating from 19 launches. Due in part to the variable number of SRBs, the characteristics of Atlas V signals can vary significantly from launch to launch. For configurations with four to five SRBs, the waveforms are similar to those shown in Figure 3.17, which are from a launch using five SRBs. At a less than 30 km range, the waveforms are characterized by impulsive onsets similar to those seen in the Artemis I data, followed shortly thereafter by a more emergent onset, longer duration signal. The expected propagation effects are observed, as is a less extreme version of the peak separation seen in the Artemis I data.




Figure 3.17 Normalized, start aligned waveforms of all signals collected from Atlas V launch AV-093. Labels on the lefthand y-axis indicate the estimated propagation distance of each signal.

In the time–frequency domain, transient signals near the aligned-start estimate of arrival time are present at all ranges, with a characteristic frequency near 6 Hz, while the power of the main launch signal is concentrated near 10, 20, and 45 Hz, with the 10 Hz band dominating at ranges greater than 30 km, as seen in Figure 3.18.




Figure 3.18 CWT power of start-aligned signals collected from launch AV-093 at 20.5 km (a), 34.1 km (b), and 51.9 km (c) from the launch pad.

Alternatively, Atlas V configurations without SRBs show fewer similarities with Artemis I data and more with data from Falcon 9 launches. Figures 3.19 and 3.20 show data from Atlas V launch AV-104, which used a configuration without SRBs. The 6 Hz transient leading the main launch signal in the previous example does not appear.




Figure 3.19 Normalized, start-aligned waveforms of all signals collected from Atlas V launch AV-104. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.




Figure 3.20 CWT power of start-aligned signals collected from launch AV-104 at 21.6 km (a), 34.2 km (b), and 53.9 km (c) from the launch pad.

3.3.7 Small-lift launch vehicles

Two small-lift launch vehicles are represented in the dataset: Relativity Space’s Terran 1 and Astra Space’s Rocket 3.3. Both are propelled in the first stage by a single liquid-fuel engine. Only one launch of each is included in ASTRA, and from those launches, few signals have high enough signal-to-noise ratios to visually identify characteristics, likely due to the relatively low energy of these launches compared to launches of heavier-lift rockets. Preliminary analysis indicated, however, that the time–frequency characteristics of the two rockets are similar. The waveforms of the two small-lift signals with the highest signal-to-noise ratios (both from Terran 1) are plotted in Figure 3.21 and their time–frequency representations are shown in Figure 3.22.




Figure 3.21 Normalized, start-aligned waveforms of signals collected from Terran 1. Labels on the left-hand y-axis indicate the estimated propagation distance of each signal.




Figure 3.22 CWT power of start-aligned signals collected from Terran-1 at 17.1 km (a) and 26.9 km (b) from the launch pad.

3.4 Discussion

Analysis of the acoustic rocket launch dataset presented in this paper shows both similarities and differences with other studies on rocket launch acoustic signatures. In this dataset, the lowest characteristic frequency of the main launch sequence observed is at 10 Hz, while previous studies (Balachandran & Donn, 1971; Blom et al., 2016; Donn et al., 1968; Evers et al., 2018; Kaschak et al., 1970; Pilger, Hupe, et al., 2021; Tenney et al., 2003) have found the characteristic energy of rocket launch signals to be significantly lower. This is unsurprising due to the diminished frequency response of smartphone microphones compared to traditional infrasound microphones, which makes detecting frequencies much lower than 5 Hz with on-board smartphone microphones challenging at low amplitudes. Nevertheless, the 10 Hz peak is sufficiently low-frequency to remain detectable over the full range of this dataset, and studies on far-field data (Balachandran & Donn, 1971; Blom et al., 2016; Donn et al., 1968; Evers et al., 2018; Kaschak et al., 1970; Pilger, Hupe, et al., 2021; Tenney et al., 2003) can be referenced for insight into the very-low-frequency features of acoustic rocket signatures.

In addition to the detected characteristic main launch frequency, interesting features were also identified earlier in the launch sequence. Acoustic data from some launches showed a transient signal at approximately 6 Hz leading the 10 Hz signal. Previous work found in conference proceedings also noted the existence of this transient and concluded that it was an indication of solid fuel ignition (Olson, 2012). Preliminary analysis of the data in this work seems to lend some support to that conclusion, as the 6 Hz transient appeared for 9 out of the 19 launches using solid fuels (47%) and only 17.3% of launches using only liquid fuel. The lack of an observed transient in the other 10 solid fuel launches may be due to any number of environmental conditions on the specific day and does not necessarily contradict the solid fuel ignition hypothesis. However, since the 6 Hz transient appears in data from a number of launches using only liquid fuel (such as the Falcon Heavy launch shown in

Figures 3.8–3.10, as well as more than two dozen Falcon 9 launches), we believe the origin of the 6 Hz transient signal is either more complex or it has similar characteristics to another signal with a different source mechanism. The exact origin of the 6 Hz transient, while interesting, is beyond the scope of this paper and will be left for future studies.

ASTRA was curated and released to the public with the intention of facilitating our own research as well as that of professionals, students, and citizen scientists alike who may be interested in replicating and improving upon the presented results. Preliminary research using data included in ASTRA to train machine learning models to detect acoustic rocket launch signatures in the field in near-real-time has shown promising results and further research into this application is ongoing. The dataset can be accessed and downloaded as a pandas DataFrame (McKinney, 2010), with each recording having its own index and named columns containing the following data: raw acoustic timeseries; epoch time of the first sample; sampling frequency in Hz; smartphone make, model number, identification number, and location; launch identification number, reported epoch time in seconds, and location; estimated source-to-receiver propagation distance in kilometers; time-of-arrival estimates; flight number of the launch; and rocket make, type, model number, and number of SRBs used. To replicate the arrival time estimates, the ERA5 2 m temperature and 10 m u- and v-components of wind data (Hersbach et al., 2018) are required. The open-source software packages (redvox-pandas, quantum-inferno, and pandas) used for data curation and analysis in this work are available online.

3.5 Acknowledgements

The authors are grateful for the support of the U.S. Department of Energy, National Nuclear Security Administration, Office of Defense Nuclear Nonproliferation, Research and Development. They would also like to thank the citizen scientists who operated the network during the collection

period, Samuel Kei Takazawa for his advice and support, and all those who supplied feedback on this project at MTV and ETI conferences. The temperature and wind data used in this study were downloaded from the Copernicus Climate Change Service (2023) (Hersbach et al., 2018). The results contain modified Copernicus Climate Change Service information from 2020. Neither the European Commission nor ECMWF is responsible for any use that may be made of the Copernicus information or data it contains.

3.6 Data Availability Statement

ASTRA is available as a pandas DataFrame and can be found in the Harvard Dataverse open-access repository under the Digital Object Identifier doi: 10.7910/DVN/ZKIS2K. The ERA5 temperature and wind data used in this study are available through the Copernicus Climate Change Service.

Chapter 4: Rocket Launch Detection with Smartphone Audio and Transfer Learning

In preparation to submit to Signals.

