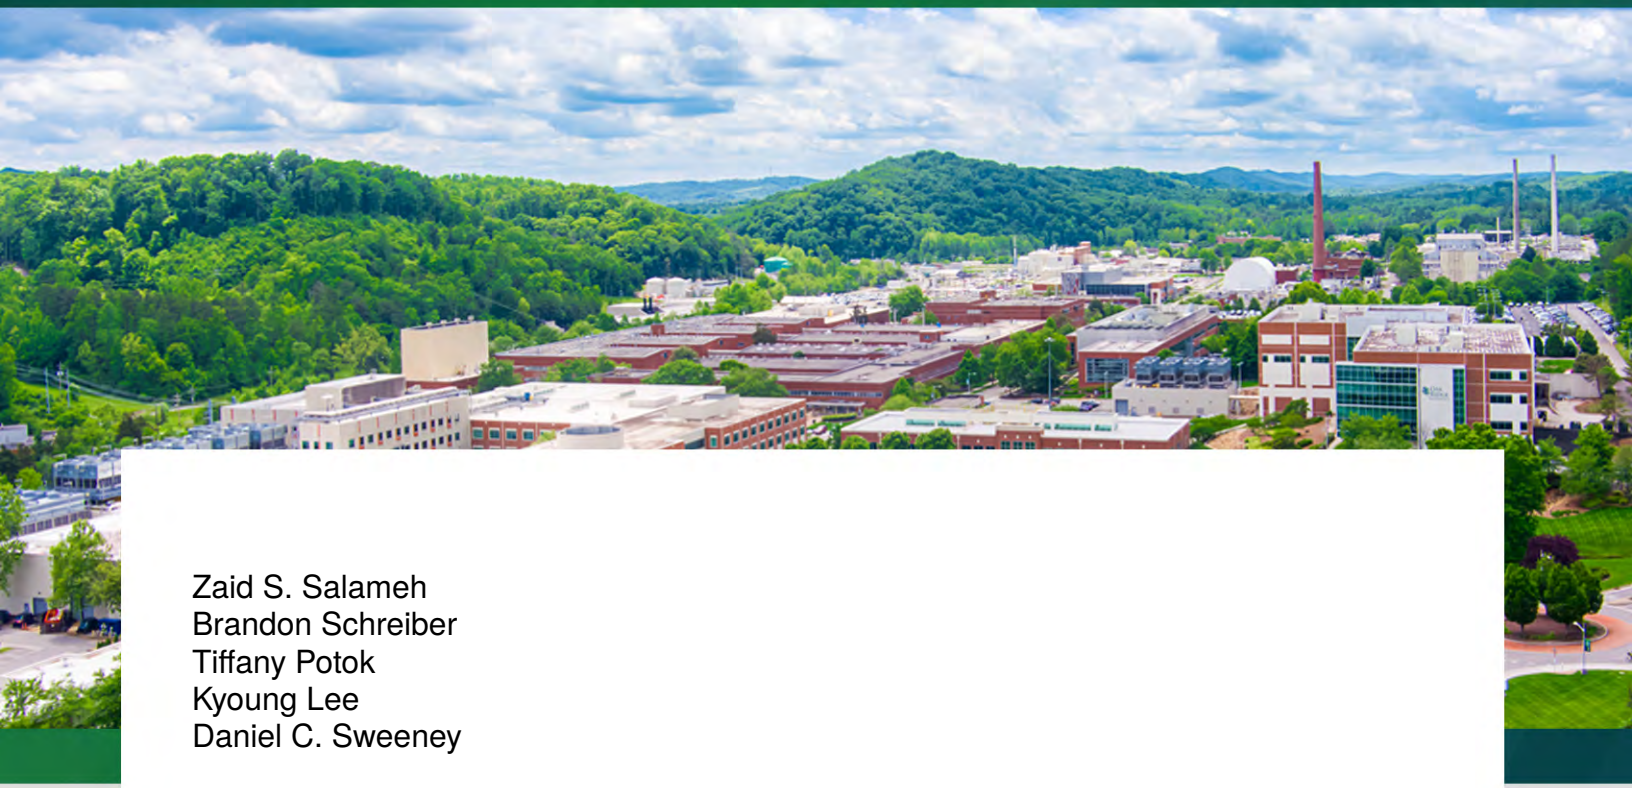
ORNL/TM-2026/4784

Machine Learning–Assisted Distributed Acoustic Sensing for Microreactors



Zaid S. Salameh
Brandon Schreiber
Tiffany Potok
Kyoung Lee
Daniel C. Sweeney

September 2026

ORNL is managed by UT-Battelle LLC for the US Department of Energy

OAK RIDGE National Laboratory

DOCUMENT AVAILABILITY

Online Access: US Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via https://www.osti.gov/.
The public may also search the National Technical Information Service's National Technical Reports Library (NTRL) for reports not available in digital format.
DOE and DOE contractors should contact DOE's Office of Scientific and Technical Information (OSTI) for reports not currently available in digital format:

US Department of Energy
Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831-0062
Telephone: (865) 576-8401
Fax: (865) 576-5728
Email: reports@osti.gov
Website: https://www.osti.gov/

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

ORNL/TM-2026/4784

Nuclear Energy and Fuel Cycle Division

