This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

SAND2019-5288J

Classification of Local Seismic Events in the Utah Region: A Comparison of Amplitude Ratio Methods with a Spectrogram-Based Machine Learning Approach

Rigobert Tibi[1], Lisa Linville[1], Christopher Young[1], and Ronald Brogan[2]

[1]Sandia National Laboratories,

P.O. Box 5800,

Albuquerque, NM 87185-0404

Email: rtibi@sandia.gov

Phone: 505-844-9843

[2]ENSCO, Inc.

5400 Port Royal Road

Springfield, VA 22151

To be submitted to *Bulletin of the Seismological Society of America*

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.

Abstract

The capability to discriminate low-magnitude earthquakes from low-yield anthropogenic sources, both detectable only at local distances, is of increasing interest to the event monitoring community. We used a dataset of seismic events in Utah recorded during a 14-day period (1–14 January 2011) by the University of Utah Seismic Stations (UUSS) network to perform a comparative study of event classification at local-scale using amplitude ratio (AR) methods and a machine learning (ML) approach. The event catalog consists of 7,377 events with magnitudes M_C ranging from -2 and lower up to 5.8. Events were subdivided into six populations based on location and source type: tectonic earthquakes (TE), mining-induced events (MIE), and mining blasts from four known mines (WMB, SMB, LMB, and CQB). The AR approach jointly exploits *P*g-to-*S*g phase amplitude ratios and *R*g-to-*S*g spectral amplitude ratios in multivariate quadratic discriminant functions (QDFs) and was able to classify 370 events with high signal quality from the three groups with sufficient size (TE, MIE, and SMB). For that subset of the events, the method achieved success rates between about 80–90%. The ML approach used pre-trained convolutional neural network (CNN) models to classify the populations. The CNN approach was able to classify the subset of events with accuracies between about 91–98%. Because the neural network approach does not have a minimum signal quality requirement, we applied it to the entire event catalog, including the abundant extremely low-magnitudes events, and achieved accuracies of about 94–100%. We compare the AR and ML methodologies using a broad set of criteria and conclude that a major advantage to machine learning methods is their robustness to low signal-to-noise ratio (SNR) data, allowing them to classify significantly smaller events.

Introduction

A variety of seismic event identification techniques using data recorded at near-regional to teleseismic distances (> 200 km) have been successfully used over four decades to discriminate moderate-size and larger events (e.g., Dowla *et al.*, 1990; Kim *et al.*, 1993; Walter *et al*., 1995; Taylor, 1996; Stump *et al.*, 2002; Bowers and Selby, 2009; and references therein). Most of these discrimination methods exploit differences in energy partitions between specific seismic phases and event types (i.e., the difference between earthquakes and explosions in the excitation of *P* and *S* waves). Recently, some of these techniques have been adapted for local distances (≤ 200 km), driven by an increased interest in the capability to discriminate low-magnitude earthquakes from low-yield anthropogenic sources, both detectable only at short distances (Kafka, 1990; Zeiler and Velasco, 2009; Koper *et al.*, 2016; O'Rourke *et al.*, 2016; Tibi *et al.*, 2018a,b; Kolaj, 2018). Other approaches developed to identify events at local distances include those based on machine learning (ML) algorithms (e.g., Del Pizzo *et al.*, 2003; Benbrahim *et al.*, 2007; Mousavi *et al.*, 2016; Linville *et al.*, 2019b). ML methods are now being increasingly used in seismology, with applications including signal detection, phase and polarity determination (Ross *et al.*, 2018ab), event location (Perol *et al.*, 2018; Kriegerowski *et al.*, 2019), aftershock location forecasting (DeVries *et al.*, 2018), and, as mentioned earlier, event classification (Del Pizzo *et al.*, 2003; Benbrahim *et al.*, 2007; Mousavi *et al.*, 2016; Linville *et al.*, 2019b). Like in other fields, the increased interest for ML methods in seismology has been encouraged by the combination of several factors, including the availability of low or no cost sophisticated ML software (e.g., Scikit-Learn, TensorFlow), the increased availability of data for model building, and the decreased cost of required computational resources (Jordan and Mitchell, 2015; Bergen *et al.*, 2019).

In this study, we used a dataset of seismic events in Utah recorded during a 14-day period (1–14 January 2011) by the UUSS network to perform a comparative study of event classification at local distances involving an amplitude ratio (AR) method and an ML approach. The dataset was compiled by an expert seismic data analyst (more than 30 years of experience), who manually scanned through continuous seismic waveform data trying to identify and locate as many events as possible. The Utah region is tectonically active and is also home to extensive mining activities, resulting in a variety of seismic sources (e.g., Whidden and Pankow, 2012). Consequently, the dataset contains tectonic earthquakes, mining-induced events (events caused by pillar or roof collapses after the extraction of mined materials), and surface blasts from several mines. In the AR method, amplitude ratios are exploited in discrimination functions that have been designed to achieve maximum separation between the populations under consideration. The spectrogram-based ML approach involves pre-trained CNN models. To our knowledge, this study is the first of its kind that directly compares/contrasts the performances—in terms of discrimination power—of a traditional, physics-based (AR) method with a data-driven (ML) approach using the same dataset of diverse local seismic sources. The takeaways from the comparison are: (1) the relative performance of the two approaches is highly dependent on the event population; (2) the ML method is robust to low SNR data, allowing it to classify significantly smaller events.

Data

The end product of the expert's analysis of UUSS data described above is a comprehensive list of events with associated phases known as the Unconstrained Utah Event Bulletin (UUEB) (Linville *et al.*, 2019a). The UUEB is characterized as "unconstrained" because it was not subject to the typical constraints on analysis time or on minimum number of detecting stations to confirm an

event that are used by event monitoring agencies that must make their bulletins available promptly and that want to limit the content of their bulletins to including only events that are relevant to their monitoring mission (e.g., for the UUSS this is assessing the potential hazard associated with the Intermountain Seismic Belt, ISB). Without such constraints, the UUEB contains many more events than the UUSS bulletin does for the same time period. Although the UUEB includes some teleseismic and regional events, this study focuses on 7,377 events that occurred within and around the state of Utah, at local to near-regional distances to the stations of the UUSS network. The dataset consists of events with magnitude M_C (magnitude from coda duration; J. C. Pechmann *et al.*, unpublished manuscript, 2010; see Data and Resources) ranging from about -2 and lower up to 5.8 (Fig. 1, left). The frequency-magnitude distribution of the dataset is linear for M_C above 2.7 (Fig. 1, right), with a maximum-likelihood *b*-value of 0.92±0.03. However, given the fact that magnitudes for a large number of smaller events (about 58% of the entire catalog) could not be estimated due to low signal quality, we believe that the magnitude of completeness is most likely lower than 2.7. A heat map of the event locations (Fig. 2, left) reveals several well-developed clusters, beside the diffuse set of events associated with the nearly NS-trending ISB (Smith and Arabasz, 1991). The most prominent of those clusters is located around 38.27°E and 112.29°W within the ISB and is associated with the January 2011 Circleville earthquake sequence (Arrowsmith *et al*., 2011). The other clusters are related to activities in several known mining areas.

We created event categories based on the event locations with respect to known mining activities. Figure 2 (middle) shows a google map depicting the locations of the events along with significant mine locations. The catalog contains six categories of events that are defined as follows.

Five event categories are associated with mining activities. Four of these event categories resulted from mining blasts occurring at specifically known mine locations: (1) Westmoreland coal mine, WMB; (2) Bingham Canyon copper mine, SMB; (3) Lisbon Valley copper mine, LMB; and (4) Cricket Mountain quarry, CQB. A fifth category consists of mining-induced events (MIE) not located at either WMB, SMB, LMB, or CQB. These mining-induced events were identified as occurring in the Wasatch Plateau–Book Cliffs region in central Utah known for its extensive coal mining activities involving long-wall mining techniques (e.g., Arabasz *et al.*, 2005).

Any event that did not fit one of our five event categories associated with mining activities was classified as a tectonic earthquake (TE). Additionally, any event with an estimated focal depth of 3 km or larger—no matter its geographic location—was considered too deep to be related to human activities and so was classified as a TE. We acknowledge that mines in Utah exist outside of the five defined mining locations, and therefore, we expect that some of the TEs are misclassified mining blasts. However, our heat-map of event distribution suggests that additional mines were not very active during the UUEB interval, so we expect there to be relatively few such misclassifications.

Based on the procedure described above, we identified the following events in each of the six categories:

- 1,040 TEs located mostly within the ISB. This represents a significantly larger number of TEs than what one would get in a typical two-week period, due to the Circleville sequence (Arrowsmith *et al*., 2011), which accounts for the majority of TEs;

- 6,286 events induced by mining activities in the permitted coal mining regions in central Utah (MIEs). These events are caused by pillar or roof collapses from long-wall mining (Arabasz *et al.* 2005; Whidden and Pankow, 2012);
- 27 ripple-fired mining blasts from the Bingham Canyon open-pit copper mine southeast of the Great Salt Lake (SMBs, Pankow *et al.*, 2014; Tibi *et al.*, 2017);
- 5 events associated with mining activities in Lisbon Valley copper mine, near the Utah/Colorado border (LMBs);
- 16 events associated with the Westmoreland coal mine in southwestern Wyoming (WMBs); and
- 3 quarry blasts from the Cricket Mountain quarry in west central Utah (CQBs).

