

An Autoencoder-Based Deep Learning Model for Enhancing Noise Characterization and Microseismic Event Detection in Underground Longwall Coal Mines Using Distributed Acoustic Sensing Monitoring

Tourei, A., and Martin, E.R.

Colorado School of Mines, Golden, Colorado, USA

Ankamah, A.T., and Hole, J.A.

Virginia Tech, Blacksburg, Virginia, USA

Chambers, D.J.A.

National Institute for Occupational Safety and Health, Spokane, Washington, USA

Copyright 2024 ARMA, American Rock Mechanics Association

This paper was prepared for presentation at the 58th US Rock Mechanics/Geomechanics Symposium held in Golden, Colorado, USA, 23-26 June 2024. This paper was selected for presentation at the symposium by an ARMA Technical Program Committee based on a technical and critical review of the paper by a minimum of two technical reviewers. The material, as presented, does not necessarily reflect any position of ARMA, its officers, or members. Electronic reproduction, distribution, or storage of any part of this paper for commercial purposes without the written consent of ARMA is prohibited. Permission to reproduce in print is restricted to an abstract of not more than 200 words; illustrations may not be copied. The abstract must contain conspicuous acknowledgement of where and by whom the paper was presented.

ABSTRACT: The longwall mining method is designed to optimize coal extraction through controlled roof caving, which inevitably induces seismicity. This research employs a distributed acoustic sensing (DAS) system incorporating a fire-safe fiber-optic cable strategically installed underground within an operational longwall coal mine. Despite lower sensitivity than traditional seismometers, DAS sensing technology benefits from dense sensor spacing and close proximity to the active face, where many microseismic events occur. To automatically detect seismic events within the voluminous DAS data records, we employ convolutional autoencoder deep learning models that can be used for anomaly (potential seismic event) detection in power spectral density (PSD) images of DAS recordings. The kernel density estimation (KDE) technique is used to calculate the probability density function (PDF) for the density scores of the latent space (representation of compressed data). We then use this calculated parameter as a threshold to distinguish between the PSD associated with background noise and with potential seismic events. The DAS monitoring system in conjunction with the developed deep learning model could enhance longwall coal mining safety and efficiency by offering valuable data from its densely deployed multichannel sensors near mining operations.

1. INTRODUCTION

Longwall mining is an efficient underground mining method for extracting a variety of stratified resources including coal, potash, and soda ash and represents a considerable advancement over conventional methods (Peng, 2019). A modern longwall primarily consists of hydraulic shields that support the roof and floor, a cutting device (e.g., a shearer or plow) that travels along the face extracting slices of coal, and an armored conveyor belt, which transports the resource to a larger mine haulage system. Normally longwall mining is safe and efficient, but a variety of ground control-related hazards are possible, especially in deep mines. One of the most significant of these hazards is a class of dynamic failures associated with induced seismicity and damage to mine workings, generally referred to as coal bursts or mine bumps. Much like tectonic earthquakes, mining-induced seismicity is difficult to predict and can have devastating consequences. For example, over the past several

decades, coal bursts have killed hundreds of miners (Zhang et al., 2017).

The mechanisms and severity of coal bursts and mine bumps can vary widely, including localized failures occurring in the coal or near-seam strata, failure of competent strata in the overburden, and catastrophic chain failure of pillars, which can span large areas (Mark, 2016). Although significant advancements have been made in the past 100 years of research, many aspects of coal bursts remain “enigmatic” (Mark, 2018). There are a variety of options for managing coal bursts risk (Wei et al., 2018), but selecting and applying appropriate measures for dealing with coal bursts depends on an adequate understanding of the source, geology, and geomechanics associated with the bursts. For this, a variety of information sources are useful, including seismic monitoring. Apart from helping to address coal bursts, seismic monitoring can be useful for a variety of other safety applications in underground coal mining.

For underground coal mines, seismic monitoring is conducted using surface or in-mine sensors. Surface networks are usually less expensive and easier to maintain and install, but in-mine networks provide higher quality data in terms of event detection and location accuracy, especially event depth constraints (Swanson et al., 2016). In addition to much greater costs, regulations designed to help avoid fires and explosions in coal mines restrict the use and placement of electronics, including many seismic sensors and most digitizers.