Abstract:

Rocket launches generate infrasound signatures that have been detected at great distances. Due to the sparsity of the networks that have made these detections, however, most signals are detected tens of minutes to hours after the rocket launch. In this work, a method of near-real-time detection of rocket launches using data from a network of smartphones located 10-70 km from launch sites is presented. A machine learning model is trained and tested on the open-access Aggregated Smartphone Timeseries of Rocket-generated Acoustics (ASTRA), Smartphone High-explosive Audio Recordings Dataset (SHAReD), and ESC-50 datasets, resulting in a final accuracy of 97% and false positive rate of <1%. The performance and behavior of the model are summarized, and its suitability for persistent monitoring applications is discussed.

4.1 Introduction

Since the launch of Sputnik on October 4, 1957, humanity has successfully launched over six thousand rockets into orbit. As science and technology have advanced in the decades since, the specifications and capabilities of rocket engines have changed and grown while the underlying principles have remained constant. To escape Earth's atmosphere, rockets must be accelerated to great velocities. To achieve this acceleration, propellants are burned in a combustion chamber. The resulting exhaust is the working fluid; it is accelerated through a propelling nozzle and expelled at hypersonic velocities to produce thrust. The nature and state of the propellants used can vary between makes and models of rockets, but these basic principles remain constant.

During the launch sequence of a rocket, acoustic waves are generated by a number of source mechanisms, resulting in a complex acoustic signature. For convenience, we will divide these acoustic waves into three categories: waves generated by the engine (engine noise), waves generated by exhaust (exhaust noise), and waves generated by turbulent flow excitation (jet noise). Engine noise, while the least intense of the three, is well-studied due to combustion instabilities leading to catastrophic engine failure, material fatigue, etc. generating distinct acoustic signatures. Through ignition and liftoff, exhaust noise is the primary source of acoustic energy. As exhaust is expelled from the rocket, it collides with ambient air, creating shock waves. For many rockets, the intensity of these waves is high enough to damage the rocket and/or nearby structures, and thus water-based suppression of exhaust noise has been frequently employed on launch pads since the Space Shuttle program. As the rocket accelerates into supersonic velocities, jet noise generated by the increasing shear flow and resulting turbulent eddies overtakes exhaust noise and becomes the dominant source of acoustic energy.

Like other large-scale events such as earthquakes and explosions, rocket launches generate low-frequency (<300 Hz) sound as well as infrasound (<20 Hz) (Schwardt et al., 2022). These

components of acoustic rocket signatures can remain detectable over great distances due to the frequency dependence of atmospheric attenuation. Infrasound sensors are employed by the International Monitoring System (IMS) of the Comprehensive Nuclear-Test-Ban Treaty Organization to detect and monitor large-scale events, including rocket launches (Hupe et al., 2022). IMS infrasound stations have successfully collected signatures from a variety of events, including numerous rocket launches (Pilger, Hupe, et al., 2021), but the sparsity of the network necessitated by its global coverage and the expense of traditional infrasound sensors results in these signatures being collected only after propagating vast distances. The result is that detections can only be made long after a rocket has launched, at which point much of the information the signature once carried about its source has been lost to attenuation.

Rapid detection of rocket launches may be possible, however, through the use of non-traditional sensors such as smartphones. Smartphone microphones have been used to successfully collect infrasonic and low-frequency signatures (Popenhagen et al., 2023; Takazawa, Popenhagen, Ocampo Giraldo, Cárdenas, et al., 2024), and smartphone audio data has been used to train machine learning models to accurately detect explosion signatures (Takazawa, Popenhagen, Ocampo Giraldo, Hix, et al., 2024; Thandu et al., 2015, 2017) in previous studies. In this work, we present a machine learning model designed and trained to detect acoustic rocket signatures in smartphone audio data from three open-access datasets using the machine learning method transfer learning. In the following sections, we will detail the methods used, evaluate performance under different conditions, and discuss the suitability of the solution for persistent monitoring.

4.1.1 Transfer Learning with YAMNet

Transfer learning is a machine learning method first developed by Stevo Bozinovski and Ante Fulgosi at the University of Zagreb in 1972 (Bozinovski, 2020; Bozinovski & Fulgosi, 1976). In theory,

the technique improves learning efficiency by reusing knowledge gained from one task for a second, related task. The popularity of transfer learning has risen in recent years (Ashurov et al., 2022; Brusa et al., 2021; Hyun, 2023; Takazawa, Popenhagen, Ocampo Giraldo, Hix, et al., 2024; Tsalera et al., 2021) due in part to the existence of publicly available models pre-trained on very large datasets for general tasks which can be used to boost performance on a related, specific task if the volume of available training data for the specific task is limited (Pan & Yang, 2010).

One such publicly available model is Google’s Yet Another Mobile Network (YAMNet) (Plakal & Ellis, 2029), a deep neural network using Mobilenet_v1, a depthwise-separable convolution architecture (Howard et al., 2017). YAMNet is pre-trained on AudioSet, a human-annotated collection of more than 2 million 10-second audio clips pulled from YouTube videos (Gemmeke et al., 2017), to predict 521 classes of audio events. The intended input of the model is audio data sampled at 16 kHz. A stabilized log Mel spectrogram is computed from the input with a frequency range of 125 Hz to 7500 Hz, after which the results are divided into 0.96 s segments with 50% overlap. These segments are fed into the Mobilenet_v1 architecture. The output of Mobilenet_v1 is then averaged-pooled, resulting in 1024 embeddings. Up to this point, the procedures we have described are those performed by YAMNet for all applications. For the general case of using YAMNet for audio classification, a final output layer would then take in the embeddings and output the class scores, from which predictions would be made. For our chosen method of transfer learning, however, the final layer must be removed, and the embeddings used as the input of our smaller, more specific rocket detection model.

4.2 Data and Methods

4.2.1 Aggregated Smartphone Timeseries of Rocket-generated Acoustics (ASTRA)

The acoustic rocket signatures used to train and test the models in this work are from Aggregated Smartphone Timeseries of Rocket-generated Acoustics (ASTRA), a publicly available

dataset containing 1089 smartphone audio recordings of 243 rocket launches (Popenhagen, 2024). The data in ASTRA were recorded at 800 Hz sampling rate by Android smartphones stationed 10-70 km from launch pads at Cape Canaveral Space Force Station and the Kennedy Space Center. Along with the recordings themselves, the estimated signal start and peak times at each station, as well as the unique launch identification strings were used. The process used to calculate these estimates is detailed in Popenhagen & Garcés (2025).

In preparation for machine learning, a subset of ASTRA was compiled by removing any recordings for which confidence in the alignment was reduced due to missing data or high environmental noise. The remaining 789 recordings from 233 launches were upsampled to 16 kHz, and 4.8 s windows centered on the estimated launch signal peak were selected. The selected 4.8 s windows were then divided into 0.96 s segments, resulting in 5 samples. The unique launch identification strings associated with each recording in ASTRA were used to ensure that recordings from the same launch were not split between the training, validation, and test sets. If a recording contained data taken from >120 s before the associated estimate of the arrival time of the beginning of the signal, that segment was separated and divided into up to 50 samples for additional noise data. These samples were processed similarly to the rocket signature samples, being divided into 0.96 s segments and split by associated event.