The event locations reported in the UUEB were obtained with the location technique described by Bratt and Bache (1988), using the AK135 velocity model (Kennett *et al.*, 1991). More accurate locations could have been determined by using different 1D velocity models for the different tectonic provinces within Utah or even using a 3D velocity model, but for our purposes the AK135 locations were sufficient. We verified that the reported anthropogenic events are clustered within known mine craters or areas of active coal-mining, as expected. In Figure 2 (right) only events for which we could estimate the coda duration magnitude (M_C) values are shown. The size of each symbol in that map is proportional to the M_C value of the event it represents. For the amplitude-ratio-based event identification method, we restricted our investigations to events with high-quality data by imposing minimum magnitudes of 2 and 2.8 for the abundant TEs and MIEs, respectively; we also required that the SNR estimated from the vertical component waveform be $\geq$ 6 dB for a least three recording stations. This resulted in a dataset of 124 events for the TE

group. The resulting MIE dataset consists of 219 events. Due to the limited number of the SMBs in the UUEB catalog (total of only 27 events), we refrained from applying minimum magnitude cutoff to that group. The sizes for the SMBs range from M_C 1.6 to ~3. We found no indication that our AR results have been compromised because of the small size for some of these events. We did not investigate the LMBs, WMBs and CQBs using the AR method because the limited number of events for each of those populations precludes sound statistics. Since, as described below in "Methods" section, the phase and spectral amplitudes were measured on the vertical, radial, and/or tangential components, according to the expected particle motion of the phase of interest, only three-component stations were used. On the other hand, for the ML approach, both single and three-component stations were used (zero-filled spectrograms were used when horizontal channels were not available). We used event origin time and hypocenter information from the catalog and retrieved waveform data for each event-station pair from the Incorporated Research Institutions for Seismology (IRIS) Data Management Center.

Methods

Phase and Spectral Amplitude Ratio Estimation

The procedures followed to estimate the phase and spectral amplitude ratios exploited for the event discrimination are discussed thoroughly in Tibi *et al.* (2018a,b). Only key aspects of these procedures are described in this section. For each event investigated, broadband and short-period stations at epicentral distances from 25 km to 150 km were used for both *Pg*-to-*Sg* and *Rg*-to-*Sg* ratios. The minimum distance cutoff of 25 km allows for sufficient time separation between *Pg* and *Sg* and between *Sg* and *Rg*. The maximum distance of 150 km was chosen to avoid interferences at larger distances due to *Pg*-*Pn* crossover. Three-component (Z, N and E) waveform data were

instrument-corrected to true ground velocity. The (Z) component waveforms of the instrument-corrected data are first filtered with a four-pole Butterworth bandpass filter 1–15 Hz (for *Pg*-to-*Sg* ratios) or 0.5–10 Hz (for *Rg*-to-*Sg* ratios) for the sole purpose of estimating the SNR. To ensure that results are not unduly influenced by background noise, only stations with SNR ≥ 6 dB are used for estimating the ratios.

For *Pg*-to-*Sg* ratios, the horizontal components of the instrument-corrected waveforms are rotated to radial (R), and tangential (T) components. For each frequency from 1 to 15 Hz in 1-Hz increment, *Pg*-to-*Sg* ratios are calculated as follows. The Z-, R-, and T-component waveforms are filtered using a bandpass filter with a 1-Hz bandwidth centered at the frequency of interest. *Pg* amplitude at that frequency is defined as $A_{P_g} = (Pg_Z^2 + Pg_R^2)^{1/2}$, the vector sum of the root mean square of the amplitudes in *Pg* window for the Z and R components (Fig. S1). Similarly, *Sg* amplitude is defined as $A_{S_g} = (Sg_Z^2 + Sg_R^2 + Sg_T^2)^{1/2}$, the vector sum of the root mean square of the amplitude in *Sg* window for the Z, R and T components. Details on how we determined the *Pg* and *Sg* windows are given in Tibi *et al.* (2018b). The calculated A_{Pg} and A_{Sg} values are then corrected for attenuation and geometrical spreading. The *Pg*-to-*Sg* ratios are obtained by dividing the corrected A_{Pg} by the corrected A_{Sg}. The estimated ratios show no correlation with magnitude. For that reason, there was no need to correct them for event size. For each event and frequency, the reported ratio consists of the average value over three or more stations, depending on the number of recording stations for that event that satisfy the selection criteria.

To estimate the *Rg*-to-*Sg* ratios, for each event-station pair that satisfies the data selection criteria described above, a mean time series of spectral amplitudes for the frequency range of 0.5–2 Hz is generated from the Z-component spectrogram to capture *Rg* energy (Fig. S2, left). A

similar time series for the frequency range of 0.5–8 Hz is generated from the T-component spectrogram to capture *Sg* energy (Fig. S2, right). $max\{Rg_{[0.5-2Hz]}\}$ is defined as the maximum amplitude of the smoothed Z component mean time series in a window of 20 seconds, starting from the *Rg* arrival time predicted for a group velocity of 1.8 km/s (Tibi *et al.*, 2018a). Similarly, $max\{Sg_{[0.5-8Hz]}\}$ is defined as the maximum amplitude of the smoothed T component mean time series in a window of 20 seconds or less (depending on the time separation between *Sg* and *Rg)*, starting from the predicted *Sg* arrival time. Both $max\{Rg_{[0.5-2Hz]}\}$ and $max\{Sg_{[0.5-8Hz]}\}$ are corrected for the propagation effect. The *Rg*-to-*Sg* ratios we used are obtained by dividing the corrected $max\{Rg_{[0.5-2Hz]}\}$ by the corrected $max\{Sg_{[0.5-8Hz]}\}$ values. Similar to the *Pg*-to-*Sg* ratios discussed in the previous paragraph, the estimated *Rg*-to-*Sg* ratios show no obvious trend with respect to magnitude. The procedure for estimating *Rg*-to-*Sg* spectral amplitude ratios is described in detail by Tibi *et al.* (2018a). The only change compared to that earlier study is that—to improve the SNR of the filtered waveforms—we increased the low-cut frequency of the applied band-pass filter from 0.1 to 0.5 Hz. Like the *Pg*-to-*Sg* ratios, the *Rg*-to-*Sg* ratio reported for each event is obtained by averaging over three or more recording stations.

Multivariate Quadratic Discriminant Analysis

The multivariate quadratic discriminant analysis followed the formalism described by Tibi *et al.* (2018b), which was adopted from Johnson and Wichern (2007). For a dataset containing *n* event populations, let $f_j(\boldsymbol{r})$ be the probability density of the measured amplitude ratios for the *j*-th population, where $\boldsymbol{r}$ is the *k*-dimensional column vector of the amplitude ratios. The characteristics of $f_j(\boldsymbol{r})$ are inferred from a subset of the population, which is referred to as learning set. Under the assumptions that (1) the sample ratios are normally distributed; (2) the

populations in the event dataset have equal prior probabilities of occurrence ($\pi_1 = \pi_2 = \ldots = \pi_n = 1/n$); and (3) the *n* event types have equal misclassification costs (the cost for allocating an event to *l*-th population, when, in fact, it belongs to the *j*-th population, $c(l|j) = 1$ for $l \neq j$ and $c(j|j) = 0$), the quadratic discriminant score for the *j*-th population, $D_j(\boldsymbol{r})$, is defined as:

$$D_j(\boldsymbol{r}) = -\frac{1}{2}\left(\boldsymbol{r} - \boldsymbol{\mu_j}\right)^{\boldsymbol{T}} \boldsymbol{S_j}^{-1} \left(\boldsymbol{r} - \boldsymbol{\mu_j}\right) - \frac{1}{2}\ln\left|\boldsymbol{S_j}\right| + \ln\left(\frac{1}{n}\right), \tag{6}$$

where $\boldsymbol{\mu_j}$ and $\boldsymbol{S_j}$ are the ratio vector mean and ratio covariance matrix of the learning events for the *j*-th population, respectively. An event of interest with the ratio vector $\boldsymbol{x}$ is allocated to the *l*-th population if $D_l(\boldsymbol{x})$ is the largest of the score values $D_1(\boldsymbol{x})$, $D_2(\boldsymbol{x})$, …, $D_n(\boldsymbol{x})$.

For a scenario involving only two populations *p* and *q* (*n* = 2, i.e., the case of two-category or binary classification), we define the score difference $d_{pq}(\boldsymbol{x})$ as:

$$d_{pq}(\boldsymbol{x}) = D_p(\boldsymbol{x}) - D_q(\boldsymbol{x}). \tag{7}$$

An event of interest is classified as *p*-type if $d_{pq}(\boldsymbol{x})$ is positive, and as *q*-type if $d_{pq}(\boldsymbol{x})$ is negative, with $d_{pq}(\boldsymbol{x}) = 0$ representing the classification line.