Deploying seismic sensors underground enables the characterization of various noises and also allows detection of smaller events since the sensors are closer to where these events occur (i.e., where the signal is strongest). One example of noise characterization is the characterization of machinery noise which is important for both operational efficiency and worker safety (Peng et al., 2020). Machinery noise cannot only be a significant occupational hazard but also a critical indicator of equipment condition and operational anomalies. Accurate noise characterization helps identify impending machinery failures, enabling preventative maintenance and a reduction in downtime. Furthermore, it is crucial to differentiate between mechanical noise and seismic events to get a proper understanding of the rockmass response to mining.

One promising technology to improve in-mine seismic monitoring of underground coal mines is distributed acoustic sensing (DAS) (Ankamah et al., 2023; Chambers & Shragge, 2023; Wang et al., 2018; Zhang et al., 2017). A DAS system is composed of an interrogator unit with optical and electronic components plugged into a fiber-optic cable. The interrogator probes the cable with light to measure a vibration time series at each position along the cable. Unlike most traditional seismic systems, MSHA-approved optical fibers pose no risk of causing a fire and so can be placed anywhere in coal mines. Another challenge encountered by underground networks of traditional seismic sensors or nodes is the need to maintain precise time synchronization underground, but the channels (i.e. sensing locations) along a DAS fiber are all automatically synchronized throughout data collection. As mining progresses, old cables can be cut and new cables surrounding the current region of interest can be connected to the system, which can measure tens of kilometers of fiber.

The spatially and temporally dense DAS data enables the detection of subtle seismic events that might otherwise go unnoticed with conventional monitoring systems. This enhanced detection capability is particularly advantageous in the context of underground coal mines, where early identification of minor seismic activities can be useful for delineating weak zones and identifying progressive failures quickly. DAS cables installed in boreholes have been previously tested for detecting

microseismicity (Luo & Duan, 2021), as well as deployed on a longwall for monitoring face bursts (Chambers & Shragge, 2023). This study differs by focusing on deploying fiber throughout the more easily accessible entries of the mine, which could yield a more cost-effective strategy that is less intrusive to operations.

One conventional approach to seismic event detection involves using the short-term average/long-term average (STA/LTA) technique, which calculates the ratio of energy in a signal's short and preceding longer time windows (Trnkoczy, 2012). However, this method often leads to false or missed seismic event identification due to its dependency on background noise levels. Hence, there has been growing interest in using deep learning techniques for seismic event detection, which have proven effective even for small-magnitude events (Huang et al., 2018; Shaheen et al., 2021; Zhu & Beroza, 2019). The two primary methods used for implementing deep learning in this context are supervised (Birnie & Hansteen, 2022; Mahmoudian et al., 2023) and unsupervised learning (Zipfel et al., 2023). In supervised learning, data are labeled as either seismic event or background noise (Mousavi et al., 2019), whereas in unsupervised learning, such labels are not required for training (Seo et al., 2024). Each of these approaches has its advantages, but supervised approaches require numerous labeled seismic events, which can be difficult to acquire. Despite the potential of unsupervised learning, there are limited examples of its application in distinguishing seismic events. Leveraging artificial intelligence to improve seismic event detection, this study employs an unsupervised approach to train a deep learning model for identifying seismic activities over continuous DAS recordings.

Anomaly detection is one of the leading applications of unsupervised learning that can be used for seismic event detection (Seo et al., 2024). Autoencoders are a type of convolutional neural network that is often used for unsupervised learning and are particularly useful for anomaly detection (Jiang et al., 2022; Mirzaee et al., 2023; Mousavi et al., 2019; Seydoux, L. et al., 2020; Shomal Zadeh et al., 2024). They can also be beneficial for studies involving seismic event clustering and seismic hazard assessment of triggered events (Nam & Wang, 2019; Seydoux, L. et al., 2020; Yaghmaei-Sabegh et al., 2022). Autoencoders aim to identify statistical outliers and are useful in seismic analysis as data containing seismic events are much less frequent than data containing only background noise. Unsupervised learning-based autoencoder methods leverage this fact to train deep neural network models, often leading to superior performance compared to traditional methods in event detection. In this study, we aimed to develop a seismic event detection system that uses DAS multichannel measurements and improves existing

seismic catalogs generated from surface seismic data. DAS can easily generate terabytes of data per day, and therefore, the use of an automated tool for anomaly (i.e., seismic event) detection is necessary. Hence, this research provides an unsupervised deep learning model that helps with the detection of seismic events on multichannel DAS recordings. After verifying the trained deep learning model on a sample of event-free background noise and achieving a satisfactory performance using seismic events from the surface seismometer network, we ran the model on DAS recordings to find seismic events that were not previously recorded. This anomaly-detection algorithm for seismic event detection could function as a warning system, potentially enhancing safety in coal mines by providing advanced notices of seismic activity.

