

# AVALANCHE INFRASOUND SIGNAL CLASSIFICATION USING MACHINE LEARNING

Jeffrey B. Johnson<sup>1\*</sup>, Scott Havens<sup>2</sup>, Clark Corey<sup>2</sup>, Evi Ofekeze, Skyler Chase<sup>1</sup>, Owen Walsh<sup>1</sup>, Jerry Mock<sup>1</sup>, Hans Peter Marshall<sup>1</sup>

<sup>1</sup> Department of Geosciences, Boise State University, Boise, ID, USA

<sup>2</sup> Snowbound Solutions, Boise, ID, USA

**ABSTRACT:** Infrasound monitoring is an established technology for detecting and locating snow avalanches using array processing techniques. Infrasound arrays are collections of three (or more) closely-spaced (10-20 m) sensors deployed in network configurations typically used for operational avalanche monitoring systems. Multiple arrays are collectively used to both identify avalanche signals and map their progression over time. Properly tuned networks of arrays, such as those operating in Little Cottonwood Canyon (Utah, USA), can be used to reliably identify avalanches with destructive indices as low as D2 and occurring at distances greater than 1 km from the sensors. Arrays are used to identify signals as avalanches and also minimize false positive detections from other signals like munitions, vehicles, aircraft, and earthquakes. An infrasound network dataset from Little Cottonwood Canyon spans three years of continuous records, starting in 2021-2023, and contains more than a thousand avalanche signals.

A priority moving forward is to robustly identify avalanches using data from a single infrasound sensor. We are using the rich dataset from Little Cottonwood Canyon, and tens of thousands of avalanche signal windows, to train automated classifiers using speech recognition methodologies. Our efforts are to identify the smallest possible feature variable space to reliably identify an avalanche signal and discard non-interesting signals or noise. We show that the use of machine learning techniques to classify avalanches may permit more rapid (real-time) identification of avalanches without the need to integrate multiple channels of data. Although a single sensor is not capable of avalanche source localization it can be valuable for low-cost installations and mobile avalanche alert detection systems.

**KEYWORDS:** infrasound, snow avalanche dynamics, avalanche hazard mitigation, machine learning

## 1. INTRODUCTION

Snow avalanches produce infrasound (sound energy below 1 Hz), which can propagate long distance and be used as a remote sensing tool (e.g., Bedard et al., 1988). Low bandwidth infrasound data recorded at ~100 Hz can then be used to detect gravity driven mass wasting, including snow avalanches and these detections can be relayed to civil authorities using cell or radio telemetry (e.g., Marchetti et al., 2015; Marchetti & Johnson, 2023). Infrasound monitoring in Little Cottonwood Canyon (LCC), Utah has been operational since 2006 (Vyas et al., 2009), but recent improvements in hardware, telemetry and signal processing techniques have improved its functionality as a monitoring tool. Boise State University works to develop signal detection algorithms together with Boise-based industry partner Snowbound Solutions (e.g., Johnson et al., 2018; Johnson et al., 2021; Johnson et al., 2023). Snowbound has developed a web interface tool that rapidly posts avalanche detections as alerts (<https://snowboundsolutions.com>).

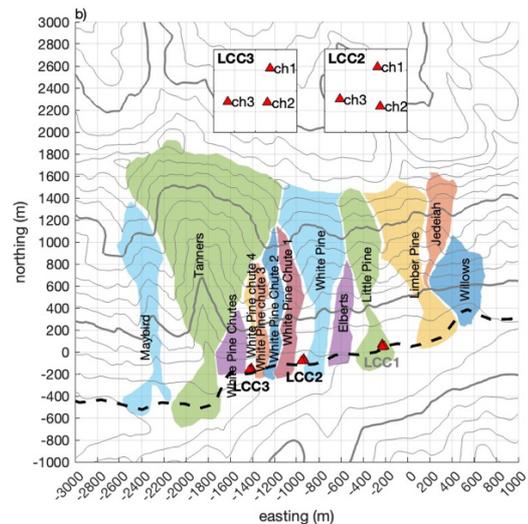


Figure 1 – Map of LCC in Utah and named slide paths near to the LCC infrasound network (red triangles). Black dashed line indicates the heavily trafficked highway U-210, which leads to Alta and Snowbird Ski resorts. Inset maps show detail of array element locations (grid spacing is 10 m).