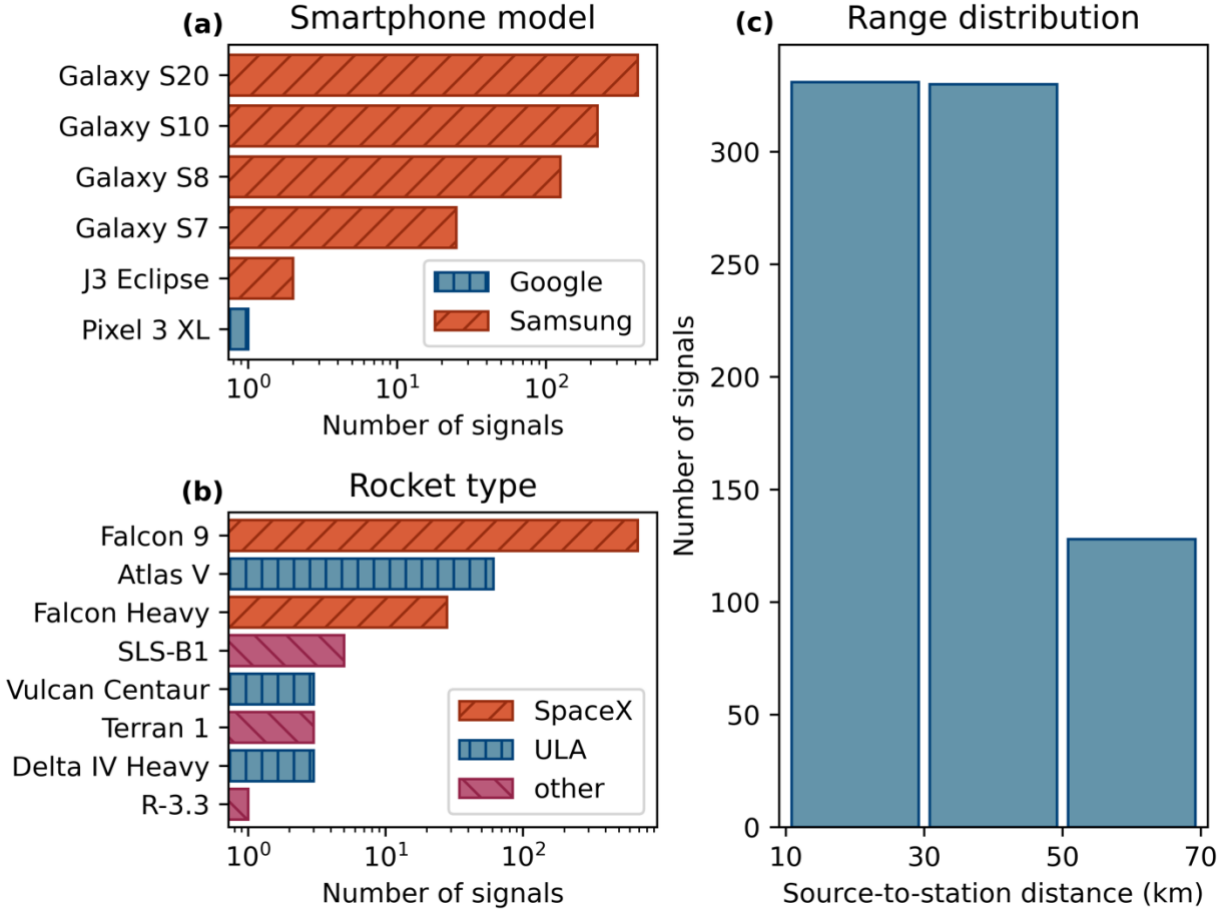


Figure 4.1 Bar plots of the distribution of signals in the high-confidence subset of ASTRA (a) recorded on different makes and models of smartphones, (b) originating from different types of rockets, and (c) collected in different range categories.

4.2.2 Smartphone High-explosive Audio Recordings Dataset (SHAReD)

The Smartphone High-Explosive Audio Recordings Dataset (SHAReD) (Takazawa, 2024) is an open-access dataset containing 326 multi-modal smartphone data collected during 70 surface high explosive events. In this work, only the smartphone microphone data were used. The explosion signals in SHAReD all have signal-to-noise ratios of >3 , effective yields from 1 to 100 kg TNT, and ranges between 10g and 4 tons. SHAReD audio were collected at sampling rates of either 800 Hz or 8000 Hz. For an in-depth overview of SHAReD, we direct the reader to its accompanying papers (Takazawa, Popenhagen, Ocampo Giraldo, Cárdenas, et al., 2024; Takazawa, Popenhagen, Ocampo Giraldo, Hix,

et al., 2024). For each station during each event, SHAReD includes one 0.96 s long ‘explosion’ waveform and one 0.96 s long ‘silence’ waveform (sampled from before the explosion signal’s arrival). Like the ASTRA data, the SHAReD data were split by event so that the model was never trained and tested on data from the same event.

4.2.3 Environmental sound data from ESC-50

The ESC-50 dataset is a collection of 2000 environmental sound recordings originally recorded as part of the Freesound project (Piczak, 2015). Each of the 2000 clips is 5 s in duration, has a sampling rate of 44.1 kHz, and is labeled as belonging to 1 of 50 semantical classes. There are exactly 20 clips from each class (ex. ‘dog’, ‘snoring’, ‘chainsaw’, etc.), with some classes (‘thunderstorm’, ‘fireworks’, ‘airplane’, etc.) including low-frequency and/or infrasonic sources. ESC-50 is widely used as a benchmarking dataset for environmental sound classification models, and its inclusion has been shown to improve robustness and performance of transfer learning audio classification models (Takazawa, Popenhagen, Ocampo Giraldo, Hix, et al., 2024). 5 0.96 s long samples were taken from each ESC-50 clip. These samples were kept together when splitting the dataset into training, validation, and testing sets.

4.2.4 Preprocessing and splitting

All audio samples recorded at >800 Hz sampling rate were downsampled to 800 Hz to match the ASTRA audio samples. For each recording in ASTRA, the mean signal-to-noise ratio (SNR) of the ‘rocket’ and ‘noise’ samples was calculated, and those with ratios of <2.0 were removed to prevent training on misaligned or exceptionally noisy samples. All data except the ‘rocket’ data from ASTRA were labeled as ‘noise’. For the proposed use-case of persistent monitoring, it is beneficial to prioritize a low false positive rate over other metrics, and thus the training data was left intentionally unbalanced, including approximately 10 ‘noise’ samples for every ‘rocket’ sample in the training set. This imbalance

was intended to increase robustness by exposing the model to a greater amount and variety of ‘noise’ data, as well as encouraging the model to favor false negatives over false positives.

The data from ASTRA, SHAReD, and ESC-50 were each split randomly into training, validation, and testing sets, with target percentages of 80%, 10%, and 10%, respectively. This random splitting was performed 15 times, in preparation for training and testing the model 15 separate times on each of the 15 randomly split datasets. The target distributions were not always met exactly, due to the constraint of isolating the data from each individual launch, explosion, or Freesound clip to one of the three sets.

4.2.5 Model design

Before building, training, and testing the transfer learning model, the full dataset was run through YAMNet alone, and the resulting predictions were analyzed. The distributions of the classes predicted by YAMNet for data from ASTRA, SHAReD, and ESC-50 are shown in Figure 4.2. Due to the data having been upsampled from 800 Hz and thus being devoid of any content above 400 Hz, the highly inaccurate predictions we observe are unsurprising, even for the ESC-50 data, which contains types of sounds similar to those YAMNet was trained to classify. What is interesting, however, is that the distributions of predicted classes vary significantly between subsets of the dataset, indicating that YAMNet is able to see some differences between these subsets despite having no context for data collected at 800 Hz.