2. FIELD STUDY DESIGN AND DAS DATA ACQUISITION

To determine if DAS could be a practical tool to improve event detection in underground coal mines, we carried out a study in a longwall mine in Virginia, USA, which has a history of seismicity concerns (Van Dyke et al., 2023). Based on our experience in a small pilot test around one pillar, we decided to focus on surrounding the active panel with fiber optic cables. The optical interrogator unit was stored in a building on the surface with power and climate control adjacent to an elevator and ventilation shaft. The interrogator unit was connected to a fiber optic patch panel, which led to a fiber cable in a borehole that descended into the mine. At the bottom of the shaft, a new fiber was spliced to a cable deployed in the mains leading to the active panel (~1490 meters of fiber from the shaft). From the mains on the west side of the panel, the fiber was spliced to another cable that extended ~1790 meters along the headgate (Figure 1). That headgate fiber was spliced to another fiber in the same cable at the east end of the segment to double the density, and then that fiber's west end was connected near the mains to another fiber that extended to the east along the tailgate ~1580 meters. Unfortunately, the splice to the tailgate did not receive enough backscattered light, either due to a splice stored with too tight a bend in its protective case, or simply too poor quality of a splice. Alternatively, the signal in the tailgate was reduced too much by traveling through multiple splices. Therefore, only the mains and headgate (forward and reverse) sections of the fiber were used for data processing, as shown using a blue line in Figure 1.

The fiber in the headgate and tailgate was installed in entries that were least likely to collapse early, and which had less machinery activity to reduce the likelihood of fiber breaks. In the headgate and tailgate, the majority of fiber was pushed against the rib (wall) and covered with mud or rock dust (non-combustible dust that mines apply to suppress potential explosions) where possible to

increase the coupling to the ground. In the mains, most fiber was strung up on hooks along with other cables, and along the entry to the headgate and tailgate it was connected to rib and roof bolts using zip ties. The cable was not coupled well to the rock in these regions because mine traffic would likely break cable on the floor. The data were acquired for 46 days from May 11 – June 27, 2022, during which time the orange region in Figure 1 was mined. The data were acquired at a channel spacing of 5.7 meters (i.e., distance between new measurements), a gauge length of 11.4 meters (i.e., distance over which average strain rate is measured), and a sampling rate of 2000 Hz, resulting in 33 T.B. of data.

Figure 2a illustrates the full fiber route (i.e., main, headgate, and tailgate) and a couple of the seismic events detected by both surface seismometers (noted as yellow triangles in Figure 1) and DAS cable. The DAS recording of the 1.2 magnitude event shown in Figure 2a is presented in Figure 2b. Clear P- and S-wave arrivals appear in the headgate forward and reverse DAS channels.




Fig. 1: A plan view of the mine showing the locations of seismic events (red dots) detected by the surface network and mined panel (orange rectangle) during the experiment, surface seismic stations (yellow triangles), and DAS fiber in headgate (blue line)




Fig. 2: DAS array channel locations with three seismic events (a) and the 1.2 magnitude event's P- and S-wave arrivals detected by DAS (b)

3. METHODOLOGY AND DAS DATA ANALYSIS WORKFLOW

3.1. Multichannel Spectral Analysis

The DAS technique enables us to record seismicity with a dense channel spacing, yielding higher resolution measurements compared to traditional seismic arrays. Although the signal-to-noise ratio of DAS measurements is lower than traditional seismometers, the spatially dense sampling and continuous recording enables characterization of various machinery noises and detection of subtle seismic events that might otherwise go unnoticed with conventional seismic monitoring systems. We use spectral methods to detect and classify both machinery noise and seismicity, which offer significant advantages over the time-domain alternative for several reasons. Primarily, the spectral approach allows for improved identification and analysis of different frequency components which often shed light on the nature of the sources. The frequency content of seismic signals can help differentiate between various types of seismic events, such as those caused by natural tectonic processes versus those induced by machinery or human activities. Furthermore, spectral analysis enhances the detection of low-amplitude signals that may be obscured by noise in the time domain, potentially improving the sensitivity and reliability of seismic monitoring systems. This capability is particularly beneficial in environments such as coal mines where the background noise level is high.