Detection of events with low rates of false positives is a primary goal of operational infrasound avalanche monitoring systems at LCC and elsewhere (e.g., Mayer et al., 2020). Toward this end machine learning (ML) event classification using single-sensor waveform attributes (i.e., features) should be able to improve the robustness of avalanche signal detection and also – potentially – allow single sensor infrasound recordings to be used to reliably detect infrasound. Deep learning ML techniques have previously been applied to (non-avalanche) infrasound data for global and regional detection objectives [e.g., Bishop et al., 2022]. In that promising study the classes were broadly grouped as transient, noise, moving sources, and persistent sources. Here we use somewhat more flexible supervised classification to distinguish between noise, explosions, and avalanches using locally recorded infrasound data.

## 2. METHODOLOGY

An avalanche infrasound classifier begins with a training dataset used to build a model for signal identification. In section 2.1 and 2.2 we briefly summarize a technique to identify signal using array processing [Johnson et al., 2024] and then assign it a label as either explosion or avalanche. Signal feature extraction is then implemented using established audio feature libraries and used to test various ML models as described in section 2.3. Finally in section 2.4, the best model is then tested with independent data for which class labels are known.

### 2.1 Signal Identification Using Arrays

Snow avalanche signals are identifiable using cross correlation analysis of infrasound data recorded on spatially separated microphones. Waveforms in LCC are recorded with a 100 Hz sample rate and with GPS precision timing. Comparison of signals on three different sensor pairs (for three channels of data) are quantified as cross-correlation functions applied to finite overlapping time windows (5 or 10 s) with 1 s time steps. These cross-correlation functions can be displayed graphically as heat maps with hot colors indicating high correlation (Figure 2). Such *correlograms* can then be visually inspected to identify snow avalanches based upon extended duration (>20 s) signals and (often) shifting peak cross correlation lag times that indicate a moving source whose infrasound back azimuth changes over time [e.g., Marchetti & Johnson, 2023].

For any time window the signal *consistency* may be calculated as the sum of lag times  $\Delta t$  for all three sensor pairs. Consistent cross correlation lags mean that  $C = |\Delta t_{21} + \Delta t_{32} + \Delta t_{13}| \rightarrow 0$ .

Our avalanche detection system in LCC currently makes use of signal identification from a network of two arrays, with three elements each. An array, referred to as an *infrasound station*, has three or more spatially distributed infrasonic microphones with inter-sensor spacing of about 15-20 m. A network of stations integrates two or more arrays separated by relatively large distances, which in the case of LCC is more than 500 m. The network topology is optimized for sound source triangulation using back azimuth information from at least two arrays. In LCC the network of three stations is located along the U-210 highway corridor (Figure 1). This study makes use of station LCC2's data to build a training dataset of avalanche, and explosion signals following detection methods outlined in Johnson et al. (2023). The model is then applied to LCC3 test data to validate ML capabilities for single-sensor detections.

Signal quality is also quantified as the average of the peak normalized cross correlation ( $r$ ) for all three sensor pairs, e.g.,  $\bar{r} = (r_{21} + r_{32} + r_{13})/3$ .

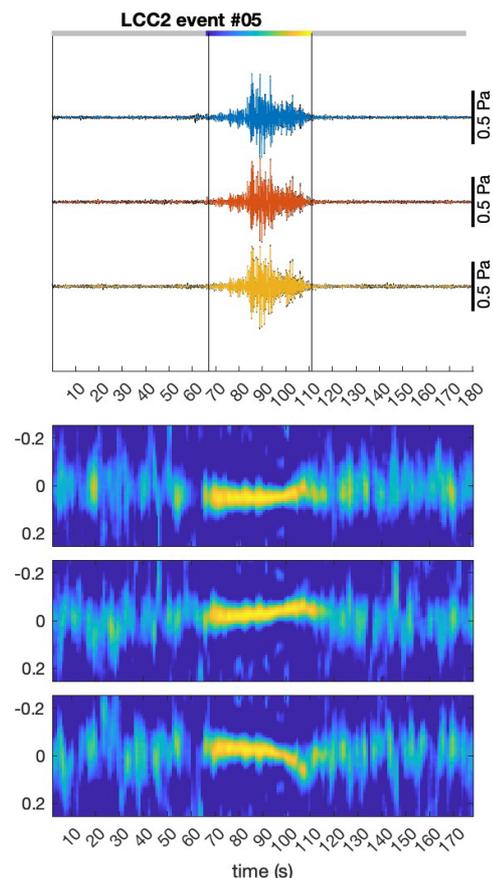
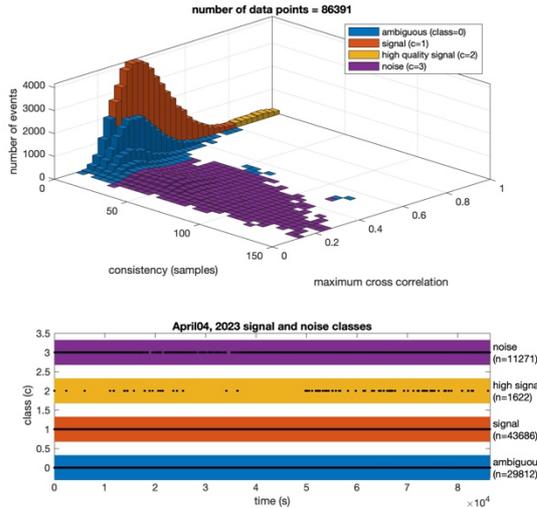