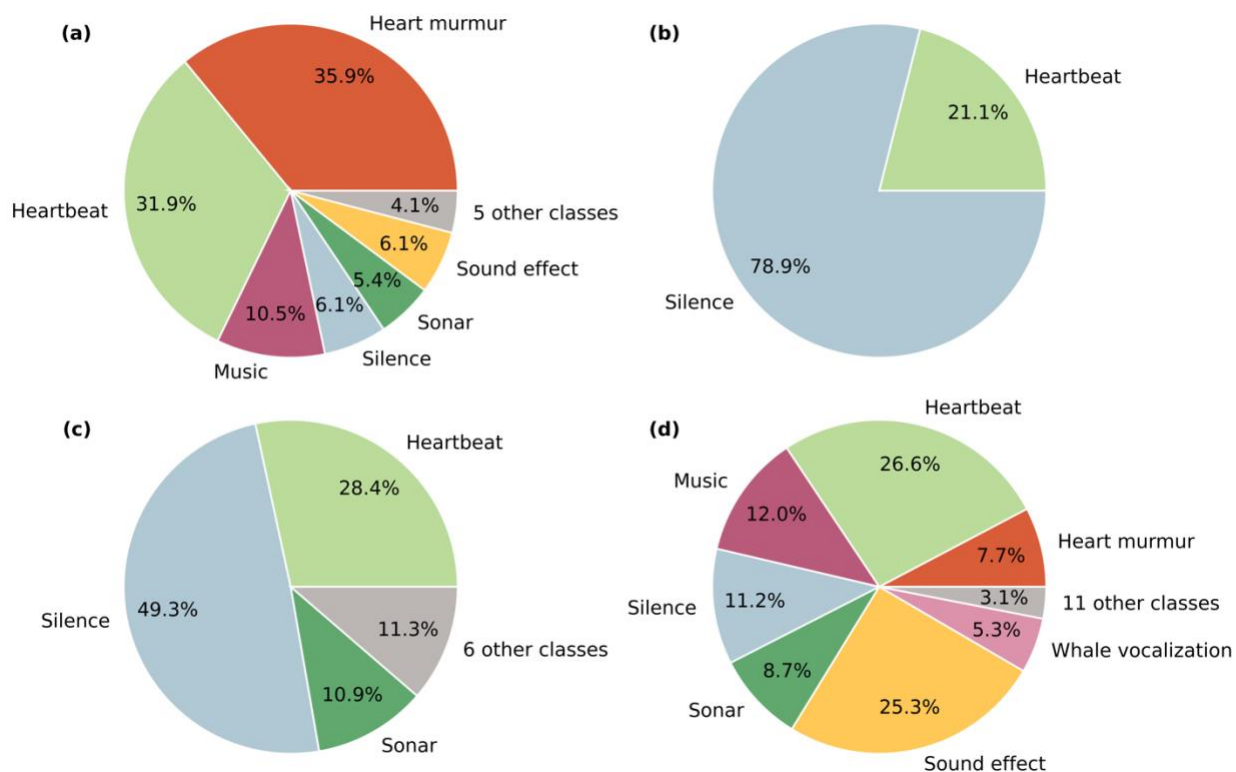


Figure 4.2 Pie charts showing the distributions of YAMNet’s predicted classes for (a) rocket launch audio from ASTRA, (b) ‘explosion’ (50%) and ‘silence’ (50%) audio from SHAReD, (c) pre-event audio from ASTRA, and (d) environmental audio recordings from the ESC-50 dataset. All data were upsampled to 16 kHz from 800 Hz sampling rate.

Transfer learning was chosen to compensate as much as possible for the limited quantity of rocket audio data. In theory, we can transfer some of the knowledge YAMNet has on the general problem of audio classification to our more specific problem of rocket detection. To construct the transfer learning model, the final output layer of YAMNet was removed and replaced with three new layers. The first of these was a fully connected layer with 32 nodes, utilizing leaky rectified linear unit (leaky ReLU) activation with an alpha value of 0.01. The second was a dropout layer added to minimize overfitting, and the final layer was an output layer with 2 nodes corresponding to the two classes (‘rocket’ and ‘noise’). Sparse categorical cross-entropy was used for loss, and the Adamax optimizer was used. The number of nodes, activation function, alpha value, and optimizer were selected after

analyzing the results of cross-validation, with lower numbers of nodes being given preference to avoid overfitting. The number of epochs was set to 300, with early stopping implemented to further reduce the likelihood of overfitting.

4.3 Results

The model was evaluated using the samples assigned to the training set for each run, consisting of roughly 10% of the data from each of the three open-access datasets. All testing samples were taken from events (rocket launches, explosions, or Freesound clips) that were not included in the training or validation sets. The overall performance of the model is presented, along with the performance of the best-performing split. In Figures 4.3 and 4.4, results are shown using normalized confusion matrices, which show the true positive rate in the upper left corner and the true negative rate in the lower right corner. Since the categories are unbalanced, the confusion matrices are all normalized for clarity, but the original counts were preserved and are included in parentheses on the figures.

4.3.1 Overall performance

The accuracy of the model was calculated according to

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{FP} + \text{FN} + \text{TN}}, \quad (4.1)$$

and ranged between 93.53% and 97.00%, with the mean over all 15 splits lying at 95.40% accuracy. The mean performance of the model is shown as a confusion matrix in Figure 4.3. In Figure 4.4, a confusion matrix of the best-performing split is shown. When trained on this split, the model performed slightly better, with no indications of overfitting. The best split accuracy was 97.00%, with a false positive rate of 0.98%.




Figure 4.3 The confusion matrix of the model on all 15 test sets. Each quadrant shows the relevant rate as a percentage, as well as the total number of samples in the category in parentheses.




Figure 4.4 The confusion matrix of the best performing split. Each quadrant shows the relevant rate as a percentage, as well as the total number of samples in the category in parentheses.

4.3.2 Misclassification analysis

To better understand the model’s behavior, it is important to remember that 5 positive samples were taken from each rocket recording, and that there are multiple recordings for most launches. In Table 4.1, the ASTRA events with false negatives in the best split model are listed, along with the other classifications from the event, which show that for every event with false negatives, there are always more true positives. For every recording in the best split model’s test set, there is at least one ‘rocket’ classification within 10 seconds of the estimated peak arrival time included in ASTRA. From

an application perspective, then, there are no false negatives, only delayed true positives, and we can focus on reducing the false positive rate.

Table 4.1 All events with ‘rocket’ samples misclassified as ‘noise’ by the best split model.

Event	No. false negatives	No. positive samples	No. true positives
ASTRA_218	7	30	21
ASTRA_219	3	15	12
ASTRA_220	1	25	24
ASTRA_221	1	10	9
ASTRA_222	10	25	15
ASTRA_224	5	20	15
ASTRA_225	4	25	21
ASTRA_226	2	20	18
ASTRA_227	3	10	7
ASTRA_229	5	15	10
ASTRA_230	1	20	19
ASTRA_232	4	20	16

The simplest method of reducing the false positive rate is to increase the threshold for positive ‘rocket’ classifications slightly, from 0.5 to 0.6. With the increased threshold, the best split model’s false positive rate decreases from 0.98% to 0.76%. This improvement in the false positive rate is accompanied by an increase of the false negative rate from 15.59% to 19.66%. However, there is still at least one ‘rocket’ classification within 10 seconds of the estimated peak arrival time of each recording, thus the practical effect of the increased false negative rate is minimal. The best split model’s remaining false positives after increasing the threshold are shown in Table 4.2. Upon examination, it

is clear that false positives tend to appear chronologically isolated or in small clusters of 5 or less samples.