To monitor seismic activities using spectral analysis, we apply the Fast Fourier Transform to each trace recorded by a DAS channel, but a selection of a time window length to compute the power spectral density (PSD) is crucial. A long time window may not clearly represent the seismic event, especially for small-magnitude events far from the cable. On the other hand, if a time window is too short, the low-frequency content of the event signal may not be captured. A time window that includes the P- and S-wave arrivals, but not much more, is ideal. Figure 3 demonstrates the capability of spectral analysis in identifying a seismic event that was not previously detected by the surface network. Figure 3a shows the seismic event recorded by DAS in the time domain along with the PSD results in the frequency domain over different time windows in Figure 3b-d.




Fig. 3: A detected seismic event's P-wave arrival using DAS (a) and its power spectral density (PSD) plots across 275 channels during (b) 250, (c) 50, and (d) 0.5 seconds. Each title in b-d indicates the start and end time of the window that captures this event.

3.2. Autoencoder Deep Learning Model

An autoencoder is designed to compress (encode) the input data into a lower-dimensional representation (latent space) and then reconstruct (decode) the compressed data back to its original form using only the reduced representation. In the context of anomaly detection in images, such as those of Figure 3b-d, the autoencoder is trained exclusively on data representing the "normal", anomaly-free state of the data. During this training phase, we train the autoencoder to capture the essential characteristics of the normal images in the compressed representation. When the trained autoencoder is exposed to new images, it attempts to reconstruct them based on the model weights that best fit the training data. Images similar to the training data are reconstructed with relatively minor errors. Anomalous images, which contain patterns or features, such as seismic events that are not present in the training data, result in significantly higher reconstruction errors. By quantifying this error, anomalies can be detected.

Quantifying reconstruction error solely with simple metrics like root mean squared error often works well; however, it does come with certain limitations. One primary challenge is the autoencoder's potential to over-generalize from the training data, which can lead to lower-than-expected reconstruction errors for anomalous images. An alternative approach is to use kernel density estimation (KDE) as a non-parametric estimator to model the probability density function (PDF) of the encoded latent space, which we refer to as a density score (Chen, 2000; Chen, 2017). A density score, attributed to a specific point within the data space, is the PDF at that point. This quantification can serve as a direct measure of the local density or crowding of the data space in the vicinity of the point in question. Elevated density scores are indicative of the point's location within a region characterized by a significant concentration of data points, denoting a high probability density region. In contrast, reduced density scores signal the point's placement in a region of a sparse data point distribution, suggesting its association with low probability density or outlier regions. Such a distinction offers a granular perspective on the spatial distribution of data points, enabling the identification of anomalies based on deviations from established density norms within the multidimensional data space.

Figure 4 presents a schematic illustration of the autoencoder model. The autoencoder comprises an input layer designed to accommodate batches of 64 images of size 512x512 pixels. The encoder segment of the model

consists of a series of convolutional layers with 64, 32, and 16 filters, respectively, to progressively reduce the spatial dimensions while retaining essential information about the original images. These layers utilize the ReLU activation function for non-linearity (Agarap, 2018) and are arranged to ensure minimal information loss up to the bottleneck layer, which serves as the latent space representation of the input data. The decoder mirrors this with upsampling and convolutional layers, culminating in an output that is a reconstructed version of the input. An Adam optimizer and mean squared error loss were used during training to ensure the model emphasizes the precision of reconstruction discrepancy for anomaly identification.




Fig. 4: Schematic representation of the autoencoder model, highlighting its multilayer architecture that consists of fully connected layers.

To train the unsupervised deep learning model, we randomly selected 1000 two-second time windows of the multichannel DAS recordings which do not contain any events according to the surface seismic catalog. Then, we plotted their normalized PSD in RGB format and visually excluded any spectrum plot with anomaly. Finally, we trained the model on 960 spectrum plots of background DAS noise (normal data), using 768 (80%) for training and 192 (20%) for testing, over 250 training cycles (epochs). Figure 5 shows the deep learning model performance on training and testing datasets based on the loss value, calculated as the mean squared error of the difference between the predicted output and the actual output across both training and testing datasets. As the epochs progress, the convergence of loss values suggests that the model is effectively capturing the underlying patterns within the training data.




Fig. 5: Training and validation loss over 250 epochs. The training and validation loss strictly decreasing suggests convergence without overfitting.

After analyzing records with known seismic events and noting that the duration of all seismic events is between 0.5 and 1 second, we selected two-second time windows (chunks) with a one-second overlap between chunks. The overlap ensures every event is fully captured by at least one window. We implemented the proposed workflow on 26 hours of continuous DAS recording, generating 2,140,710 two-second chunks with 1 second of overlap and calculating their PSD plots. Subsequently, we employed the trained autoencoder on the PSD plots to detect seismic events.