For each population investigated in this study, we used a learning set of 30 randomly selected events to infer the QDFs. Because the UUEB dataset contains only 27 SMBs, the learning set for that population was taken from the dataset of 142 events analyzed by Tibi *et al.* (2018b). The ratio vector $\boldsymbol{r}$ is 5-dimensional, with the first four coordinates of that vector representing *Pg*-to-*Sg* ratios for four discrete frequencies, and the fifth coordinate representing *Rg*-to-*Sg* ratios obtained for the frequency bands discussed in "Phase and Spectral Amplitude Ratio Estimation" section. In contrast to an earlier work using the same approach, in which the selected frequencies for the *Pg*-to-*Sg* ratios consisted simply of the consecutive values that provided the largest average separation between the populations investigated (Tibi *et al.*, 2018b), in the current study, the

optimal frequency set for *Pg*-to-*Sg* ratios in each scenario was chosen based on a trial-and-error process involving 1,365 combinations of four different frequencies, and corresponds to the set associated with the lowest total misclassification rate for the populations under consideration. Hence, as described later, different frequency sets were chosen for the three-category classification and each of the three cases of binary classification. This allows for the capacity of the designed discriminants to be fully exploited. We jointly used *Pg*-to-*Sg* and *Rg*-to-*Sg* ratios because a previous study by Tibi *et al.* (2018b) and test analyses with the current dataset have both shown that combining the two ratio types results in improved discrimination power for the QDF approach compared to any of the ratio types alone.

Spectrogram-Based Machine Learning Approach

The classification models for the ML approach followed the waveform processing and CNN architecture from Linville *et al.* (2019b). For model input, we used 90-second long spectrograms (starting 10 seconds before the predicted first *P* arrival time) from three-component (when available) or single-component stations, with zero-filled horizontal channels for single-component stations. The spectrograms were generated for the frequency range from 1 to 20 Hz after filtering the raw seismograms with a 1-Hz high pass filter. The input spectrograms are passed through 4 layers of (2×2) learned filters (filter counts: 18, 36, 54, 54) to a three-node softmax activated output. The output for each sample represents the probability for each class, and the event class is determined by summing the probabilities for each class over all the stations used. The classification is then based on the largest summation value.

The model training relied on both the UUEB and UUSS event catalogs and excluded the subset of events analyzed with the AR method to allow for an independent test set shared

between the methods. The UUSS catalog that covers a 5-year period (Linville *et al.*, 2019b) was used in conjunction with the UUEB listing to obtain a sufficient number of labeled events of each population. The UUSS catalog contained abundant representation of surface mining blasts, while the UUEB catalog contained a large number of MIEs, which are comparatively underrepresented in UUSS catalog. The use of both catalogs provided a large set of labeled events, which is generally required for effective model building with deep learning. In the original UUSS catalog, events were assigned to either the earthquake or mining blast class by UUSS analysts. We follow the category assignment procedure described in "Data" section and create a MIE class in the UUSS catalog. Events were then randomly assigned to 10 event partitions and models were trained using 10-fold cross validation (10% validation, 10% test) with early stopping when the validation accuracy ceased improving over 8 epochs. Since the subset of events investigated by the AR method did not participate in any model training, classes for these events were predicted using 10 different models and the standard deviation between model predictions guided our understanding of the variation in accuracy for each event using the ML method. For events that are not part of that subset (this constitutes the bulk of the UUEB), the class for each event is predicted with one model (the model in which the event under consideration did not go into the training set); and we report the average accuracies and associated standard deviations over all the events.

Results and Discussions

Depth Discrimination Using Rg-to-Sg Ratios

In our dataset, the depth for the TEs ranges from near the surface to about 10 km. We leverage the known dependence of *Rg* amplitudes on event depth by using *Rg*-to-*Sg* ratios to discriminate deeper TEs from shallow events (e.g., Båth, 1975; Kafka, 1990), with the latter

consisting of shallow TEs and anthropogenic events. Figure 3 shows the likelihood of equality (LE) between the median *Rg*-to-*Sg* ratios of shallow and deeper TEs for different depths of separation. For a depth of separation d_{sep} (in km) the population of shallow TEs consists of TEs with depth $\leq d_{sep}$, and the population of deeper TEs consists of TEs with depth $> d_{sep}$. The LE values reported in Figure 3 were estimated using Mood's statistics (Moods, 1950). For the depths of separation ≤ 6 km, the LE values are under 2%, implying that for any of those depths of separation shallow and deeper TEs are statistically distinct. Conversely, for any depth of separation > 6 km, the populations of shallow and deeper TEs are only marginally different or even the same, as the LE values in this case are no longer negligible. In other words, the distinction between the two groups in terms of *Rg*-to-*Sg* ratios breaks down for the depths of separation > 6 km.

The distribution of *Rg*-to-*Sg* ratios for the populations of MIEs, SMBs, shallow TEs (STEs), and deeper TEs (DTEs) are displayed in Figure 4. The subdivision of TEs into STEs and DTEs along the depth of 5 km is based on the results discussed in the previous paragraph. Also shown in Figure 4 is the median value of the ratios and the corresponding one-median absolute deviation (1-MAD) for each population. In particular, the histograms for the two populations of TEs each exhibit broader tails, resulting in higher 1-MAD values. Nevertheless, for the six population pairs, Mood's median tests indicate probabilities of nearly zero that the median *Rg*-to-*Sg* ratios are the same. This suggests that for each population pair, there is a statistically significant difference between the populations. These results provide further argument for the notion that using local distance *Rg*-to-*Sg* spectral amplitude ratios one can not only discriminate shallow from deeper events, but is able to discriminate among different populations of shallow events (Tibi *et al.*, 2018a). This finding is also consistent with the results reported by Kolaj (2018), according to which

low-frequency *Sg*-to-*Rg* ratios successfully discriminate low-magnitude shallow earthquakes and road construction blasts recorded at distances less than 50 km in New Brunswick, Canada.

Event Classification Using the AR Method

The average *Pg*-to-*Sg* ratios for the frequencies of 1 to 15 Hz in 1-Hz increment for the populations investigated are shown in Figure 5. The two-standard error for each measurement is also indicated. The three populations are well separated from one another in the frequency range from 8 to 15 Hz, but are less so for frequencies below 8 Hz. To illustrate this observation further, we chose 12 and 5 Hz as representative values for the two frequency ranges, respectively, and generated histograms of the *Pg*-to-*Sg* ratios calculated for these frequencies (Fig. S3). Mood's tests suggest that at 12 Hz the populations are all statistically distinct; and at 5 Hz, the SMBs are statistically different from the TEs, and the MIEs from the TEs. At that frequency, there is, however, a non-negligible likelihood of ~7% that the SMB and MIE populations are the same.

To improve the discrimination power of the AR method, *Pg*-to-*Sg* ratios for four discrete frequencies are exploited jointly with the *Rg*-to-*Sg* ratios in the QDFs, according to the formalism described above in "Multivariate Quadratic Discriminant Analysis" section. Results for the two-category, pairwise classification for the three possible population pairs are summarized in Figure 6. The QDFs were tailored for each pair; and the optimum set of frequencies for the *Pg*-to-*Sg* ratios were selected for each pair using a trial-and-error process as described earlier. With a misclassification rate (error rate, ER) of only 1.6%, the QDF approach performs best for the population pair MIE/SMB, followed by the pair SMB/TE (ER = 6%). For the former and latter population pairs, we used the frequency sets for *Pg*-to-*Sg* ratios of 1, 2, 10 and 15 Hz and 1, 2, 6, and 8 Hz, respectively. For the population pair MIE/TE, the best-performing frequency set for *Pg*-

to-*Sg* ratios of 6, 8, 14, and 15 Hz is shifted more toward higher frequencies. In terms of the exploited ratios, the two populations from that pair share some characteristics, which is reflected in the relatively high ER value of 14% (Fig. 6). This is particularly the case between the MIEs and the shallow (depth $\leq$ 5 km) TEs, as implied by the fact that the ER value drops to 4.5% if only deeper TEs are considered in this binary classification (Fig. 6). In which case, only four TEs, all from the Circleville sequence, are misclassified as MIE.