Figure 2- (>2 Hz) signal from three infrasound channels and cross-correlation correlograms for an event recorded on LCC2 in 2023. Cross-correlation matrices (bottom) show comparisons of infrasound data from the three distinct pairs for overlapping 5 s windows. Y-axes indicate time delays (or lags) of sensor comparisons. (reproduced from Johnson et al., 2023)



Deterministic classifications can be accomplished using  $C(t)$  and  $\bar{r}(t)$ . For a full day's worth of data each data window for a three-element array has its own consistency and average maximum cross-correlation coefficient. These may be plotted as a two-dimensional histogram whose bar heights correspond to the number of data points in each bin (figure 3).

Figure 3 (left) – 24-hour signal-versus-noise classification using array processing parameters for 10 s overlapping windows with 1 s time steps. (top) 2D histogram showing the distribution of time windows as a function of consistency and normalized cross correlation coefficients. (bottom) four classes assigned on the basis of consistency and maximum cross correlation features. High signal quality events are further analyzed to build training dataset of explosions and avalanches.

## 2.2 Labeling and Signal Feature Extraction

We arbitrarily define classes for the processed array data as *signal* (red or orange in Figure 2) if consistency is less than or equal to 5 samples ( $C \leq 0.05$  s) and if average cross-correlation is also greater than  $\bar{r} \geq 0.2$ . High quality signal corresponds to  $C \leq 0.05$  s and  $\bar{r} \geq 0.8$ . Those time windows representing 13% of all data from LCC2023 on April 4, 2023 are subsequently used to train the model for either explosion or avalanche designation. Noise is another class label identified by poor consistency ( $C \geq 0.25$  s) as well as low correlation ( $\bar{r} \leq 0.5$ ). Ambiguous signals

(blue in Figure 2; class=0) and low-correlation consistent 'signal' (red in Figure 2; class=1) are not used in our model training. Instead, we limit our classification in this study to those time windows with either high quality signal (explosion or avalanche) and definite noise. Candidate explosion and avalanche signals are initially picked automatically by calculating peak cross correlations over both short (10 s) and longer (30 s) time windows. Correlation duration is an effective criterion for differentiating snow avalanches, which are extended duration in time, from explosions, which are manifested as short (~1 s) transient infrasound pulses.

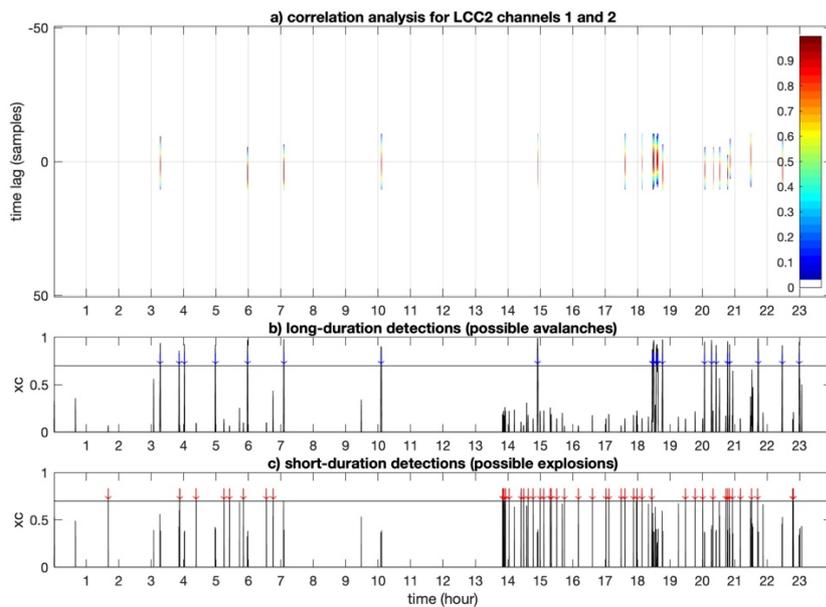


Figure 4 – 24-hour record showing training dataset candidate classifications for high quality signal (class #2 in Figure 3). (a) correlation analyses plotted as a function of time of day and time lag (y-axis) using 10 s windows and 1 s time steps. (b) long-duration classifications correspond to 30 second running average of maximum cross-correlation values from panel a. (c) short-duration classifications show 10 s maximum value cross correlations. Blue and red arrows indicate peak values of  $r > 0.7$  and are displayed in Figures 5 and 6.