MACHINE LEARNING–ASSISTED DISTRIBUTED ACOUSTIC SENSING FOR MICROREACTORS

Zaid S. Salameh
Brandon Schreiber
Tiffany Potok
Kyoung Lee
Daniel C. Sweeney

September 2026

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-BATTELLE LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

LIST OF FIGURES iv
LIST OF TABLES v
LIST OF ABBREVIATIONS vi
ABSTRACT 1
1. INTRODUCTION 2
2. METHODS 4
2.1 THEORY OF OPERATION 4
2.2 CONSTRUCTING A DAS DATASET 4
2.3 YOLO26n TRAINING 5
2.4 FIBER DEPLOYMENT AND GAUGE LENGTH 6
2.5 MODIFIED Z-SCORE ANOMALY DETECTION ON TIME SERIES, MEL, AND MFCC FEATURES 7
2.6 INDUSTRIAL ENVIRONMENT EVENT CLASSIFICATION 8
2.7 DEPLOYMENT AT THE SPALLATION NEUTRON SOURCE 9
2.8 SNS FREQUENCY-BAND ENERGY ANALYSIS 9
3. RESULTS 12
3.1 EFFECTS OF CROSS-DOMAIN MODEL TRAINING ON VALIDATION ACCURACY . 12
3.2 FIBER DEPLOYMENT AND GAUGE LENGTH IMPACT 14
3.3 DEVELOPMENT OF AN ANOMALY DETECTION PIPELINE FOR DAS MONITORING 14
3.4 EXTENDED DAS DATASET CLASSIFICATION WITH INDUSTRIAL SOUNDS 16
3.5 PROOF-OF-CONCEPT DEPLOYMENT AT THE SPALLATION NEUTRON SOURCE . . 17
4. DISCUSSION 20
5. CONCLUSIONS 22