Results of three-category classification, in which all the three populations are considered simultaneously, are summarized in Figure 7. Since the number of events varies greatly among the populations, the confusion matrix reported in Table 1 represents a better measure of the overall performance of the classifier than the error rate alone. For the three-category classification, the optimum frequency set for *Pg*-to-*Sg* ratios was 1, 2, 7, and 8 Hz. On average, about 85% of the events in the dataset are classified correctly. The population of SMBs is associated with the lowest error rate (ER = 11.1%) followed by the MIEs (14.6%). The TE group exhibits the highest error rate (19.4%) of any population. The majority of the misclassified TEs are assigned to the MIE group. Similarly, most of the misclassified MIEs are assigned to the TE population. This is consistent with the outcome of the two-category classification discussed above, and the results from an earlier study by Tibi *et al.* (2018a), which suggests that the two populations have some properties in common, which make it difficult to discriminate between them using the AR approach. The apparent strong commonality between the TEs and the MIEs in terms of the investigated ratios is intriguing. In general, the source mechanisms for the TEs consists predominantly of shear dislocation (i.e., negligible non-double couple components). In contrast, MIE sources should include significant non-double couple components, consistent with the events being associated

with pillar or roof collapses from long-wall mining (Arabasz *et al.* 2005; Whidden and Pankow, 2012). As discussed in detail by Tibi *et al.* (2018a), the difference in source mechanisms between the two event types results in differing excitations for the phases under consideration. Furthermore, compared to the TEs, signals from the MIEs are generally depleted in high-frequency energy (Tibi *et al.*, 2018a). Curiously, despite these differences the discriminant is relatively less efficient than between other groups.

Some of the TEs that occurred in the coal mining district of central Utah overlap geographically with the MIEs (Fig. 2, right). These TEs were labeled as such only because the reported depths (> 3 km) are deemed too large for anthropogenic events. We investigated the possibility that some of these events might be contributing to the high misclassification rate for the TE population because they might actually be MIEs that were mis-labeled due to the inaccuracies of the reported depths. Only two TEs located in the coal mining region are misclassified, and they are all assigned to the SMB population. This observation, combined with the outcome of similar investigations carried out for the SMBs and the MIEs, suggest that mislabeling of the events in the UUEB dataset is not a major contributing factor for the misclassification of any of the groups. The majority of the misclassified TEs are lower magnitude (M_C < 2.5) events, suggesting that the performance of the AR method deteriorate at low SNRs. Many of these events are from the Circleville sequence, which is also the largest subgroup in that population.

Comparison of the AR with the ML Method

The comparison between the AR and the ML method is based on the subset of events investigated in “Event Classification Using the AR Method” section. However, with the ML

approach only 120 TEs were investigated, four fewer than with the AR method. Because the IRIS data fetch failed for those four events. To emulate the ML approach that uses 10 different training models for class prediction (see “Methods” section), the classification in this case was performed 10 times with the AR method. Each time, the learning events used to construct the discriminants functions for each of the three populations (see “Multivariate Quadratic Discriminant Analysis” section) were randomly selected. This allows us to assess the performance variability associated with the random selection of events for the learning sets.

Classification results for the ML method are summarized in Figure 8; and the associated confusion matrix is reported in Table 2. The reported values represent averages over the 10 prediction runs. These results are contrasted in Figure 9 (left) against those obtained from the AR method. For each classification method and population, the standard deviations estimated from the 10 prediction runs are indicated. The high variability in classification success rates (accuracy) for the SMB group is mainly the consequence of the limited number of events in that group. For that population, within the margin of errors, the performance of the AR method in terms of success rates is comparable with that of the ML approach. For the TE and MIE groups, however, the latter method outperforms the former by an average of about 14% (Fig. 9, left). For each population, we also assessed the precision achieved by each classification method. The ML approach shows precisions of 95–97%, which are 16% on average higher than those from the AR method.

Using the ML approach, we were able to classify 7,350 of the 7,377 events from the UUEB dataset, including the abundant, extremely low-magnitude events. For the ML models, however, the SMBs, WMBs, LMBs, and CQBs were all grouped together to form a class referred to as

"mining blast" or MB. The results summarized in Figure 9 (right) and Table 3 indicate success rates of about 94–100%, with the highest rate associated with the population of MIEs. The standard deviation for the prediction accuracy of the MB population is high (±11.8%). For that population, eight of the ten cross-validation folds show accuracies of 100%, while the accuracies of about 67–75% for the two remaining cross-validation folds are significantly lower. The low accuracies for these folds may be misleading; there is only one misclassified event in each case, implying that the apparent high variability in accuracy is due to the limited number of events in the MB group.

Like any other approach that relies on measurements of signal amplitudes, the AR method would fail if for any reason those amplitudes can no longer be estimated accurately. This would be the case for the multitude of extremely low-magnitude events in the UUEB dataset due to the low-SNR nature of their signals. The purpose of this work was to compare/contrast the performances of the AR and ML methods in terms of their discrimination power at local distances, using the same set of well-recorded low-magnitude events. The choice of the minimum cutoff magnitudes to ensure high quality data for the TE and MIE populations in the subset investigated by the AR method was guided simply by previous experience (Tibi et al., 2018a,b). Thus, the limitations of the AR method with respect to the event size were not explicitly investigated. Nevertheless, for the TE population, we noticed a deterioration of the performance of the method at low SNRs. The fact that the ML approach can classify our dataset in almost its entirety (7,350 events out of 7,377), including extremely low-magnitude events, with the reported high success rates, suggests that—compared with the AR method—the ML approach is more robust to low SNR data and is therefore an attractive approach, given the right circumstances. One potential drawback to the ML approach is that it is completely data-driven, not relying on any knowledge of seismic phenomena

(sources types, energy partitioning into phases, signal propagation, etc.). Another limitation of the ML method is that models that generalize well to new sources generally require large training sets. Hence, building a robust ML model for event classification requires datasets with 100s or even 1000s of meticulously labeled events, and such large datasets of labeled events are not always available, restricting the application of the approach only to cases for which such datasets exist. It is of course possible to apply an ML model that has been developed in an area with abundant labeled data to another area, and some promising results have been shown doing this (e.g., Ross *et al.*, 2018b), but this has high risks given our lack of knowledge about exactly how the ML method works. Another way to deal with this problem is by applying semi-supervised learning methods which require much less labeled data (e.g., Bergen *et al.*, 2019; Chapelle *et al.*, 2006 for a review). In terms of computational demands for real-time event classification, the AR and ML methods are comparable and either would only require modest computational resources. Developing ML models can be computationally expensive, of course, but that is an up-front expense.

Conclusions

Using the same set of tectonic and anthropogenic events in the Utah region that were well-recorded at local distances by the UUSS network, we compared two methods of classifying the source types of the events: a traditional method, in which amplitude ratios of different phases are exploited in the QDFs, and a spectrogram-based machine learning approach. A summary of the observed advantages and limitations of each of the two methods is provided in Table 4. In the AR method, the QDFs are tailored to achieve maximum separation between the populations under consideration. Also, the set of frequencies for the *Pg*-to-*Sg* ratios exploited in the QDFs are

selected using a trial-and-error procedure and consists of values that minimize the misclassification rates. With this method, we achieved classification success rates of about 80–90%. The ML approach used convolutional neural network models based on 90-second spectrograms to classify the populations and achieves success rates above 90%. Importantly, this method is more robust to low SNR data, allowing it to classify extremely low-magnitude events, and hence classify a much larger portion of the Utah data set (99.6% compared by 5.0% for AR). We suggest that the complex deep learning architecture involved in the ML method is able to expose and exploit characteristics that are specific to each event population to the level that the traditional AR approach cannot, allowing ML to achieve higher accuracies and precisions, and classify significantly smaller events. An advantage of the AR method, is that the physical basis is well understood (i.e., energy partitioning between specific phases). Conversely, the lack of insight into the physical basis for ML classifications has traditionally been a major reason why seismic event monitoring agencies have been reluctant to use ML methods, and make the transportability of models developed from one region to a different region problematic. We agree with Kong *et al.* (2019) that the best way to take advantage of the strengths of each of the two methods would be to design a hybrid approach that combines the traditional physical modeling from the AR method with the data-driven ML approach. This is a research topic we intend to focus on in the near future.

Data and Resources

Waveform data were retrieved from IRIS database (http://www.iris.edu, last accessed April 2019). The coda duration magnitude (M_C) was calculated using the equation for the Utah region proposed in the unpublished manuscript *Correction of Systematic Time-Dependent Coda*

Magnitude Errors in the Utah and Yellowstone National Park Region Earthquake Catalogs, 1981–2001 by J. C. Pechmann, J. C. Bernier, S. J. Nava, and F. M. Terra.

Acknowledgments

The authors acknowledge the support of the National Nuclear Security Administration Office of Defense Nuclear Nonproliferation Research and Development for funding this work. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. 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 Tim Draelos was reading the manuscript.

References

Arabasz, W. J., S. J. Nava, M. K. McCarter, K. L. Pankow, J. C. Pechmann, J. Ake, and A. McGarr (2005). Coal-mining seismicity and ground-shaking hazard: A case study in t, he Trail Mountain area, Emery County, Utah, *Bull. Seismol. Soc. Am.* **95,** 18–30, doi: 10.1785/0120040045.

Arrowsmith, S. J., R. Burlacu, K. Pankow, B. Stump, R. Stead, R. Whitaker, and C. Hayward (2012). A seismoacoustic study of the 2011 January 3 Circleville earthquake. *Geophys. J. Int.* **189,** 1148–1158.