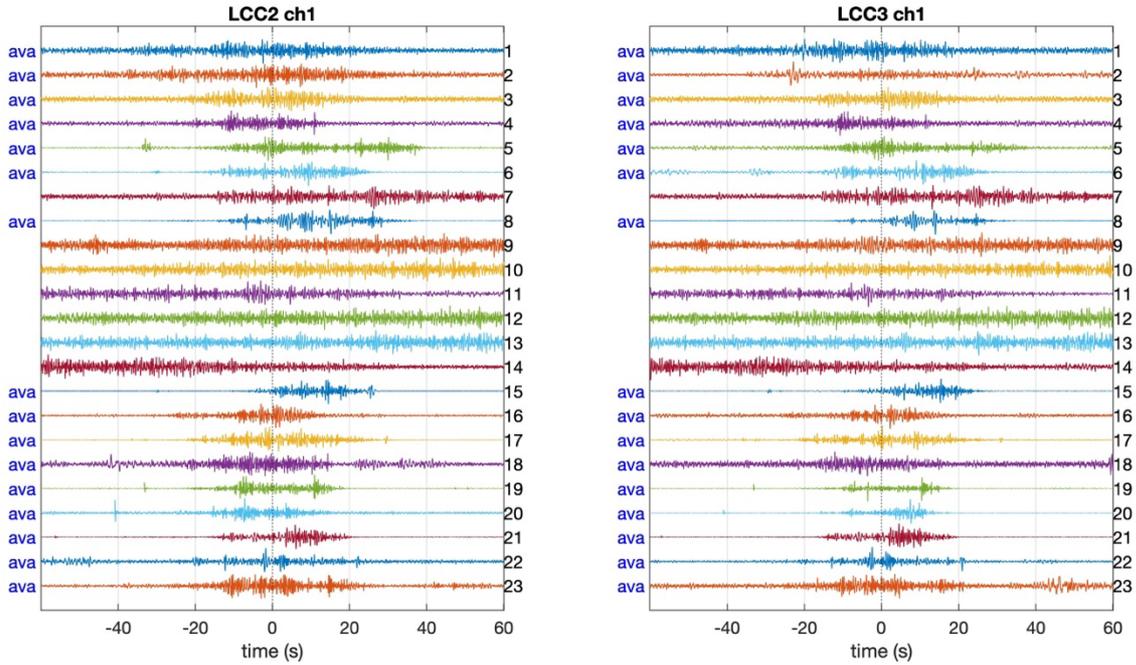


Figure 5 – 120-s time windows for automatically-picked long-duration detections (from Figure 5) using the LCC2 array. Waveforms are shown from channel 1 of LLC2 (left) as well as channel 1 of LCC3 (right), which was not used to find events. Events manually identified as definite avalanches, according to waveform envelope are indicated by ‘ava’.

### 2.3 Feature Extraction:

Windowed events are identified as explosion signal (class = 4) or avalanche signal (class = 5) based upon whether they are short-duration or long-duration and consistent with highly correlated signal. This processing identifies candidate events well, but additional manual validation is still needed to identify signals.

For instance, out of the 23 candidate avalanche signals shown in Figure 6, only 16 were identified as definite avalanche based upon visual inspection. These waveforms are from channel 1 and from the two different arrays separated by ~500 m. Signals reveal signal similarities, in terms of waveform envelope, but they have differing amplitudes and signal-to-noise quality. Of the 46 candidate explosion waveforms in Figure 6, 14 are discarded based following visual inspection and non-identification of short impulsive transients.