LIST OF FIGURES

Figure 1. The DAS dataset was created by playing the source files over a speaker 1.2 m away from the fiber-optic spool. 5
Figure 2. ML-enhanced DAS acoustic classification pipeline using mel spectrograms. 5 s source files were classified and played through a speaker to the DAS system. The DAS- and source-derived spectrograms were used to train an object detection model. 6
Figure 3. Effects of time and frequency masking implemented as part of SpecAugment on spectrograms. Red arrow points to a horizontal mask. Yellow arrow points to a vertical mask 7
Figure 4. Modified z-score event detection pipeline to condense high-rate DAS stream to usable data packets for streamlined O&M. 8
Figure 5. Fiber-optic sensing deployed on industrial valves to capture industrially relevant acoustic sources: (a) a PRV and (b) a vacuum oven. 9
Figure 6. (a) DAS fiber-optic cable along equipment in the klystron building at SNS at ORNL. (b) Instantaneous OTDR trace of the fiber run. 11
Figure 7. Confusion matrices with spectrogram augmentation from same-domain validation of the DAS file trained model validated with (a) DAS recorded data and (b) the source recording trained model validated against source files. 13
Figure 8. Effect of RMS window length on the strain rate envelope. 15
Figure 9. MAD-based modified z-score computed on two stitched DAS clips with a 0.25 s RMS window, comparing two reference-frame choices on a single axis. 15
Figure 10. FNs (blue) and FPs (orange) versus z-score threshold for the DAS trigger detector, scored against the manifest with a ± 1 s match tolerance. 16
Figure 11. Event detection on the clean DAS record. 16
Figure 12. Event detection on the corrupted distributed acoustic sensing (DAS) record (100 Hz tone added). 17
Figure 13. Mel spectrograms of acoustic signature across six operating states of the pressure vacuum. 18
Figure 14. DAS-derived spectrogram of auxiliary machinery. 18
Figure 15. FBE from spectrograms measured in the klystron gallery of SNS. 19

LIST OF TABLES

Table 1. Training configurations evaluated, varying training data format, SpecAugment, and fine-tuning data format 6
Table 2. Summary of model development and live testing results. 12
Table 3. Validation accuracy results for models trained with source files or DAS recordings under different SpecAugment and fine-tuning conditions. 13
Table 4. Comparison of classification model accuracy based on gauge length and fiber layout topology 14

LIST OF ABBREVIATIONS

DAS	distributed acoustic sensing
DTL	drift tube linac
FBE	frequency band energy
FN	false negative
FP	false positive
MAD	median absolute deviation
MFCC	mel-frequency cepstral coefficient
ML	machine learning
MR	microreactor
O&M	operation and maintenance
ORNL	Oak Ridge National Laboratory
PRV	pressure relief valve
PRV	pressure relief valve
RFQ	radio-frequency quadrupole
RMS	root mean square
ROI	region of interest
SNR	signal-to-noise ratio
SNS	Spallation Neutron Source
STFT	short-time Fourier transform
TP	true positive
YOLO	You Only Look Once

ABSTRACT

This report describes an end-to-end pipeline for monitoring and classifying acoustic events in a microreactor (MR) using distributed acoustic sensing (DAS). The objective of this pipeline is to realize autonomous supervision of MRs to reduce their operational footprint. MRs promise modular, compact nuclear power, but their economic viability depends on minimizing operation and maintenance costs through remote, autonomous monitoring. To address this limitation, we investigated how optical fiber–based remote sensing could be used to support compact monitoring systems for energy facilities. Distributed optical fiber sensing, and DAS in particular, is well suited to this role because silica optical fibers are immune to electromagnetic noise, are tolerant to radiation and high temperatures, and are chemically inert. In this work, we establish the feasibility of a novel application of established machine learning (ML) event classification models and statistical anomaly detection algorithms to DAS-derived vibroacoustic data for industrial system health monitoring. We developed an end-to-end classification pipeline built on the You Only Look Once (YOLO) object recognition model and evaluated cross-source transfer between microphonic- and DAS-sourced data. The raw strain rate data were processed using an anomaly detection algorithm, which leveraged a modified z-score using the median absolute deviation (MAD) metric. A commercial DAS system was used to collect acoustic data at the Spallation Neutron Source (SNS) at Oak Ridge National Laboratory. The model trained on DAS data was effective in classifying unseen DAS data, cross-source classification was inaccurate, and fine-tuning source-derived weights with DAS spectrograms recovered strong DAS classification. Multiple anomaly detection or "trigger" approaches were explored, including a time-amplitude root mean square (RMS) trigger and a frequency-based detector. The RMS-based detector detects stationary events but fails under persistent low-frequency noise interference, and the frequency-based anomaly detector is robust to the elevated background noise. Deployment at SNS demonstrated full operation in a nuclear environment, validating the approach end to end as a proof of concept for MR monitoring and validating DAS hardware as viable for collecting vibroacoustic data for the detection and classification of MR operation noise. Dedicated DAS-derived datasets are required to achieve reliable accuracy. For future work, investigating advanced signal processing to bridge existing acoustic datasets to DAS data is a promising direction worth pursuing.