4. RESULTS AND DISCUSSION

4.1. Noise Characterization and Analysis

Noise characterization of mining operations offers the potential for insights into equipment conditions and operations monitoring. Figure 6 showcases the PSD of different noises averaged across a group of 10 adjacent DAS channels. Figure 6a highlights a discernible daily variation in operational noise levels, with a pronounced increase during day shifts—a time of heightened activity—compared to the quieter night shifts. Figure 6b reveals distinct on-and-off noise patterns associated with machinery over a 9-hour timeframe. The longwall shearers are a probable noise source based on the power distribution of the channels near the longwall. The noise is broadband and the oscillation at the high frequency is apparent. Monitoring this noise could inform predictive equipment maintenance, potentially reducing downtime and resulting in cost savings on equipment maintenance.




Fig. 6: PSD of DAS recording showing mine's working shifts daily pattern (a) and machinery on-off pattern 245 m from the active face (b).

4.2. Enhanced Event Detection

We employed the trained autoencoder on the training dataset containing PSD images that only include background noise and on the PSD images from known seismic events. This approach enabled us to analyze the distribution of the KDE for the latent space's density score. A higher density score indicates a higher likelihood that the PSD data corresponds to the normal dataset, facilitating the detection of seismic activity. This unsupervised strategy, augmented with the incorporation of some labeled anomalous data, enhances the model's reliability, particularly in reducing false detections.

During the autoencoder model's development, a parametric study revealed the significant influence of input image size on the model's performance. The size parameter mandates resizing all input images to square dimensions defined by this parameter before either training or processing. This uniformity in image size is a restriction of the model's architecture. The selected image size directly affects the model's computational demands, its precision in capturing details, and its overall efficacy in differentiating between normal and anomalous seismic activities based on visual data. Although enlarging the size parameter enhances the model's anomaly detection capabilities, it substantially escalates computation time and memory usage. Figure 7 illustrates the density score for both normal (PSD data containing background noise) and anomalous (PSD data containing seismic events) datasets across varying image input sizes. The histograms, based on different size parameters, represent the distribution of the density scores for anomalous PSD data (left Y-axis). In contrast, the lines indicate the mean density score of the normal PSDs. The results highlight that increasing the size spreads out the density scores for PSDs with anomalies and also amplifies the disparity between the density scores for anomalous and normal PSD data. This amplification boosts the model's discriminative capability. For a size parameter of 512, the density score for the normal PSD is shown as a solid line at $\sim 45,000$, which contrasts distinctly with the distribution and histogram for the anomalous PSD data. Therefore, a size parameter of 512 was selected for the autoencoder model. Leveraging the outcomes from this primarily unsupervised approach, we set the 95th percentile of the

anomalous PSDs' density scores at $\sim 16,000$ as the threshold for event detection. Thus, any PSD data with a density score below this threshold is classified as an anomaly, indicating a seismic event. However, although adopting the 95th percentile rather than the maximum density score from known events may overlook a couple of very weak events with high density scores (around 40,000 and potentially events that are far from the DAS cable), it effectively reduces false positives. This methodology may ensure that the model is appropriately calibrated to avoid excessive sensitivity to minor anomalies in the PSD data, such as those caused by nearby vehicles.

Figure 8 illustrates the application of the developed autoencoder model for the detection and labeling of seismic events. In this figure, two PSD images are fed into the model as inputs: one representing background noise (Figure 8a) and the other depicting a seismic event (Figure 8b). The model processes these inputs to generate outputs that are reconstructions of their encoded latent space representations. The noise image produced a density score of 45,246, which is not anomalous. The reconstructed image of the noise image is a faithful but smoothed reproduction of the actual image. In contrast, the seismic event is identified as anomalous based on its density score of $-12,373$. The reconstructed image of the seismic event exhibits noticeable discrepancies from the original PSD image, as highlighted by the ovals in Figure 8b.




Fig. 7: Histogram of density score for normal (containing only background noise) and anomalous (containing a seismic event) data based on different input image sizes along with the mean density score of the normal PSDs.