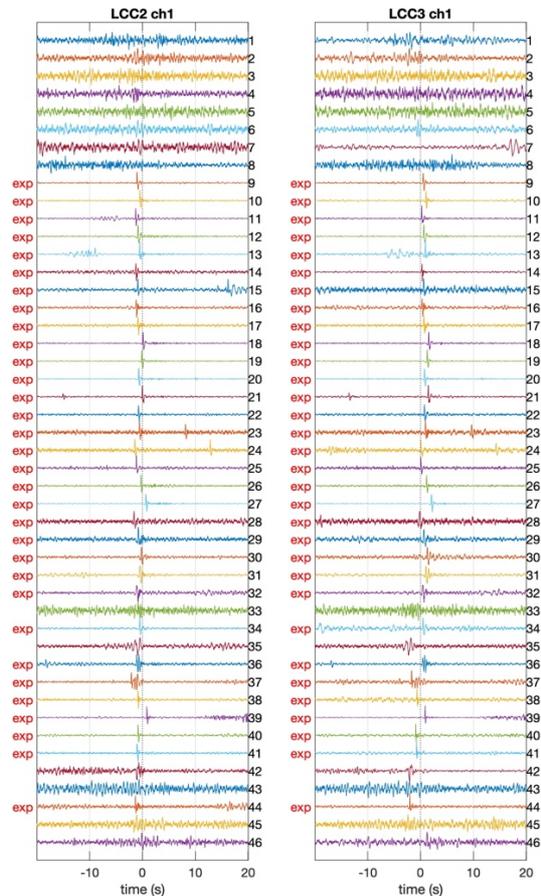


Figure 6 (right) – 40-s time windows for automatically-picked short-duration detections (from Figure 4) using array processing applied to LCC2. Waveforms are shown from channel 1 of LLC2 (left) and channel 1 of LCC3 (right). Those events visually labeled as explosions are indicated by ‘exp’.

The training dataset for one day thus consists of 32 explosions and 16 avalanches. Because each data point is a 10 s window (and avalanches are generally 30 to 60 s in duration) we are able to train our model with 21 overlapping time windows such that there are 336 avalanche data points. Similarly, we use 7 overlapping time windows to train the explosion class such that there are 224 data points corresponding to 32 identified explosions. These events are in addition to the 11,271 data windows labeled as noise from Figure 3.

Each time window of 10 seconds (or 1000 samples) is then quantified in terms of a 54-dimension feature space. We borrow from feature space metrics used in audio voice identification ML in identifying the most relevant features. We make use of MATLAB's audioFeatureExtractor recognizing that infrasound is simply sound signals of low frequency and presumably can be similarly characterized using a similar number of samples (10 s = 1000 samples). A thousand samples is similar to the data length used in audio recognition. The feature space explored in this study consists of those variables shown in Table 1.

Feature name	#values	MRMR rank
mfcc (cepstral frequencies)	1-13	12 (#3), 5 (#8)
mfcc derivative	14-26	14 (#2), 20 (#7)
gtcc (gammatone frequencies)	27-39	30 (#1),39,35,33,37
spectralCentroid	40	
spectralCrest	41	
spectralDecrease	42	
spectralEntropy	43	
spectralFlatness	44	
spectralFlux	45	
spectralKurtosis	46	
spectralRolloffPoint	47	
spectralSkewness	48	
spectralSlope	49	
spectralSpread	50	
pitch	51	51 (#9)
harmonicRatio	52	
zerocrossrate	53	
shortTimeEnergy	54	

Features shown in Table 1 are defined in MATLAB's documentation at <https://www.mathworks.com/help/audio/ref/audiofeatureextractor.html>. These features, alternatively known as predictors, are used to train the ML models along with class designations, known as responses. In this work response is limited to noise, explosions, and avalanches.

### 2.4 Avalanche Classification:

MATLAB's *classificationLearner* is used to determine the best classifier model along with feature optimization. Accuracy was initially estimated using all 54 features and for models including *Fine Tree*, *Linear Discriminant*, *Naïve Bayes*, *Linear Support Vector Machine*, *Nearest Neighbor*, and *Narrow Neural Network* models. Validation gave accuracies between 97.1% for Naïve Bayes and 99.8% for Nearest Neighbor. These 'worst' and 'best' confusion matrices are shown in Figure 7.

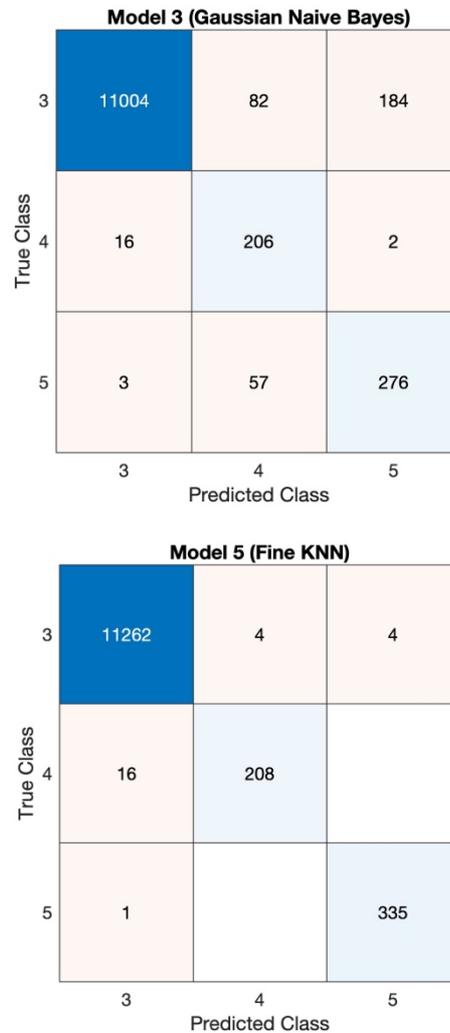


Figure 7 – Example confusion matrices for two different models using all 54 features. The fine KNN model offers the best predictors of class. In these matrices class 3 corresponds to 'noise', class 4 is explosions, and class 5 corresponds to avalanche windows.