1. INTRODUCTION

Microreactors (MRs) are compact nuclear reactors that are factory manufactured, transportable, and produce power outputs up to 20 MWe [1]. They serve as an energy source for applications poorly suited to conventional generation, including remote communities and industrial operations independent from the main electric grid [2]. To advance the commercialization of MRs, they must demonstrate economic competitiveness against relatively inexpensive fossil fuel alternatives and established terrestrial nuclear power plants [3]. Shifting from staffed, scheduled maintenance to a remote, largely unmanned model could reduce operation and maintenance (O&M) costs by more than 80% [4]. Therefore, advanced instrumentation and machine learning (ML) algorithms that detect and classify irregularities of system health would advance the economic viability of MRs [3].

Distributed acoustic sensing (DAS) is a well-characterized technology for system health monitoring in railways [5], pipelines [6], and highways [7] as well as geological movement [8] and ocean applications [9]. DAS instruments use a narrow-linewidth laser to inject coherent optical pulses (5 to 50 ns, or 1 to 10 m) into an optical fiber routed near the equipment of interest and measure the reflection returned from intrinsic scattering sites along the fiber. This Rayleigh backscatter is recorded by the DAS instrument, and the time of flight between pulse launch and signal arrival identifies the fiber segment from which each backscatter profile originates. Acoustic and vibrational disturbances strain the fiber locally, producing phase changes in the backscattered light that can be resolved to recover the strain rate along a single fiber with no distributed electronics, local power, or active components (outside of the interrogator itself). The fiber's immunity to electromagnetic interference, adaptive algorithms for extreme temperature environments [10], and tolerance to irradiation [11] make DAS especially attractive for acoustic event classification in nuclear environments.

Acoustic classification models generally use time-frequency representations, such as mel spectrograms and mel-frequency cepstral coefficients (MFCCs), to recognize acoustic events [12]. The short-time Fourier transform (STFT) decomposes each analysis frame into its constituent frequencies. For classification tasks, the linear-frequency spectrogram is often transformed to the mel scale, which spaces frequency bands to approximate the nonlinear resolution of human hearing [13]. The mel spectrogram has become a standard input to modern acoustic classifiers, including the convolutional neural networks that dominate environmental sound and acoustic scene recognition [14]. The You Only Look Once (YOLO) family of object detectors formulates detection as a single-stage regression problem, in which one convolutional passes over the whole image and returns class label and confidence score [15]. This is important because collapsing classification into a single pass yields the real-time speed required for deployment on streaming data and edge hardware. However, labeled training data for irregular fault conditions are scarce, and models trained on microphonic source data have little application to DAS-derived data. A goal of this report is to demonstrate that standard audio classification ML techniques can be leveraged with DAS instrumentation for streamlined O&M in the context of nuclear MRs.

In industrial system health monitoring, vibration signals are continuously surveilled to catch the onset of bearing wear, cavitation, or structural fatigue before failure [16]. There is a large body of work on this topic [17], and there are commercially available devices [18] that monitor and detect anomalies for predictive maintenance with traditional sensors. However, there is no currently available software that automates anomaly detection using DAS instrumentation. The challenge in DAS-based monitoring is that normal behavior must be characterized, often from a limited or noisy reference. A single fiber-optic cable (up to 100 km) connected to the interrogator, which can sample thousands of serial channels with ~5 m spatial resolution at kilohertz rates, will stream $\geq$ 1 TB/day. Extracting useful data from DAS therefore

requires advanced signal processing techniques to reduce the size of the data while preserving the relevant information. Statistics such as the root mean square (RMS) envelope summarize the amplitude of the signal over a sliding window, and thresholding these quantities yields a lightweight trigger suitable for real-time operation on high-rate data streams [19]. As environmental conditions evolve and low-frequency background events occur, the detector must remain sensitive to anomalies while tolerating shifts to the baseline.