Båth, M. (1975). Short-period Rayleigh waves from near-surface events, *Phys. Earth. Planet. Inter.* **10,** 369–376.

Benbrahim, M., A. Daoudi, K. Benjelloun, A. Ibenbrahim (2007). Discrimination of seismic signals using artificial neural networks, *Int. J. Comput. Inform. Eng.,* **1,** no.4, 984–987.

Bergen, K. J., P. A. Johnson, M. V. de Hoop, and G. C. Beroza (2019). Machine learning for data-driven discovery in solid Earth geoscience, *Science* **363,** doi: 10.1126/science.aau0323.

Bowers, D., and N. Selby (2009). Forensic seismology and the Comprehensive Nuclear-Test-Ban Treaty, *Ann. Rev. Earth Planet. Sci.* **37,** 209–236.

Bratt, S.R., T. C. Bache (1988). Locating events with a sparse network of regional arrays, *Bull. Seismol. Soc. Am.* **78,** 780–798.

Chapelle, O., B. Schölkopf, and A. Zien (Editors) (2006). *Semi-Supervised Learning*, The MIT Press, Cambridge, Massachusetts, 508 pp.

Del Pezzo, E., A. Esposito, F. Giudicepietro, M. Marinaro, M. Martini, and S. Scarpetta (2003). Discrimination of earthquakes and underwater explosions using neural networks, *Bull. Seismol. Soc. Am.* **93,** no. 1, 215–223.

DeVries, P. M. R., F. Viégas, M. Wattenberg, and B. J. Meade (2018). Deep learning of aftershock patterns following large earthquakes, *Nature* **560,** no. 7720, 632–634.

Dowla, F. U., S. R. Taylor, and R. W. Anderson (1990). Seismic discrimination with artificial neural networks: Preliminary results with regional spectral data, *Bull. Seismol. Soc. Am.* ***80,*** *no.*5, 1346–1373.

Johnson, R. A., and D. W. Wichern (2007). *Applied Multivariate Statistical Analysis*, Pearson Prentice Hall, Upper Saddle River, New Jersey, 773 pp.

Jordan, M. I., and T. M. Mitchell (2015). Machine learning: Trends, perspectives, and prospects, *Science* **349,** 255–260.

Kafka, A. L. (1990). *Rg* as depth discriminant for earthquakes and explosions: A case study in New England, *Bull. Seismol. Soc. Am.* **80,** no. 2, 373–394.

Kennett, B. L. N., E. R. Engdahl, and R. Buland (1995). Constraints on seismic velocities in the Earth from traveltimes, *Geophys. J. Int.* **122,** 108–124, doi: 10.1111/j.1365-246X.1995.tb03540.x.

Kim, W.-Y., D. W. Simpson, and P. G. Richards (1993). Discrimination of earthquakes and explosions in the eastern United States using regional high-frequency data, *Geophys. Res. Lett.* **20,** 1507–1510.

Kolaj, M. (2018). Discrimination between low-magnitude shallow earthquakes and road construction blasts near Big Salmon River, New Brunswick, Canada, *Seismol. Res. Lett.* **89,** 1966–1976, doi: 10.1785/0220180077.

Kong, Q., D. T. Trugman, Z. E. Ross, M. J. Bianco, B. J. Meade, and P. Gerstoft (2019). Machine learning in seismology: Turning data into insights, *Seismol. Res. Lett.* **90,** 3–14, doi: 10.1785/0220180259.

Koper, K. D., J. C. Pechmann, R. Burlacu, K. L. Pankow, J. Stein, J. M. Hale, P. Roberson, and M. K. McCarter (2016). Magnitude-based discrimination of man-made seismic events from naturally occurring earthquakes in Utah, USA, *Geophys. Res. Lett.* **43**, 10638-10645, doi: 10.1002/2016GL070742.

Kriegerowski, M., G. M. Petersen, H. Vasyura-Bathke, and M. Ohrnberger (2019). A deep convolutional neural network for localization of clustered earthquakes based on multistation full waveforms, *Seismol. Res. Lett.*, doi: 10.1785/0220180320.

Linville, L., R. Brogan, C. Young, and K. A. Aur (2019a). Global to local high-resolution event catalogs for algorithm testing and source studies, Submitted to *Seismol. Res. Lett.*

Linville, L., K. L. Pankow, T. Draelos (2019b). Deep learning models augment analyst decisions for event discrimination, *Geophys. Res. Lett.* **46,** doi: 10.1029/2018GL081119.

Mood, A. M. (1950). *Introduction to Theory of Statistics*, McGraw-Hill, New York, New York, 394–399.

Mousavi, M. S., S. P. Horton, C. A. Langston, and B. Samei (2016). Seismic features and automatic discrimination of deep and shallow induced-microearthqakes using neural network and logistic regression, *Geophys. J. Int.* **207,** 29–46.

O'Rourke, C. T., G. E. Baker, and A. F. Sheehan (2016). Using P/S amplitude ratios for seismic discrimination at local distances, *Bull. Seismol. Soc. Am.* **106,** 2320–2331.

Pankow, K. L., J. R. Moore, J. M. Hale, K. D. Koper, T. M. Kubacki, K. M. Whidden, and M. K. McCarter (2014). Massive landslide at Utah copper mine generates wealth of geophysical data, *GSA Today* **24,** 4–9, doi: 10.1130/GSATG191A.1.

Perol, T., M. Gharbi, and M. Denolle (2018). Convolutional neural network for earthquake detection and location, *Sci. Adv.* **4,** no. 2, e1700578.

Ross, Z. E., M.-A. Meier, and E. Hauksson (2018a). P-wave arrival picking and first-motion polarity determination with deep learning, *J. Geophys. Res*. **123,** 5120–5129.

Ross, Z. E., M.-A. Meier, E. Hauksson, and T. H. Heaton (2018b). Generalized seismic phase detection with deep learning, *Bull. Seismol. Soc. Am.* doi: 10.1785/0120180080.

Smith, R. B., and W. J. Arabasz (1991). Seismicity of the intermountain seismic belt, in *Neotectonics North America*, D. B. Slemmons, E. R. Engdahl, M. D. Zoback, M. L. Zoback, and D. Backwell (Editors), Vol. 1, Geological Society of North America, 185–228.

Stump, B. W., M. A. H. Hedlin, D. C. Pearson, and V. Hsu (2002). Characterization of mining explosions at regional distances, *Rev. Geophys.* **40,** 1011, doi: 1998RG000048.

Taylor, S. R. (1996). Analysis of high-frequency Pg/Lg ratios from NTS explosions and Western U.S. Earthquakes, *Bull. Seismol. Soc. Am.* **86,** no. 4, 1042–1053.

Tibi, R., K. D. Koper, K. L. Pankow, and C. J. Young (2018a). Depth discrimination using *Rg*-to-*Sg* spectral amplitude ratios for seismic events in Utah recorded at local distances, *Bull. Seismol. Soc. Am.* **108,** 1355–1368, doi: 10.1785/0120170257.

Tibi, R., K. D. Koper, K. L. Pankow, and C. J. Young (2018b). Discrimination of anthropogenic events and tectonic earthquakes in Utah using a quadratic discriminant function approach with local distance amplitude ratios, *Bull. Seismol. Soc. Am.* **108,** 2788–2800, doi: 10.1785/0120180024.

Walter, W. R., K. M. Myeda, and H. J. Patton (1995). Phase and spectral ratio discrimination between NTS earthquakes and explosions. Part I: Empirical observations, *Bull. Seismol. Soc. Am.* **85,** 1050–1067.

Whidden, K. M., and K. L. Pankow (2012). A catalog of regional moment tensors in Utah from 1998 to 2011*, Seismol. Res. Lett.* **83,** doi: 10.1785/0220120046.

Zeiler, C., and A. A. Velasco (2009). Developing local to near-regional explosion and earthquake discriminants, *Bull. Seismol. Soc. Am.* **99,** 24–35, doi: 10.1785/0120080045.

Full mailing addresses

Rigobert Tibi

Sandia National Laboratories

P.O. Box 5800

Albuquerque, NM 87185-0404

Lisa Linville

Sandia National Laboratories

P.O. Box 5800

Albuquerque, NM 87185-0758

Christopher J. Young

Sandia National Laboratories

P.O. Box 5800

Albuquerque, NM 87185-0401

Ronald Brogan

ENSCO, Inc.

5400 Port Royal Road

Springfield, VA 22151

Tables

Table 1. Confusion matrix for the AR method and the subset of 370 events.

		Actual Class		
		SMB	**MIE**	**TE**
Predicted Class	**SMB**	24	3	8
	MIE	2	187	16
	TE	1	29	100

For each population, the diagonal elements in the gray cells indicate the number of events classified correctly (true positives), and the off-diagonal elements the number of misclassified events (false negatives and false positives).

Table 2. Confusion matrix for ML approach and the subset of 370 events.

		Actual Class		
		SMB	**MIE**	**TE***
Predicted Class	**SMB**	26	1	2
	MIE	1	214	2
	TE	0	4	116

Description of the elements is the same as in Table 1. The reported values represent averages from the 10 prediction runs. * Note that in this case only 120 TEs were investigated, compared with 124 TEs for the AR approach (Table 1).