The KNN model appears best and is further analyzed to see how accuracy changes by limiting the feature space. Feature importance scores are estimated using the MRMR algorithm in the classificationLearner application, which allows ranking of relevant predictors. The top ten features are indicated in Table 1. Using only these features the resulting confusion matrix (with an accuracy of 99.6%) is shown in Figure 8. True positive rates (TPRs) are 99.4% for avalanches and 85.3% for explosions while False Discovery Rates (FDRs) are 5.4% and 2.9% for explosions and avalanches respectively.

Figure 8 (right) – Confusion matrix for fine KNN model using only the ten most important features indicated in Table 1. Here class 3 is ‘noise’, class 4 is explosions, and class 5 corresponds to avalanche windows.

**Model 8 (Fine KNN)**

		3	4	5
3	11255	10	5	
4	28	191	5	
5	1	1	334	
		3	4	5
		Predicted Class		

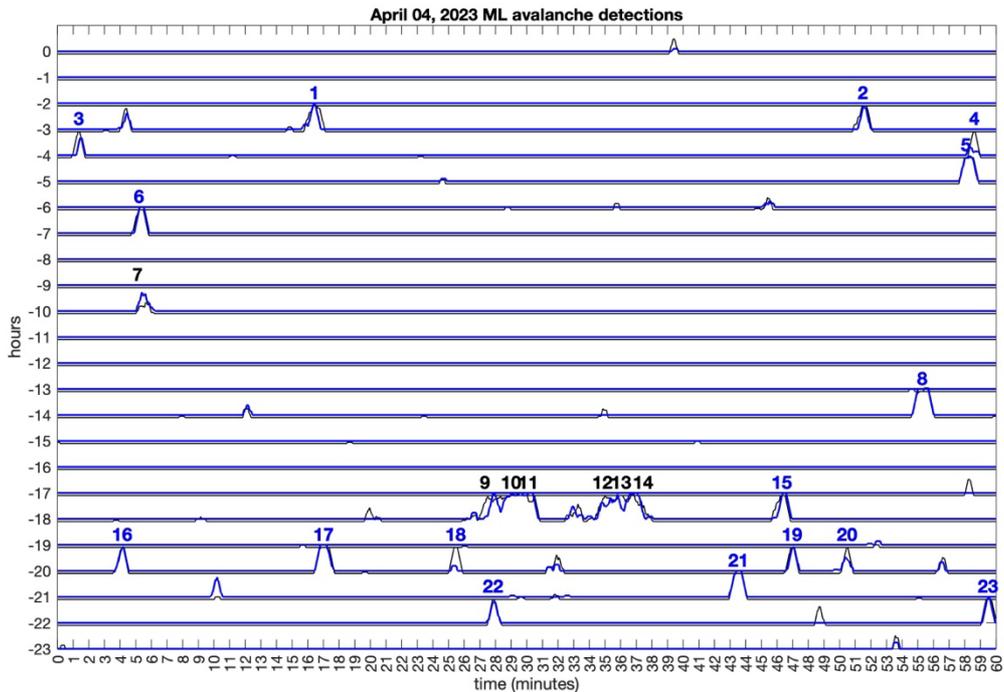


Figure 9 – 24-hour record showing classification detections using array LCC2 ch1 (training dataset; black) and LCC3 ch1 (test dataset; blue) identifying avalanche class events. Each line is one hour of data. Avalanche classifications have been convolved with a 30-s running mean for visualization purposes. Numeric values correspond to waveforms numbering shown in Figure 5.

### 3. RESULTS

While it is reassuring that simple KNN models using only ten features can be used to fit the training data well, a true test of the model must be applied to labeled data that were not part of the training process. Because the model was developed using data collected exclusively at LCC2 we are able to validate model efficacy using data recorded on any of the channels of the LCC3 array. Utilizing the ten most important features for all 86,391 time windows (i.e., ~1 day with 10 s

windows, and 1 s steps) we attempt to classify all events as noise, explosions, or avalanches.