Table 4.2 All events with ‘noise’ samples misclassified as ‘rocket’ by the best split model.

Event	No. false positives	No. negative samples	No. true negatives
ASTRA_225	5	50	45
ESC50_1660	3	5	2
ESC50_1667	1	5	4
ESC50_1675	1	5	4
ESC50_1683	1	5	4
ESC50_1731	2	5	3
ESC50_1761	1	5	4
ESC50_1779	2	5	3
ESC50_1808	2	5	3

To visualize this behavior, we can plot the ‘rocket’ classifications as vertical green lines with a transparency inversely proportional to the model’s confidence in its classification. In Figure 4.5, all recordings from one of the launches in the testing set are plotted this way, showing a number of false positives before the arrival of the signal at two of the three stations. Since the rocket launch signal persists over at least 100 s, we can eliminate chronologically isolated positives by setting all scores less than the 0.6 threshold to 0 and applying a 4.8 s long rolling median filter. In Figure 4.6, the remaining ‘rocket’ classifications after applying the median filter. After applying this method to the scores of pre-launch segments of the ASTRA recordings included in the test set, the best split model’s false positive rate on pre-launch ASTRA data decreased from 0.76% to 0.11%. This improvement in the false positive rate has minor tradeoffs in two areas. First, the detection time is delayed from the first ‘rocket’

classification over 0.5 by a mean of 4.92 s. Second, despite the increase in the false negative rate appearing negligible on the sample level, 1 of the 100 recordings in the test set now has no ‘rocket’ classifications within 10 seconds of its estimated peak arrival time.




Figure 4.5 Plot showing the best split model’s unfiltered performance on all signals from the SpaceX Falcon 9 launch, Transporter-4. Green lines indicate ‘rocket’ classification scores >0.6.




Figure 4.6 Plot showing the best split model’s median-filtered performance on all signals from the SpaceX Falcon 9 launch, Transporter-4. Green lines indicate ‘rocket’ classification scores >0.6 .

When applied to the ESC-50 data, this method also decreases the false positive rate significantly. Figure 4.7 shows the false positive rates before and after thresholding and median filtering for each ESC-50 class containing samples the best split model misclassified, as well as the mean false positive rates over all ESC-50 classes. As can be seen in Figure 4.7, the overall false positive rate on ESC-50 data dropped to 0.05% after applying thresholding and median filtering. As the recordings from SHAReD were each only 0.96 s long, we were unable to determine the effect of median filtering on those data. However, the model performed exceptionally well on SHAReD data even before adjusting the threshold, with no misclassifications by the best split model and false positive rates $<1\%$ for all 15 models.




Figure 4.7 A horizontal bar plot representing the best split model's false positive classifications of ESC-50 data, plotted according to class labels included with ESC-50. Light colored bars represent the false positive rates before thresholding and median filtering, and dark colored bars represent the same rates after thresholding and median filtering are applied. Results for individual classes are shown in blue and labeled with black text, while the overall results are shown in orange and labeled with white text.

4.4 Discussion and Conclusion

After training and testing on 15 random train-test splits, we found that the transfer learning model's behavior is exceptionally stable, with only slight variations in the results between different splits. Analyzing the behavior of the best split model in detail, it was discovered that despite the nominal false negative rate of 15.59%, the model made at least one 'rocket' classification within 10 seconds of the estimated arrival times of 99/100 recordings in the training set. In addition, the already low false positive rate of 0.98% can be reduced even further by increasing the positive classification threshold from 0.5 to 0.6 and applying a 4.8 s long median filter. The application of these two measures resulted in the best split model's false positive rates on ESC-50 data decreasing from 1.43% to 0.05%, pre-launch 'noise' data from ASTRA decreasing from 0.76% to 0.11%, and the overall false positive rate decreasing from 0.98% to 0.072%. The increased threshold and addition of the median filter combined were found to increase the time between launch and detection by 4.92 s on average, but this delay is relatively short compared to the estimated propagation times (~60-200 s) and its negative effects would likely be outweighed by the benefits of the decreased false positive rate in most use cases. Applying the median filter also reduced the true positive rate slightly, resulting in the model failing to detect the launch within 10 seconds of the estimated peak arrival time of 1 of the 100 recordings in the dataset. It is worth noting, however, that the recording in question was 1 of 4 recordings of the same launch, and the model was able to detect the launch successfully in the other 3 recordings.

Although the initial success of this transfer learning model is encouraging, the dataset is too small to truly determine its usefulness for persistent monitoring. As more and more rockets continue to be launched every year, the amount of available data will likely increase rapidly in the next few years. Future work should include long-term deployment and testing of the model to fully evaluate its

potential for monitoring. If further decreases in the false positive rate are required in the future, there are a number of different avenues to pursue, including training against additional sources of infrasound such as earthquakes or volcanic activity, incorporating ensemble learning (Takazawa, Popenhagen, Ocampo Giraldo, Hix, et al., 2024), and combining input from multiple smartphones, which may also prove useful for trajectory modeling in the future.

4.5 Data Availability

ASTRA is available as a pandas DataFrame (McKinney, 2010; The pandas development team, 2020) and can be found in the Harvard Dataverse open-access repository under the Digital Object Identifier doi: 10.7910/DVN/ZKIS2K. The ERA5 temperature and wind data used in this study are available through the Copernicus Climate Change Service [34,35]. SHAReD can also be found and downloaded from the Harvard Dataverse, under the Digital Object Identifier doi: 10.7910/DVN/YOCYO2. The ESC-50 dataset is available online at <https://github.com/karolpiczak/ESC-50>.

Chapter 5: Conclusion

5.1 Summary

In this dissertation, acoustic explosion data collected from the stratosphere and acoustic rocket launch data collected on the surface were presented. Along with the aggregation, curation, and public release of these data, time frequency analysis was performed, and a preliminary machine learning model was developed to detect rocket launch audio signatures in near-real-time.

Chapter 2 provided a comparison of a surface chemical explosion signature recorded by a smartphone microphone and a traditional infrasound microphone from an ascending balloon in the middle stratosphere, the first successful observation of its kind. It was determined that despite the smartphone microphone's diminishing frequency response below 20 Hz frequency, high energy regions around 5 and 10 Hz were clear and recognizable in time-frequency representations of the smartphone audio data, with a signal to noise ratio comparable to that of the traditional infrasound sensor.

Chapter 3 introduced Aggregated Smartphone Timeseries of Rocket-generated Acoustics (ASTRA), a dataset containing 1089 smartphone audio recordings of 243 rocket launches. It included a brief overview of the chronology and time-frequency characteristics of the rocket signatures, identified features of interest for monitoring efforts, and provided estimates of the arrival times of the signals.

Chapter 4 presented a rocket launch detection model using transfer learning. The model was trained on data from ASTRA, as well as data from two other open-access datasets: a smartphone explosion dataset and an environmental sound dataset. A method of thresholding and median filtering the model's classification scores to minimize the false positive rate was detailed, and the model's suitability for persistent monitoring applications was discussed.