Table 3. Confusion matrix for the ML approach and the 7,350 events from the UUEB dataset.

		Actual Class		
		MB	**MIE**	**TE**
Predicted Class	**MB**	45	1	7
	MIE	3	6280	37
	TE	0	5	972

The same as in Table 2, but for the 7,350 events of the UUEB dataset.

Table 4: Advantages and limitations of the classification methods used in this study.

	AR Method	**ML Method**
Advantages	• Require only limited number of learning events for each category • Fast computation • Physical basis well understood (i.e., energy partitioning between specific phases)	• Achieve high accuracy (> 90%) • Can classify extremely low-magnitude events • Fast to apply the model once it has been built
Limitations	• Achieve moderate accuracy (~80–90%) • Require high-quality data	• Require large dataset (100s–1000s) of labeled training events for each category • Computationally intensive to build the model • Lack of insight into the physical basis

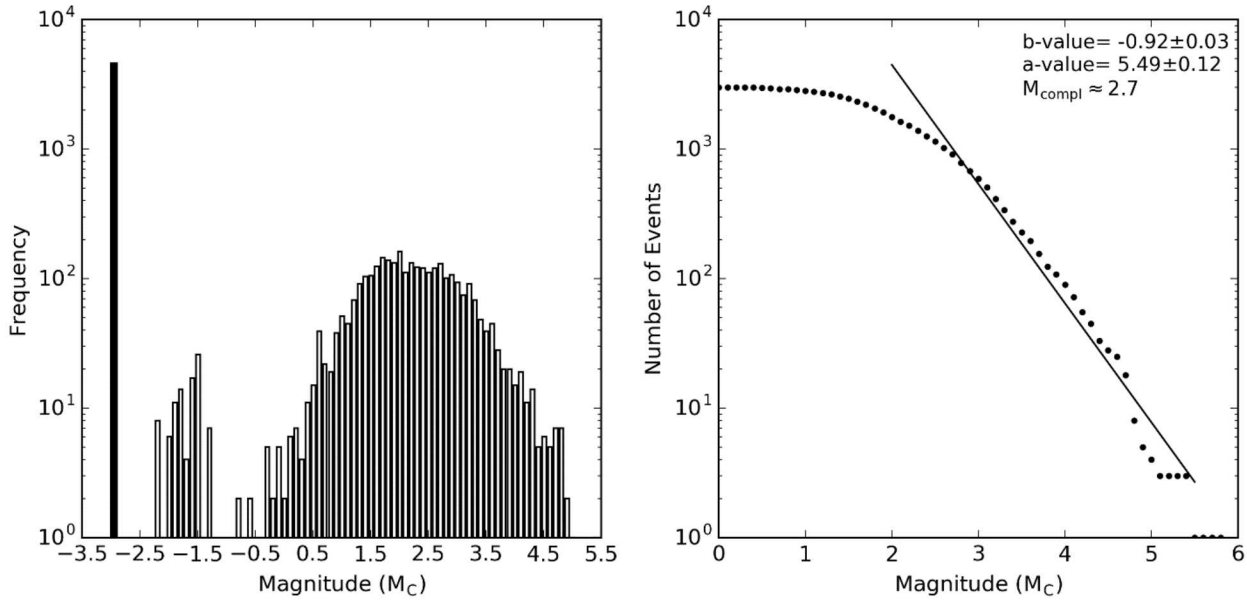

Figure 1. (Left) Magnitude (M_C) distribution for events in the UUEB catalog. The black bar at magnitude −3 represents the frequency of low SNR events for which M_C could not be estimated. (Right) Magnitude-frequency plot for the 3,118 events for which magnitude were calculated. The black line corresponds to *a*- and *b*-value of 5.49 and −0.92, respectively, obtained using a maximum-likelihood method. The magnitude of completeness, M_{compl}, is estimated to be about 2.7

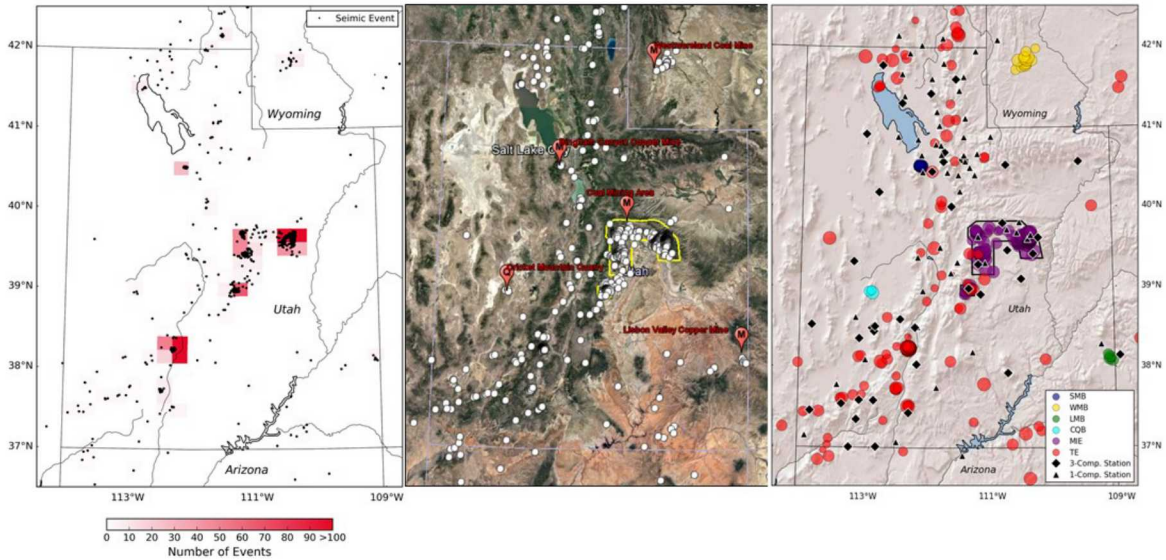

Figure 2. (Left) Heat map of the spatial density of events in the UUEB dataset. The event cluster located around 38.27°E and 112.29°W is associated with the January 2011 Circleville earthquake sequence (Arrowsmith *et al.*, 2011). Other clusters are related to mining activities. (Middle) Google map of the Utah region showing the locations of the events from the UUEB catalog (white circles). To demonstrate the relationship between some of the events and mining activities, known mining areas are also indicated. The yellow polygons outline the two regions of active coal mining (multiple mines). (Right) Relief map of the region shown in the Google map on the left. Navy circles represent shallow mining blasts from an open-pit copper mine (SMB), gold circles are events associated with the Westmoreland coal mine (WMB), green circles are blasts from the Lisbon Valley copper mine (LMB), cyan circles blasts from the Cricket Mountain quarry, purple circles are mining-induced earthquakes (MIE), and red circles are tectonic earthquakes (TE). The two regions of active coal mining are outlined with black polygons.

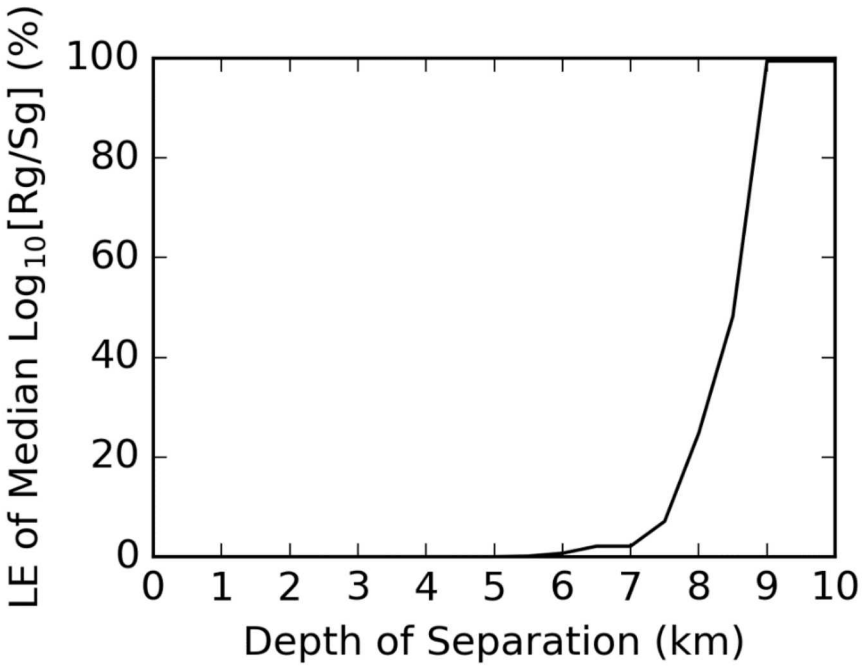

Figure 3. Likelihood of equality (LE) between the median *Rg*-to-*Sg* ratios of shallow and deeper TEs for different depths of separation. For a depth of separation d_{sep} (in km) the population of shallow TEs consists of TEs with depth $\leq d_{sep}$, and the population of deeper TEs consists of TEs with depth > d_{sep}. The LE values were estimated using Mood's statistics (Mood, 1950).