Fig. 8: Comparison of a background noise (a) and a seismic event (b) using the developed autoencoder deep learning method (X-axis: Sensors [275 channels (headgate forward) = 1570 m]; Y-axis: Frequency [0-1000 Hz. (Gradually increasing from top to bottom)]; Color: PSD amplitude in RGB)

We ran the developed model on 15.6 T.B. of DAS data collected over 26 days, utilizing 275 channels (spanning 1570 meters) of the headgate forward cable for the purpose of event detection. After processing all the two-second PSD images with a one-second overlap, the model identified 56,558 anomalies as potential seismic events, compared to the sparse surface network's seismic catalog of 528 events during the same period. Note that the surface array covers a much larger area than the DAS array, but is located farther from the active longwall face, complicating the comparison. We reviewed time series and PSD plots of 500 instances randomly selected from data labeled as background noise and 500 selected from data labeled as anomalies, comparing them with seismic events recorded by the surface seismic network. 496 of the 500 (99.2%) visually inspected data detected as background noise did not include seismic events. The remaining 4 instances exhibited very weak anomalies; however, none constituted a seismic event, resulting in zero false negatives. Of the data labeled as anomalies, at least 85 of the 500 (17.0%) are potentially seismic events. None of the 85 events existed in the surface network catalog. The absence of the surface network catalog events in the reviewed time series and PSDs is attributed to the mere small random sampling of anomalous data, 500 out of 56,558 (0.9%). These newly detected events are likely of small magnitude and close to the DAS cable. The relatively low prediction accuracy is due to the high threshold for the density score that led to many false positives (i.e., 83.0%). Based on analyzing known seismic

events detected by both the surface seismic network and DAS, we have observed that for events occurring relatively far from the cable, especially those of smaller magnitude, the anomalies are not clearly apparent in the DAS data. Therefore, by lowering the threshold of the density score, we can more effectively detect apparent anomalies, potentially those of higher magnitude or occurring closer to the fiber. However, this approach introduces a trade-off in that while we can achieve a more robust outcome for seismic event detection, we may overlook labeling other weak potential seismic activities. For future studies, a more robust outcome may be achieved by employing additional metrics, such as the reconstruction root mean squared error, alongside the density score to detect seismic events more effectively.

Results show that the model successfully identified new seismic events that are likely small but close to the DAS cable and had not been previously recorded by the surface seismic network. This showcases the enhanced sensitivity and effectiveness provided by DAS technology deployed underground in detecting seismic events that are otherwise elusive to traditional surface monitoring systems. The integration of underground DAS and surface seismometers represents a useful advance in applied seismology, suggesting the potential for improved predictive models and safety measures in mines with seismic events.

5. CONCLUSIONS

This research provides an unsupervised deep learning model that helps with the detection of anomalies on multichannel DAS recordings. 56,558 potential seismic events were detected during 26 days of coal mining, compared to the surface network's seismic catalog of 528 events, showing an increase of more than 100 times. Analyzing time series and PSD plots of 1,000 instances against seismic events recorded by the surface network indicates a 99.2% accuracy in identifying non-seismic background noise, thanks to model training on such data. Additionally, preliminary analysis suggests that at least 17.0% of the detected anomalies could be seismic events, none of which were detected by the surface network. The modest success in detecting seismic events stems from a high density score threshold, which although limits effective seismic event detection, ensures detection of distant or low-magnitude events. A further statistical study is needed for selecting a more effective threshold. The developed anomaly detection algorithm for seismic event identification could serve as a warning system, potentially enhancing safety in coal mines by offering advanced notice of seismic activities. Furthermore, this algorithm could reduce data storage requirements by a factor of 50 by recording data when the algorithm detects an anomaly, thereby facilitating a mine's ability to manage voluminous DAS data more effectively.

ACKNOWLEDGEMENTS

This research was funded by the NSF IUCRC Center to Advance the Science of Exploration to Reclamation in Mining (CASERM), including through additional support from the National Institute for Occupational Safety and Health (NIOSH). We would like to thank Virginia Tech Department of Geosciences IT, Virginia Tech ARC and Colorado School of Mines CIARC for computing resources and data management support. We appreciate the collaboration, guidance, and site access of the team at the Buchanan Mine Complex of Coronado Global Resources. We thank James Garner of Oak Ridge National Laboratory for his collaboration on the data acquisition. We thank experts from NIOSH and Compass Minerals for their advice. We thank Jake Beale of Pilot Geophysical for his collaboration related to the surface seismic network.

DASCore Python library (Chambers et al., 2024) was used for DAS PSD analysis along with TensorFlow Python library (Abadi et al., 2016) and Bhattiprolu's GitHub repository (2023) for developing the autoencoder algorithm.

The findings and conclusions in this report are those of the author(s) and do not necessarily represent the official position of the National Institute for Occupational Safety and Health, Centers for Disease Control and Prevention. Mention of any company or product does not constitute endorsement by NIOSH, CDC.