As stated in the introduction, the overarching goal of the work detailed in this dissertation was to develop a fast and accurate method of detecting acoustic signatures generated by rocket launches. Over the course of the three content chapters, we have shown that smartphones can be used to collect infrasound signatures with signal to noise ratios comparable with those of traditional infrasound sensors. As there were few publicly available data that could be used to train and test machine learning models for near-real-time detection, we aggregated, annotated, and released to the public a new dataset of 1089 audio recordings of rocket launches. Through analysis of that dataset, we gained insight into the characteristics of rocket-generated acoustic signatures and compared the features present in our data to theory and existing literature. Finally, we applied machine learning methods to our dataset, resulting in a model that successfully detected 99/100 rocket launch recordings in the test set, with a final false positive rate of 0.072% after thresholding and median filtering were applied.

5.2 Future Work

The overarching goal of this dissertation is to increase our understanding of what the characteristics of acoustic rocket launch signatures are, how they distort and attenuate during propagation, and how they can aid in monitoring efforts. Potential avenues of future work using ASTRA include ensemble learning, source localization, ignition-specific detection, trajectory modelling, and rocket classification, as well as experimentation with and improvements upon basic rocket detection as more data become available. Information on the rocket launch each recording comes from was included in ASTRA to facilitate investigation into these areas and others by any who may be interested. All data and models have been made public, and all processing, analysis, and modeling was performed using open-access software.

References

- Ashurov, A., Zhou, Y., Shi, L., Zhao, Y., & Liu, H. (2022). Environmental Sound Classification Based on Transfer-Learning Techniques with Multiple Optimizers. *Electronics*, 11(15), 2279. <https://doi.org/10.3390/electronics11152279>
- Asmar, K., Garces, M., & Williams, B. (2019). A method for estimating the amplitude response of smartphone built-in microphone sensors below 4 kHz. *Journal of the Acoustical Society of America*, 146(1). <https://doi.org/10.1121/1.51110723>
- Balachandran, N. K., & Donn, W. L. (1971). Characteristics of Infrasonic Signals from Rockets. *Geophysical Journal of the Royal Astronomical Society*, 26(1–4), 135–148. <https://doi.org/10.1111/j.1365-246X.1971.tb03387.x>
- Barat, J., Cot, C., & Sidi, C. (1984). On the Measurement of the Turbulence Dissipation Rate from Rising Balloons. *Journal of Atmospheric and Oceanic Technology*, 1(3), 270–275.
- Bellemare, M. G., Candido, S., Castro, P. S., Gong, J., Machado, M. C., Moitra, S., Ponda, S. S., & Wang, Z. (2020). Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature*, 588(7836), 77–82. <https://doi.org/10.1038/s41586-020-2939-8>
- Blom, P. (2014). *infraGA*. GitHub. <https://github.com/LANL-Seismoacoustics/infraGA>
- Blom, P., Marcillo, O., & Arrowsmith, S. (2016). Analysis and modeling of infrasound from a four-stage rocket launch. *The Journal of the Acoustical Society of America*, 139(6), 3134–3138. <https://doi.org/10.1121/1.4953817>
- Blom, P., & Waxler, R. (2017). Modeling and observations of an elevated, moving infrasonic source: Eigenray methods. *The Journal of the Acoustical Society of America*, 141(4), 2681–2692. <https://doi.org/10.1121/1.4980096>
- Bowman, D. C., & Albert, S. A. (2018). Acoustic event location and background noise characterization on a free flying infrasound sensor network in the stratosphere. *Geophysical Journal International*, 213(3), 1524–1535. <https://doi.org/10.1093/gji/ggy069>
- Bowman, D. C., & Krishnamoorthy, S. (2021). Infrasound From a Buried Chemical Explosion Recorded on a Balloon in the Lower Stratosphere. *Geophysical Research Letters*, 48(21). <https://doi.org/10.1029/2021GL094861>
- Bowman, D. C., & Lees, J. M. (2015). Infrasound in the middle stratosphere measured with a free-flying acoustic array. *Geophysical Research Letters*, 42(22), 10010–10017. <https://doi.org/10.1002/2015GL066570>
- Bowman, D. C., Rouse, J. W., Krishnamoorthy, S., & Silber, E. A. (2022). Infrasound direction of arrival determination using a balloon-borne aeroseismometer. *JASA Express Letters*, 2(5), 054001. <https://doi.org/10.1121/10.0010378>
- Bozinovski, S. (2020). Reminder of the first paper on transfer learning in neural networks, 1976. *Informatica (Slovenia)*, 44(3), 291–302. <https://doi.org/10.31449/INF.V44I3.2828>
- Bozinovski, S., & Fulgosi, A. (1976). The influence of pattern similarity and transfer of learning upon training of a base perceptron B2. *Symposium Informatica 3-121-5*.
- Brissaud, Q., Krishnamoorthy, S., Jackson, J. M., Bowman, D. C., Komjathy, A., Cutts, J. A., Zhan, Z., Pauken, M. T., Izraelvitz, J. S., & Walsh, G. J. (2021). The First Detection of an Earthquake From a Balloon Using Its Acoustic Signature. *Geophysical Research Letters*, 48(12). <https://doi.org/10.1029/2021GL093013>
- Brown, R., & Evans, L. (2011). Acoustics and the smartphone. *Acoustics*, 106.