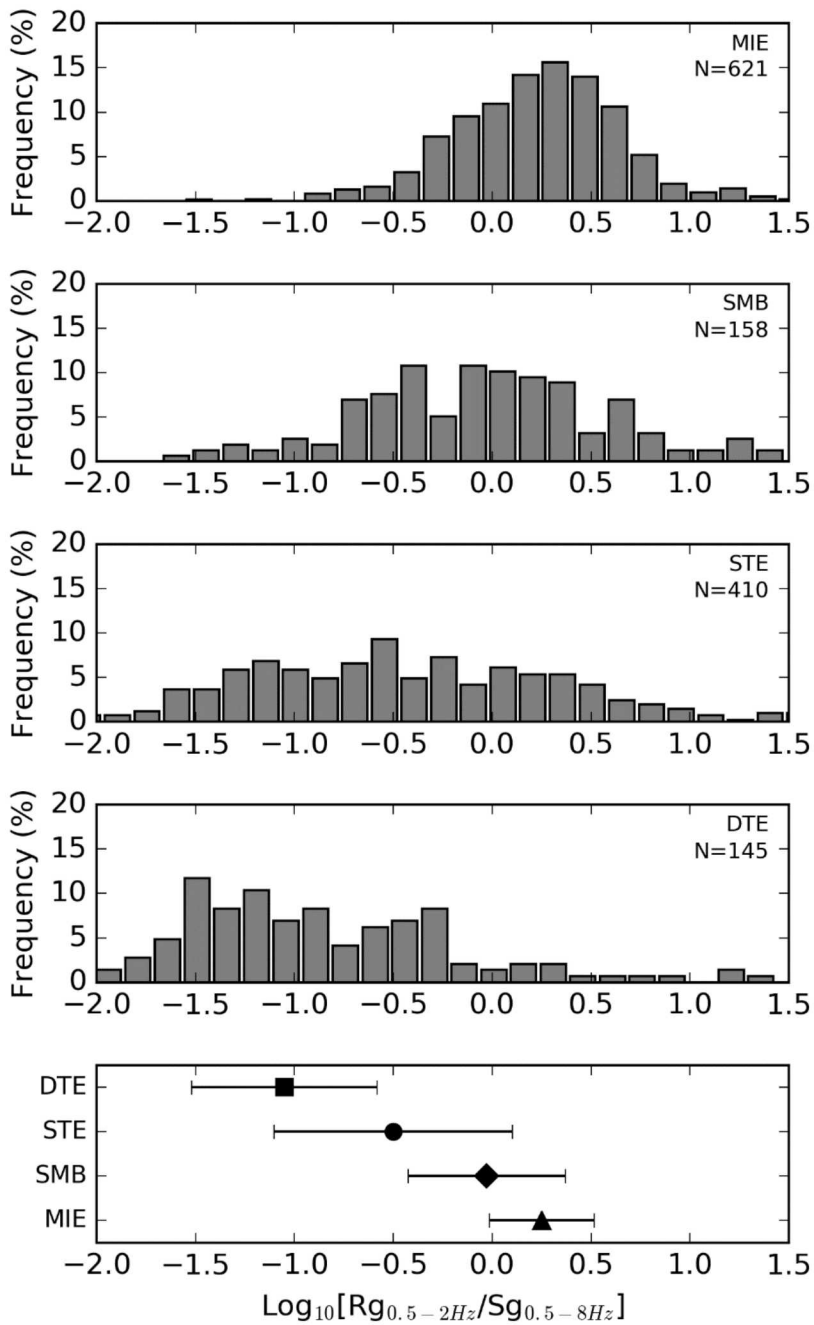

Figure 4. Histograms of the estimated *Rg*-to-*Sg* spectral amplitude ratios for the population of MIEs, SMBs, STEs (TEs with depth ≤ 5 km), and DTEs (TEs with depth > 5 km). The value of N in each plot indicates the total number of contributing single ratios for the population. For each population, the median (triangle for MIEs, diamond SMBs, circle STEs, and square DTEs) and one-median absolute deviation values of the ratios are shown in the bottom plot.

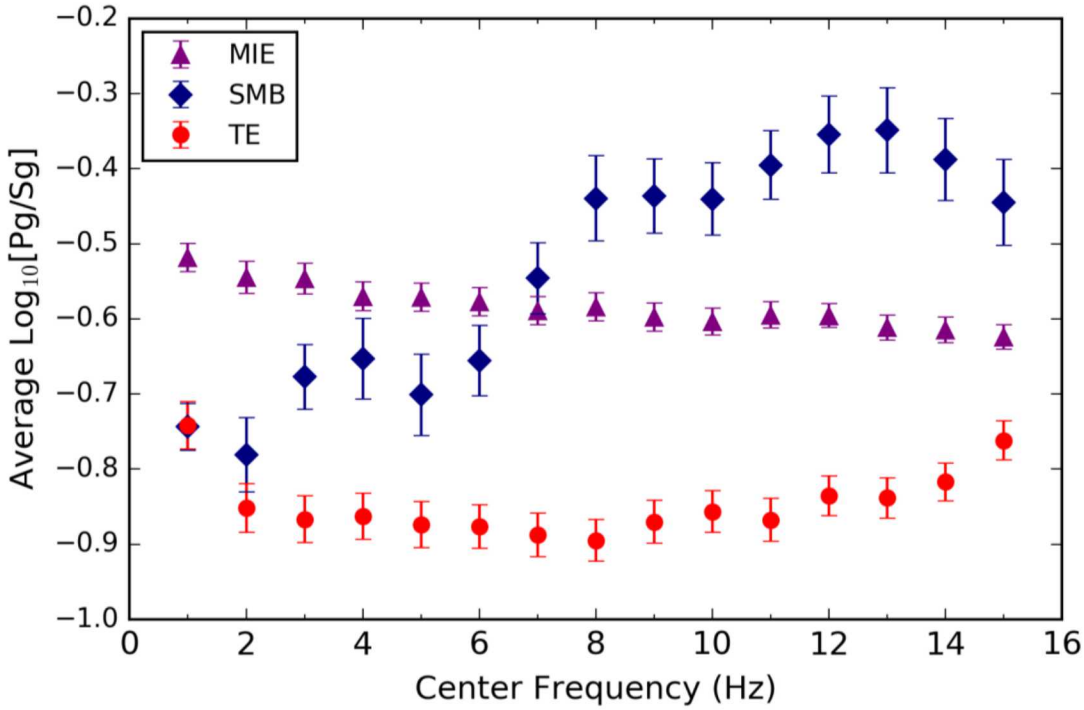

Figure 5. Average *Pg*-to-*Sg* ratios and associated two-standard errors for the discrete frequencies from 1 to 15 Hz and the population of MIEs (purple triangles), SMBs (navy diamonds), and TEs (red circles).