For visualization purposes explosion detections and avalanche detections are plotted in two separate graphs (see Figures 9 and 10) along with classifications results using the same model and the original training data. The ML classification appears to do a commendable job at finding both explosions and avalanches in the LCC3 waveform data that were ‘heard’ at station LCC2. This demonstrates the general capability of identifying avalanches and explosions using a single

isolated sensor. It is noteworthy that avalanches detected at LCC2 were also well recorded at LCC3 using only waveform features. Black numbered 'events', which were considered ambiguous as avalanche signals, are also picked using the ML model while some additional events (e.g., at 3:04 and 21:10), that were not found in the deterministic modeling appear to be actual avalanches upon visual inspection. This last observation is particularly exciting because it suggests optimistically that audio feature extraction using

single channel infrasound data might be able to identify otherwise obscured avalanches.

Explosion detection results (Figure 10) are also promising. The classifier appears to identify almost all the explosions that were classified as explosions on LCC2. Only events 15, 17, 30, and 38 were not detected on LCC3. It remains to be seen if those events were not picked because they were diminished in signal-to-noise due to ambient noise or distance-to-source differences.

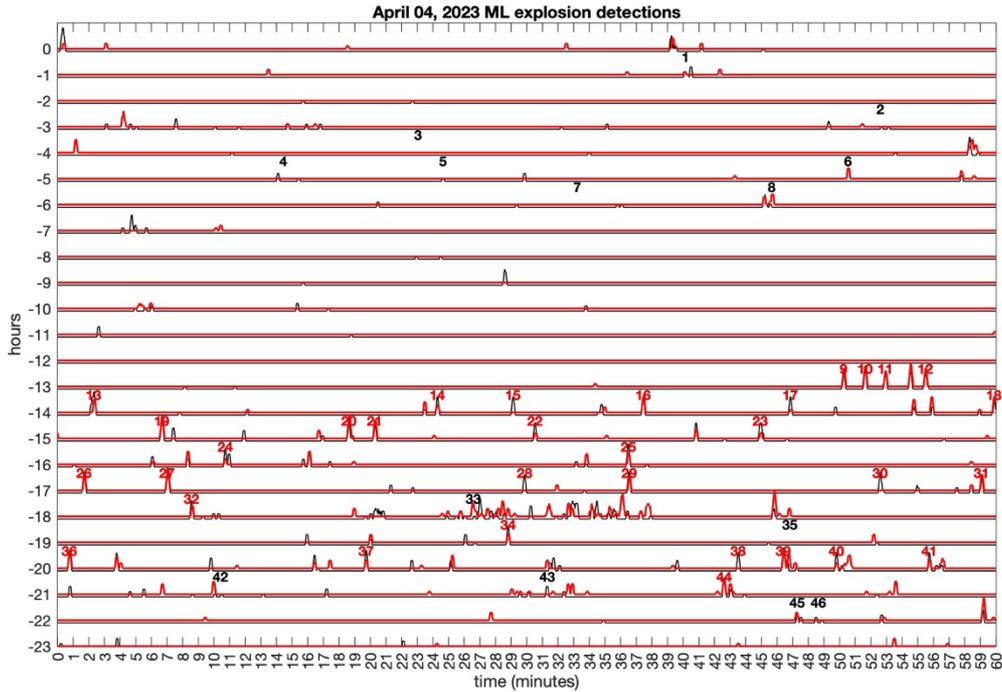


Figure 10 – 24-hour record showing classification detections using array LCC2 ch1 (training dataset; black) and LCC3 ch1 (test dataset; red) identifying explosion class events. Explosion classifications have been convolved with a 10-s running mean for visualization. Numbered values correspond to waveforms shown in Figure 6.

#### 4. CONCLUSION AND NEXT STEPS

It is unsurprising that infrasonic records can be classified using audio features, which are used extensively and successfully in the voice recognition community. It is also exciting that a basic ML KNN classifier appears to do such a good job of finding avalanches and explosion signals using a relatively small amount of training data. Moving forward additional testing is necessary. First, we plan to apply the detection algorithm to the full avalanche season in 2023, which includes many hundreds of confirmed snow avalanches. We also wish to expand the signal class designations to include other types of infrasound signal such as vehicular traffic (e.g., cars, trucks, planes, and helicopters) because we suspect that some of the classified explosions are erroneous.

We also intend to expand the training dataset by including additional time windows and more sensor data. The LCC network consists of 9 continuous data feeds and all 9 channels can and should be used for training purposes. Finally, all results presented here correspond to data from a single day. Moving forward we plan to use data from an entire season. By increasing the number of channels of data and making use of the full 2023 season we estimate we can build a future model using more than 10,000 avalanche and explosion signal data points in the training data. Applying an improved model to infrasonic data collected outside of LCC will be a final confirmation of its utility. The ultimate goal of using a single channel of infrasound to identify avalanches is possible although arrays will still have their place for those seeking to determine back azimuth.