- Brusa, E., Delprete, C., & Di Maggio, L. G. (2021). Deep Transfer Learning for Machine Diagnosis: From Sound and Music Recognition to Bearing Fault Detection. *Applied Sciences*, 11(24), 11663. <https://doi.org/10.3390/app112411663>
- Caswell, T. A., Droettboom, M., Lee, A., de Andrade, E. S., Hoffman, T., Hunter, J., & et al. (2021). *matplotlib/matplotlib: Rel: v3.5.1*. Zenoda.
- Christie, D. R., & Campus, P. (2010). The IMS Infrasound Network: Design and Establishment of Infrasound Stations. In *Infrasound Monitoring for Atmospheric Studies* (pp. 29–75). Springer Netherlands. https://doi.org/10.1007/978-1-4020-9508-5_2
- Copernicus Climate Change Service (2023): ERA5 hourly data on single levels from 1940 to present. (2023). Copernicus Climate Change Service (C3S) Climate Data Store (CDS). <https://cds.climate.copernicus.eu/datasets/reanalysis-era5-single-levels>
- Cotten, D., & Donn, W. L. (1971). Sound from Apollo Rockets in Space. *Science*, 12(3971), 565–567. <https://doi.org/10.1126/science.171.3971.565>
- Cotten, D. E., Donn, W. L., & Oppenheim, A. (1971). On the Generation and Propagation of Shock Waves From Apollo Rockets at Orbital Altitudes. *Geophysical Journal of the Royal Astronomical Society*, 26(1–4), 149–159. <https://doi.org/10.1111/j.1365-246X.1971.tb03388.x>
- Donn, W. L., Posmentier, E., Fehr, U., & Balachandran, N. K. (1968). Infrasound at Long Range from Saturn V, 1967. *Science*, 162(3858), 1116–1120. <https://doi.org/10.1126/science.162.3858.1116>
- Drob, D. P. (2019). Meteorology, climatology, and upper atmospheric composition for infrasound propagation modeling. In A. Le Pichon, E. Blanc, & A. Hauchecorne (Eds.), *Infrasound Monitoring for Atmospheric Studies* (2nd ed., pp. 485–508). Springer Nature. <https://doi.org/10.1007/978-3-319-75140-5>
- Drob, D. P., Picone, J. M., & Garcés, M. (2003). Global morphology of infrasound propagation. *Journal of Geophysical Research: Atmospheres*, 108(D21). <https://doi.org/10.1029/2002JD003307>
- Evers, L. G., Assink, J. D., & Smets, P. S. M. (2018). Infrasound from the 2009 and 2017 DPRK rocket launches. *Geophysical Journal International*, 213(3), 1785–1791. <https://doi.org/10.1093/gji/ggy092>
- Garcés, M. A. (2020). Quantized constant-Q gabor atoms for sparse binary representations of cyber-physical signatures. *Entropy*, 22(9). <https://doi.org/10.3390/E22090936>
- Garcés, M. A. (2023). Quantized Information in Spectral Cyberspace. *Entropy*, 25(3). <https://doi.org/10.3390/e25030419>
- Garcés, M. A., Bowman, D., Zeiler, C., Christe, A., Yoshiyama, T., Williams, B., Colet, M., Takazawa, S., & Popenhagen, S. (2022). Skyfall: Signal Fusion from a Smartphone Falling from the Stratosphere. *Signals*, 3(2), 209–234. <https://doi.org/10.3390/signals3020014>
- Garcés, M. A., Hansen, R. A., & Lindquist, K. G. (1998). Traveltimes for infrasonic waves propagating in a stratified atmosphere. In *Geophys. J. Int* (Vol. 135). <https://academic.oup.com/gji/article/135/1/255/570080>
- Garcés, M. A., & Le Pichon, A. (2009). Infrasound from Earthquakes, Tsunamis and Volcanoes. In R. A. Meyers (Ed.), *Encyclopedia of Complexity and Systems Science* (pp. 4839–4855). Springer.
- Gemmeke, J. F., Ellis, D. P. W., Freedman, D., Jansen, A., Lawrence, W., Moore, R. C., Plakal, M., & Ritter, M. (2017). Audio Set: An ontology and human-labeled dataset for audio events. *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 776–780. <https://doi.org/10.1109/ICASSP.2017.7952261>
- Hersbach, H., Bell, B., Berrisford, P., Biavati, G., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Rozum, I., Schepers, D., Simmons, A., Soci, C., Dee, D., & Thépaut, J.-N. (2018).

- ERA5 hourly data on single levels from 1940 to present. Copernicus Climate Change Service (C3S) Climate Data Store (CDS).
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., & Adam, H. (2017). *MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications*. arXiv.
- Hunter, J. D. (2007). Matplotlib: A 2D Graphics Environment. In *Computing in Science & Engineering* (Vol. 9, Issue 3, pp. 90–95). <https://doi.org/10.1109/MCSE.2007.55>
- Hupe, P., Ceranna, L., Le Pichon, A., Matoza, R. S., & Mialle, P. (2022). International Monitoring System infrasound data products for atmospheric studies and civilian applications. *Earth System Science Data*, 14(9), 4201–4230. <https://doi.org/10.5194/essd-14-4201-2022>
- Hyun, S. H. (2023). Sound-Event Detection of Water-Usage Activities Using Transfer Learning. *Sensors*, 24(1), 22. <https://doi.org/10.3390/s24010022>
- IRIS. (2023). Large surface explosion coupling experiment. In *IRIS*. <https://doi.org/10.7914/1p2e-0167>
- Kaschak, G., Donn, W. L., & Fehr, U. (1970). Long-Range Infrasound from Rockets. *The Journal of the Acoustical Society of America*, 48(1A_Supplement), 12–20. <https://doi.org/10.1121/1.1912102>
- Kong, Q., Martin-Short, R., & Allen, R. M. (2020). Toward global earthquake early warning with the myshake smartphone seismic network, part 1: Simulation platform and detection algorithm. *Seismological Research Letters*, 91(4), 2206–2217. <https://doi.org/10.1785/0220190177>
- Krishnamoorthy, S., Bowman, D. C., Komjathy, A., Pauken, M. T., & Cutts, J. A. (2020). Origin and mitigation of wind noise on balloon-borne infrasound microbarometers. *The Journal of the Acoustical Society of America*, 148(4), 2361–2370.
- Lamb, O. D., Lees, J. M., & Bowman, D. C. (2018). Detecting Lightning Infrasound Using a High-Altitude Balloon. *Geophysical Research Letters*, 45(14), 7176–7183. <https://doi.org/10.1029/2018GL078401>
- Marcillo, O., Johnson, J. B., & Hart, D. (2012). Implementation, Characterization, and Evaluation of an Inexpensive Low-Power Low-Noise Infrasound Sensor Based on a Micromachined Differential Pressure Transducer and a Mechanical Filter. *Journal of Atmospheric and Oceanic Technology*, 29(9), 1275–1284. <https://doi.org/10.1175/JTECH-D-11-00101.1>
- McKinney, W. (2010). Data Structures for Statistical Computing in Python. In S. van der Walt & J. Millman (Eds.), *9th Python in Science Conference* (pp. 56–61). <https://doi.org/10.25080/Majora-92bf1922-00a>
- Mclaughlin, K. L., Brown, D. J., & Australia, G. (2000). Infrasound detection of rocket launches. *22nd Annual DoD/DOE Seismic Research Symposium*. <https://www.researchgate.net/publication/228643229>
- Olson, J. (2012). Infrasound Rocket Signatures. *Adv Maui Opt Space Surveillance Technol Conf 1*, 638–645.
- Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. In *IEEE Transactions on Knowledge and Data Engineering* (Vol. 22, Issue 10, pp. 1345–1359). <https://doi.org/10.1109/TKDE.2009.191>
- Peebles, C., & Air Force History and Museums Program. (1997). *High Frontier: The United States Air Force and the Military Space Program*. U.S. Government Printing Office.
- Piczak, K. J. (2015). ESC: Dataset for Environmental Sound Classification. *The 23rd ACM International Conference on Multimedia*, 1015–1018.
- Pilger, C., Gaebler, P., Hupe, P., Kalia, A. C., Schneider, F. M., Steinberg, A., Sudhaus, H., & Ceranna, L. (2021). Yield estimation of the 2020 Beirut explosion using open access waveform and remote sensing data. *Scientific Reports*, 11(1). <https://doi.org/10.1038/s41598-021-93690-y>