REFERENCES

1. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., ... Zheng, X. (2016). TensorFlow: A system for large-scale machine learning. *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016*, 265–283.
2. Agarap, A. F. (2018). Deep learning using rectified linear units (relu). *arXiv preprint* <https://doi.org/10.48550/arXiv.1803.08375>
3. Ankamah, A., Hole, J. A., Tourei, A., Martin, E. R., Chambers, D., Beale, J., & Garner, J. (2023). Monitoring Mining-Induced Seismicity and Mine Operations Using an Underground Distributed Acoustic Sensing Array. *AGU*.
4. Bhattiprolu, S. (2023). Python for Microscopists. In *GitHub Repository*. https://github.com/bnsreenu/python_for_microscopists/blob/master/260_image_anomaly_detection_using_autoencoders/
5. Birnie, C., & Hansteen, F. (2022). Bidirectional recurrent neural networks for seismic event detection. *Geophysics*, 87(3). <https://doi.org/10.1190/geo2020-0806.1>
6. Chambers, D., Jin, G., Tourei, A., Issah, A. H. S., Lellouch, A., Martin, E. R., Zhu, D., Girard, A., Yuan, S., Cullison, T., Snyder, T., Kim, S., Danes, N., Punithan, N., Boltz, S., & Mendoza, M. M. (2024). DASCore: a Python Library for Distributed Fiber Optic Sensing. *Earth Arxiv. Preprint*.
7. Chambers, D., & Shragge, J. (2023). Seismoacoustic Monitoring of a Longwall Face Using Distributed Acoustic Sensing. *Bulletin of the Seismological Society of America*, 113(4). <https://doi.org/10.1785/0120220219>
8. Chen, S. X. (2000). Probability density function estimation using gamma kernels. *Annals of the Institute of Statistical Mathematics*, 52(3). <https://doi.org/10.1023/A:1004165218295>
9. Chen, Y. C. (2017). A tutorial on kernel density estimation and recent advances. *Biostatistics and Epidemiology*, 1(1). <https://doi.org/10.1080/24709360.2017.1396742>
10. Huang, L., Li, J., Hao, H., & Li, X. (2018). Microseismic event detection and location in underground mines by using Convolutional Neural Networks (CNN) and deep learning. *Tunnelling and Underground Space Technology*, 81. <https://doi.org/10.1016/j.tust.2018.07.006>
11. Jiang, J., Ren, H., & Zhang, M. (2022). A Convolutional Autoencoder Method for Simultaneous Seismic Data Reconstruction and Denoising. *IEEE Geoscience and Remote Sensing Letters*, 19. <https://doi.org/10.1109/LGRS.2021.3073560>
12. Luo, X., & Duan, Y. (2021). Microseismic monitoring of longwall caving process using distributed optic fiber sensing. *55th U.S. Rock Mechanics / Geomechanics Symposium 2021*, 2.
13. Mahmoudian, A., Tajik, N., Taleshi, M. M., Shakiba, M., & Yekrangnia, M. (2023). Ensemble machine learning-based approach with genetic algorithm optimization for predicting bond strength and failure mode in concrete-GFRP mat anchorage interface. *Structures*, 57. <https://doi.org/10.1016/j.istruc.2023.105173>
14. Mark, C. (2016). Coal bursts in the deep longwall mines of the United States. *International Journal of Coal Science and Technology*, 3(1). <https://doi.org/10.1007/s40789-016-0102-9>
15. Mark, C. (2018). Coal bursts that occur during development: A rock mechanics enigma. *International Journal of Mining Science and Technology*, 28(1). <https://doi.org/10.1016/j.ijmst.2017.11.014>
16. Mirzaee, H., Kamrava, S., & Tahmasebi, P. (2023). Minireview on Porous Media and Microstructure Reconstruction Using Machine Learning