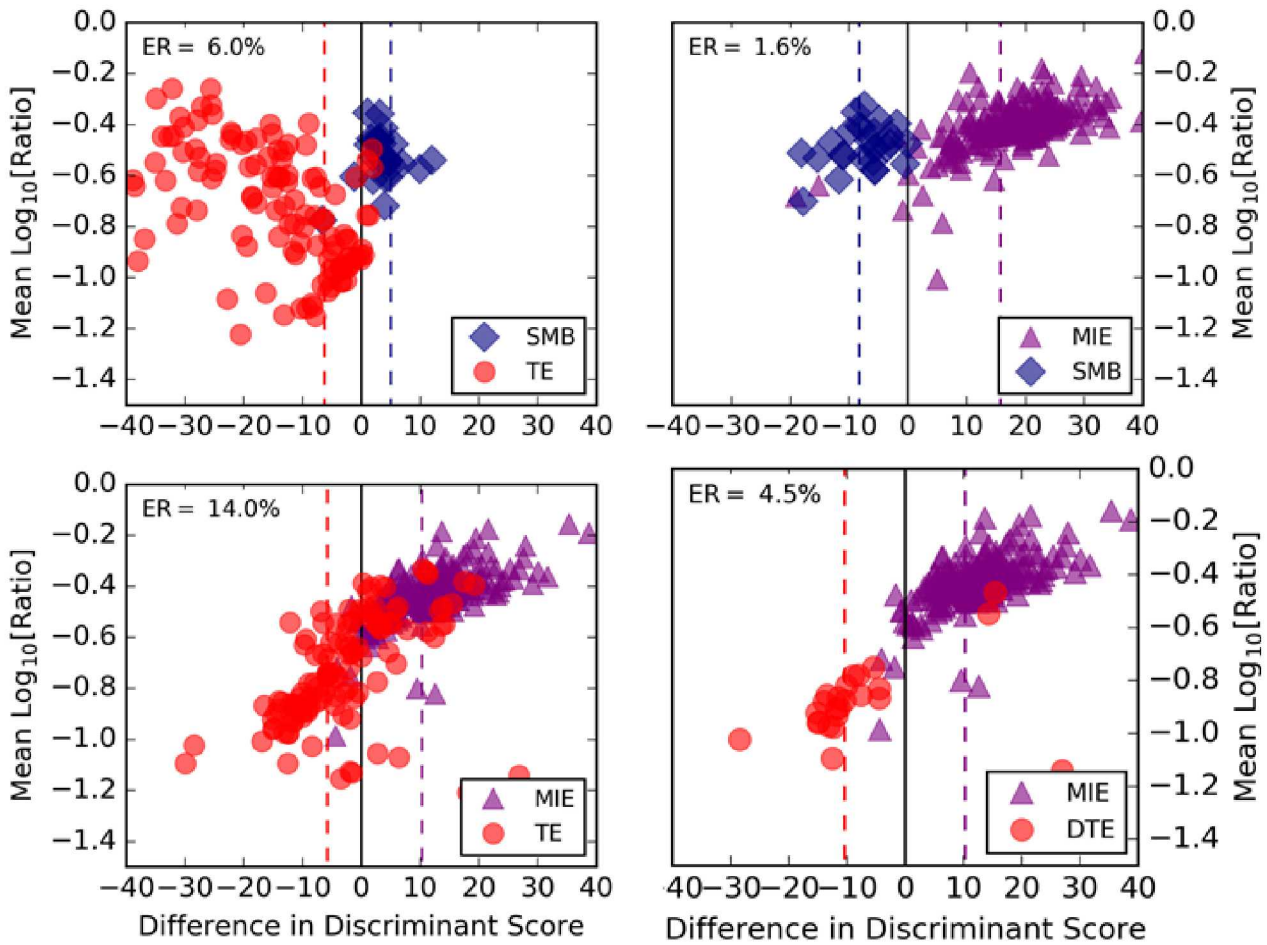

Figure 6. Difference in discriminant scores for two-category classification using both *Pg*-to-*Sg* and *Rg*-to-*Sg* ratios. The vertical axes represent the mean of the five coordinates of the ratio vectors for each event. There is one plot for each population pair. For each pair, the discriminant scores were calculated using *Pg*-to-*Sg* ratios for the optimum set of frequencies in the discriminant functions inferred for that population pair (see "Multivariate Quadratic Discriminant Analysis" section). The vertical dashed lines in each plot represent the projections of the multi-ratio means of the populations. The ER value indicates the classification error rate for each population pair. The population of DTE (deeper tectonic earthquakes) in the lower right plot consists of earthquakes with depth larger than 5 km.

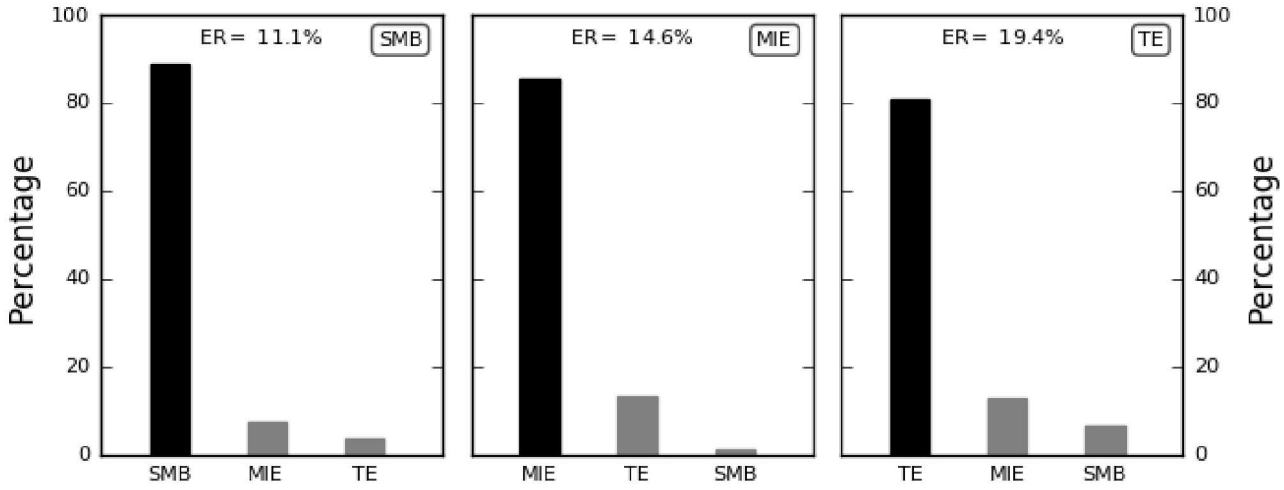

Figure 7. Bar plots of the proportion of events allocated to each group by the AR approach in three-category classification for the subset of 370 events. For each plot, the population under consideration is indicated in the upper right corner; the groups, which the events from that population are allocated to, are indicated along the horizontal axis. The black bar in each plot represents the proportion of events that are correctly classified, and the gray bars the proportion of events that are misclassified. The ER value indicates the classification error rate.

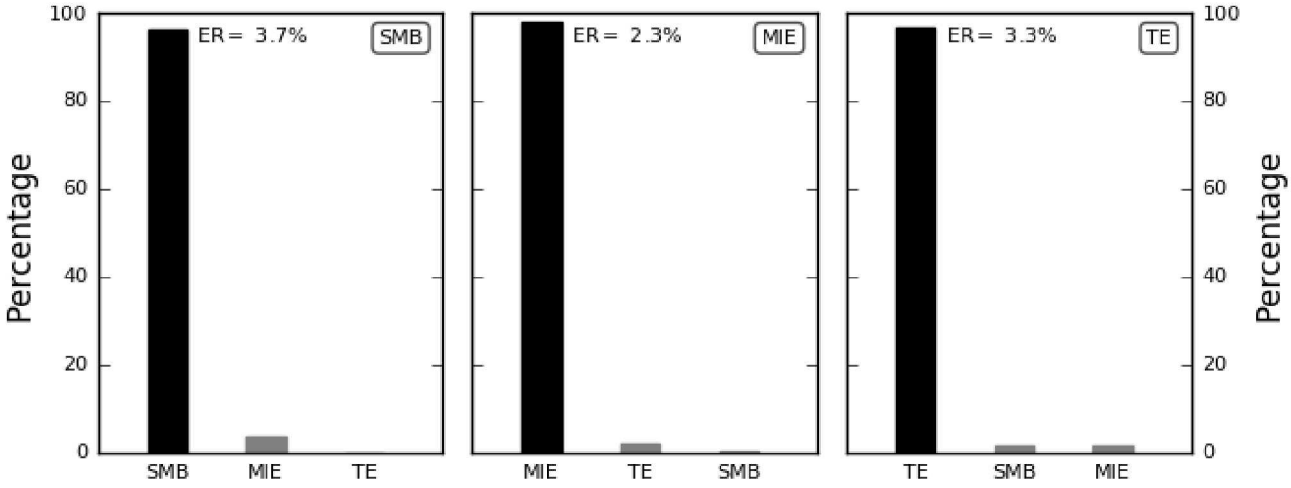

Figure 8. The same as Figure 7, but for the ML approach.

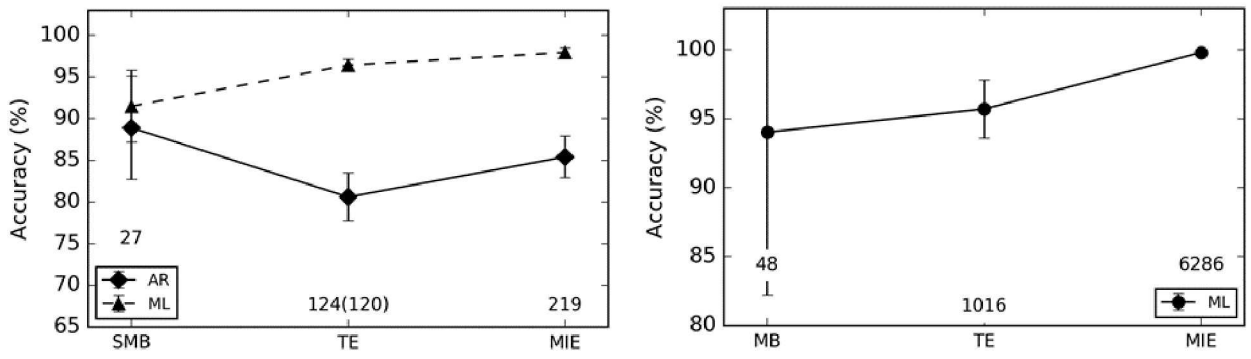

Figure 9. (Left) Classification accuracies for the AR method (diamonds) and ML approach (triangles) for the subset of 370 events. The value above each population name on the horizontal axis represents the total number of events from that population that were classified. Note that for the ML approach only 120 TEs were investigated (number in parentheses). (Right) Classification accuracies for the UUEB dataset using the ML approach. A total of 7,350 events were classified using that method.