- Pilger, C., & Hupe, P. (2024). The infrasonic signature of three exceptional rocket launches. *Proceedings of Meetings on Acoustics*, 52, 040006. <https://doi.org/10.1121/2.0001861>
- Pilger, C., Hupe, P., Gaebler, P., & Ceranna, L. (2021). 1001 Rocket Launches for Space Missions and Their Infrasonic Signature. *Geophysical Research Letters*, 48(8). <https://doi.org/10.1029/2020GL092262>
- Plakal, M., & Ellis, D. (2029). *YAMNet*. <https://github.com/tensorflow/models/tree/master/research/audioset/yamnet>
- Podglajen, A., Le Pichon, A., Garcia, R. F., G erier, S., Millet, C., Bedka, K., Khlopenkov, K., Khaykin, S., & Hertzog, A. (2022). Stratospheric Balloon Observations of Infrasound Waves From the 15 January 2022 Hunga Eruption, Tonga. *Geophysical Research Letters*, 49(19). <https://doi.org/10.1029/2022GL100833>
- Popenhagen, S. K. (2024). *Aggregated Smartphone Timeseries of Rocket-generated Acoustics (ASTRA)*. <https://doi.org/10.7910/DVN/ZKIS2K>
- Popenhagen, S. K., Bowman, D. C., Zeiler, C., & Garc es, M. A. (2023). Acoustic Waves From a Distant Explosion Recorded on a Continuously Ascending Balloon in the Middle Stratosphere. *Geophysical Research Letters*, 50(20). <https://doi.org/10.1029/2023GL104031>
- Popenhagen, S. K., & Garc es, M. A. (2025). Acoustic Rocket Signatures Collected by Smartphones. *Signals*, 6(1), 5. <https://doi.org/10.3390/signals6010005>
- RedVox Inc. (2022). *libquantum v1.4.1*. RedVox, Inc. <https://github.com/RedVoxInc/libquantum>
- RedVox Inc. (2023). *LSECE Artemis 2020-10-27 12:45:00 (90.0m) 3 stations*. <https://redvox.io/#/reports/0BB3>
- Rossi, F., Saboia, M., Krishnamoorthy, S., & Hook, J. Vander. (2023). Proximal exploration of Venus volcanism with teams of autonomous buoyancy-controlled balloons. *Acta Astronautica*, 208, 389–406. <https://doi.org/10.1016/j.actaastro.2023.03.003>
- Sagdeev, R. Z., Linkin, V. M., Kerzhanovich, V. V., Lipatov, A. N., Shurupov, A. A., Blamont, J. E., Crisp, D., Ingersoll, A. P., Elson, L. S., Preston, R. A., Hildebrand, C. E., Ragent, B., Seiff, A., Young, R. E., Petit, G., Boloh, L., Alexandrov, Yu. N., Armand, N. A., Bakitko, R. V., & Selivanov, A. S. (1986). Overview of VEGA Venus Balloon in Situ Meteorological Measurements. *Science*, 231(4744), 1411–1414. <https://doi.org/10.1126/science.231.4744.1411>
- Schwardt, M., Pilger, C., Gaebler, P., Hupe, P., & Ceranna, L. (2022). Natural and Anthropogenic Sources of Seismic, Hydroacoustic, and Infrasonic Waves: Waveforms and Spectral Characteristics (and Their Applicability for Sensor Calibration). In *Surveys in Geophysics* (Vol. 43, Issue 5, pp. 1265–1361). Springer Science and Business Media B.V. <https://doi.org/10.1007/s10712-022-09713-4>
- Silber, E. A., Bowman, D. C., & Ronac Giannone, M. (2023). Detection of the Large Surface Explosion Coupling Experiment by a Sparse Network of Balloon-Borne Infrasound Sensors. *Remote Sensing*, 15(2). <https://doi.org/10.3390/rs15020542>
- Slad, G., & Merchant, B. J. (2021). *Evaluation of Low Cost Infrasound Sensor Packages*. <https://doi.org/https://doi.org/10.2172/1829264>
- Smith, K., Solomon, M., Bryan, K. J., Smith, A. O., & Peter, A. M. (2018). Near-field infrasound classification of rocket launch signatures. *SPIE-Intl Soc Optical Eng*, 51. <https://doi.org/10.1117/12.2302680>
- Sutherland, L. C., & Bass, H. E. (2004). Atmospheric absorption in the atmosphere up to 160 km. *The Journal of the Acoustical Society of America*, 115(3), 1012–1032. <https://doi.org/10.1121/1.1631937>

- Takazawa, S. K. (2024). *Smartphone High-explosive Audio Recordings Dataset (SHAReD)*. Harvard Dataverse. <https://doi.org/10.7910/DVN/ROWODP>
- Takazawa, S. K., Popenhagen, S. K., Ocampo Giraldo, L. A., Cárdenas, E. S., Hix, J. D., Thompson, S. J., Chichester, D. L., & Garcés, M. A. (2024). A comparison of smartphone and infrasound microphone data from a fuel air explosive and a high explosive. *The Journal of the Acoustical Society of America*, *156*(3), 1509–1523. <https://doi.org/10.1121/10.0028379>
- Takazawa, S. K., Popenhagen, S. K., Ocampo Giraldo, L. A., Hix, J. D., Thompson, S. J., Chichester, D. L., Zeiler, C. P., & Garcés, M. A. (2024). Explosion Detection Using Smartphones: Ensemble Learning with the Smartphone High-Explosive Audio Recordings Dataset and the ESC-50 Dataset. *Sensors*, *24*(20). <https://doi.org/10.3390/s24206688>
- Tenney, S. M., Noble, J. M., Whitaker, R. W., & ReVelle, D. O. (2003). Acoustic/infrasonic rocket engine signatures. *The International Society for Optical Engineering*, *5090*, 30–41. <https://doi.org/10.1117/12.487119>
- Thandu, S. C., Bharti, P., Chellappan, S., & Yin, Z. (2017). Leveraging multi-modal smartphone sensors for ranging and estimating the intensity of explosion events. *Pervasive and Mobile Computing*, *40*, 185–204. <https://doi.org/10.1016/j.pmcj.2017.06.012>
- Thandu, S. C., Chellappan, S., & Yin, Z. (2015). Ranging explosion events using smartphones. *2015 IEEE 11th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, 492–499. <https://doi.org/10.1109/WiMOB.2015.7348002>
- The pandas development team. (2020, February). *pandas-dev/pandas: Pandas (2.0.1)*. Zenodo. <https://doi.org/10.5281/zenodo.3509134>
- Tsalera, E., Papadakis, A., & Samarakou, M. (2021). Comparison of Pre-Trained CNNs for Audio Classification Using Transfer Learning. *Journal of Sensor and Actuator Networks*, *10*(4), 72. <https://doi.org/10.3390/jsan10040072>
- Vergoz, J., Hupe, P., Listowski, C., Le Pichon, A., Garcés, M. A., Marchetti, E., Labazuy, P., Ceranna, L., Pilger, C., Gaebler, P., Näsholm, S. P., Brissaud, Q., Poli, P., Shapiro, N., De Negri, R., & Mialle, P. (2022). IMS observations of infrasound and acoustic-gravity waves produced by the January 2022 volcanic eruption of Hunga, Tonga: A global analysis. *Earth and Planetary Science Letters*, *591*. <https://doi.org/10.1016/j.epsl.2022.117639>
- Walter, W. R., Blom, P., & Abbott, R. (2020). *Large Surface Explosion Coupling Experiment*. International Federation of Digital Seismograph Networks. <https://doi.org/10.7914/1p2e-0167>
- Weaver, R. L., & McAndrew, J. (1995). *The Roswell Report: Fact versus Fiction in the New Mexico Desert*. U.S. Government Printing Office.
- Wermer, L., Murzyn, C., Segal, J. W., Clemenson, M. D., & Jackson, D. C. (2021). *Large Surface Explosion Coupling Experiment-SNL Remote Optical*. <https://classic.ntis.gov/help/order-methods/>
- Wescott, J. W. (1964). *Acoustic Detection of High-Altitude Turbulence*.
- Young, E. F., Bowman, D. C., Lees, J. M., Klein, V., Arrowsmith, S. J., & Ballard, C. (2018). Explosion-generated infrasound recorded on ground and airborne microbarometers at regional distances. *Seismological Research Letters*, *89*(4), 1497–1506. <https://doi.org/10.1785/0220180038>