## ACKNOWLEDGEMENT

The original work in LCC was made possible through grants to Boise State University from the Transportation Avalanche Research Pool (TARP). We appreciate the financial support as well as the discussion, field assistance, and on-the-ground efforts of Utah and Colorado Departments of Transportation (UDOT/CDOT). Data

collected in 2023 were made available by Snowbound Solutions through collaboration with Boise State University who provide the infrasound sensing hardware. This work was funded by U.S. Army CRREL, “Advancement of snow monitoring for water resources, vehicle mobility, and hazard mitigation: using optical, microwave, acoustic, and seismic techniques”, Grant Number W913E520C0017.

## REFERENCES

- Bedard, A. J., Greene, G. E., Intrieri, J., & Rodriguez, R. (1988). On the feasibility and value of detecting and characterizing avalanches remotely by monitoring radiated sub-audible atmospheric sound at long distances. In *Proceedings of A Multidisciplinary Approach to Snow Engineering, U.S. Army Corps of Engineers; Cold Regions Research and Engineering Laboratory; Santa Barbara; Calif.* (pp. 267–275).
- Bishop, J. W., Blom, P. S., Webster, J., Reichard-Flynn, W., & Lin, Y. (2022). Deep learning categorization of infrasound array data. *The Journal of the Acoustical Society of America*, 152(4), 2434–2445. <https://doi.org/10.1121/10.0014903>
- Havens, S., Marshall, H.-P., Johnson, J. B., & Nicholson, B. (2014). Calculating the velocity of a fast-moving snow avalanche using an infrasound array. *Geophysical Research Letters*, 41(17), 6191–6198. <https://doi.org/10.1002/2014GL061254>
- Johnson, J. B., Anderson, J. F., Marshall, H. P., Havens, S., & Watson, L. M. (2021). Snow Avalanche Detection and Source Constraints Made Using a Networked Array of Infrasound Sensors. *Journal of Geophysical Research: Earth Surface*, 126(3), 1–23. <https://doi.org/10.1029/2020JF005741>
- Johnson, J. B., Marshall, H. P., & Anderson, J. F. (2021). *Infrasound avalanche monitoring: Enhanced monitoring in Little Cottonwood Canyon, Utah*. Retrieved from <https://rosap.nrl.bts.gov/view/dot/60194>
- Johnson, J. B., Marshall, H. P., Loo, S. M., Nalli, B., Saurer, M., Havens, S., et al. (2018). Detection and tracking of snow avalanches in Little Cottonwood Canyon, Utah using multiple small-aperture infrasound arrays. *Proceeding of International Snow Science Workshop, Innsbruck, Austria, 7–12 October 2018*.
- Johnson, J. B., Marshall, H. P., Loo, S. M., Nalli, B., Saurer, M., Havens, S., et al. (2018). Heard but not seen: automated detection and quantification of snow avalanches using infrasound. *Proceeding of International Snow Science Workshop, Bend, OR, USA, 8–13 October 2023*.
- Marchetti, E., & Johnson, J. B. (2023). Infrasound array analysis of rapid mass movements in mountain regions. (C. Schmelzbach, Ed.), *Advances in Geophysics: Geohazards* (Advances in, Vol. 64). Zurich: Academic Press.
- Marchetti, E., Ripepe, M., Olivieri, G., & Kogelnig, A. (2015). Infrasound array criteria for automatic detection and front velocity estimation of snow avalanches: towards a real-time early-warning system. *Natural Hazards and Earth System Sciences Discussions*, 3(4), 2709–2737. <https://doi.org/10.5194/nhessd-3-2709-2015>
- Mayer, S., van Herwijnen, A., Olivieri, G., & Schweizer, J. (2020). Evaluating the performance of an operational infrasound avalanche detection system at three locations in the Swiss Alps during two winter seasons. *Cold Regions Science and Technology*, 173(December 2019), 102962. <https://doi.org/10.1016/j.coldregions.2019.102962>
- Vyas, M. (2009). *Infrasound avalanche monitoring system research evaluation*. Retrieved from <http://digitallibrary.utah.gov/awweb/awarchive?type=file&item=18921>
- Watson, L. M., Carpenter, B., Thompson, K., & Johnson, J. B. (2022). Using local infrasound arrays to detect plunging snow avalanches along the Milford Road, New Zealand (Aotearoa). *Natural Hazards*, 111(1), 949–972. <https://doi.org/10.1007/s11069>