- Techniques: Recent Advances and Outlook. *Energy & Fuels*, 37(20), 15348–15372.
<https://doi.org/10.1021/acs.energyfuels.3c02126>
17. Mousavi, S. M., Zhu, W., Sheng, Y., & Beroza, G. C. (2019). CRED: A Deep Residual Network of Convolutional and Recurrent Units for Earthquake Signal Detection. *Scientific Reports*, 9(1).
<https://doi.org/10.1038/s41598-019-45748-1>
 18. Nam, K., & Wang, F. (2019). The performance of using an autoencoder for prediction and susceptibility assessment of landslides: A case study on landslides triggered by the 2018 Hokkaido Eastern Iwate earthquake in Japan. *Geoenvironmental Disasters*, 6, 1-14.
 19. Peng, P., He, Z., Wang, L., & Jiang, Y. (2020). Automatic classification of microseismic records in underground mining: A deep learning approach. *IEEE Access*, 8.
<https://doi.org/10.1109/ACCESS.2020.2967121>
 20. Peng, S. S. (2019). Longwall Mining, 3rd Edition. In *Longwall Mining*. CRC Press.
<https://doi.org/10.1201/9780429260049>
 21. Seo, J., Kim, Y., Ha, J., Kwak, D., Ko, M., & Yoo, M. (2024). Unsupervised anomaly detection for earthquake detection on Korea high-speed trains using autoencoder-based deep learning models. *Scientific Reports*, 14(1), 639.
<https://doi.org/10.1038/s41598-024-51354-7>
 22. Seydoux, L., Balestrieri, R., Poli, P., Hoop, M. D., Campillo, M., & Baraniuk, R. (2020). Clustering earthquake signals and background noises in continuous seismic data with unsupervised deep learning. *Nature communications*, 11(1), 3972.
<https://doi.org/10.1038/s41467-020-17841-x>
 23. Shaheen, A., Waheed, U. Bin, Fehler, M., Sokol, L., & Hanafy, S. (2021). Groningenet: Deep learning for low-magnitude earthquake detection on a multi-level sensor network. *Sensors*, 21(23).
<https://doi.org/10.3390/s21238080>
 24. Shomal Zadeh, S., Khorshidi, M. & Kooban, F. (2023). Concrete surface crack detection with convolutional-based deep learning models. *SSRN Electronic Journal*.
<https://doi.org/10.2139/ssrn.4661249>
 25. Swanson, P., Boltz, M. S., & Chambers, D. (2016). Seismic Monitoring Strategies for Deep Longwall Coal Mines. In *Centers for Disease Control and Prevention, National Institute for Occupational Safety and Health*.
<https://stacks.cdc.gov/view/cdc/43944>
 26. Trnkoczy, A. (2012). Understanding and parameter setting of STA/LTA trigger algorithm. In *New Manual of Seismological Observatory Practice 2 (NMSOP-2)*. https://doi.org/10.2312/GFZ.NMSOP-2_IS_8.1
 27. Van Dyke, M., Klemetti, T., Khademian, Z., Wickline, J., & Beale, J. (2023). Evaluation of Seismic Potential in a Longwall Mine with Massive Sandstone Roof Under Deep Overburden: an Update. *Mining, Metallurgy and Exploration*, 40(5).
<https://doi.org/10.1007/s42461-023-00818-x>
 28. Wang, H., Fratta, D., Lord, N., Zeng, X., & Coleman, T. (2018). Distributed Acoustic Sensing (DAS) Field Trials for Near-Surface Geotechnical Properties, Earthquake Seismology, and Mine Monitoring. *SEG Technical Program Expanded Abstracts*. <https://doi.org/10.1190/segam2018-2997833.1>
 29. Wei, C., Zhang, C., Canbulat, I., Cao, A., & Dou, L. (2018). Evaluation of current coal burst control techniques and development of a coal burst management framework. *Tunnelling and Underground Space Technology*, 81, 129–143.
<https://doi.org/10.1016/J.TUST.2018.07.008>
 30. Yaghmaei-Sabegh, S., Shokrgozar-Yatimdar, E. & Shoaieifar, P. (2022). Temporal clustering PSHA based on smoothing: a case study in the New Madrid seismic zone of the central USA. *J Seismol* 26, 119–135. <https://doi.org/10.1007/s10950-021-10060-x>
 31. Zhang, C., Canbulat, I., Hebblewhite, B., & Ward, C. R. (2017). Assessing coal burst phenomena in mining and insights into directions for future research. In *International Journal of Coal Geology* (Vol. 179).
<https://doi.org/10.1016/j.coal.2017.05.011>
 32. Zhu, W., & Beroza, G. C. (2019). PhaseNet: A deep-neural-network-based seismic arrival-time picking method. *Geophysical Journal International*, 216(1). <https://doi.org/10.1093/gji/ggy423>
 33. Zipfel, J., Verworn, F., Fischer, M., Wieland, U., Kraus, M., & Zschech, P. (2023). Anomaly detection for industrial quality assurance: A comparative evaluation of unsupervised deep learning models. *Computers and Industrial Engineering*, 177.
<https://doi.org/10.1016/j.cie.2023.109045>