This report documents several experiments conducted to classify acoustic events identified by DAS. First, we describe a pipeline for batch conversion of acoustic data into standardized mel spectrograms, with training and classification using an image-based ML model. Second, we present an anomaly detection algorithm for time series data based on a modified z-score with the median absolute deviation (MAD) in the time-amplitude domain, with extensions to frequency-domain features. Finally, we detail a proof of concept at the Spallation Neutron Source (SNS) facility, including hardware deployment, an end-to-end software pipeline that triggers on abnormal events and saves data, and subsequent frequency band analysis. This work advances the necessary hardware, signal processing, and ML techniques for DAS-derived strain rate files for nuclear MR applications.

2. METHODS

2.1 THEORY OF OPERATION

The DAS systems developed by Terra15 rely on a double-pulse scheme in which phase shifts between the first and second pulse measured directly correspond to strain rate measurements in the optical fiber [20]. This technique has been shown to increase the sensitivity of vibroacoustic measurements using DAS and decrease measurement noise associated with time differentiation strain measurements over other phase-sensitive coherent optical time-domain reflectometry techniques [21]. To this end, let $I(t)$ and $I(t)'$ describe the transfer functions for the first and second pulses in a train, in between which the optical fiber is strained. Interfering $I(t)$ and $I'(t)$ enables the calculation of the phase between the two pulses. The duration between the two pulses is constant and enables the calculation of a local velocity at a specific portion of the optical fiber. However, this local phase velocity is not measured in isolation but as an integrated phase velocity $\dot{\varphi}$ along the length of the fiber as

$$\dot{\varphi}(t, L) = \angle\left\{ \int_0^L \left(\frac{I(t,z) \otimes I'(t,z)}{\Delta t_d} \right) dz \right\}, \tag{1}$$

where z is the distance along the fiber of length L, relative to the fiber connection to the DAS instrument, at time t, and Δt_d is the time delay between pulses. The integrated phase velocity $\dot{\varphi}$ can be converted to a local strain rate as

$$\dot{\varepsilon}(t, z_i) = \frac{1}{L_g}\left[\dot{\varphi}(t, z_i + L_g/2) - \dot{\varphi}(t, z_i - L_g/2) \right], \tag{2}$$

where L_g is the length of the optical fiber segment (gauge) considered as a single sensor, typically on the order of 1 to 10 m. The DAS data can be processed over discrete periods (chunks), which can be represented as vector $\dot{\mathbf{e}}(, z_i)$ associated with T samples:

$$\dot{\mathbf{e}}(t, z_i) = [\dot{\varepsilon}(t_0, z_i), \dot{\varepsilon}(t_1, z_i), ..., \dot{\varepsilon}(t_T, z_i)]. \tag{3}$$

It should be noted that $\dot{\mathbf{e}}(t, z_i)$ may also be represented using a spatial average across a range of sensors to improve the signal-to-noise ratio (SNR).

Time-domain vibroacoustic signals are next transformed into tensors, which can be more readily classified using widely available ML models using STFTs over $K \leq T$ samples per computation interval over each data chunk as

$$\mathbf{S}(t, f, z_i) \approx \frac{1}{K}\left| \mathrm{DSTFT}\{\dot{\mathbf{e}}(t, z_i)\} \right|^2, \tag{4}$$

where $\mathbf{S}(t, f, z_i)$ is the real-valued rank-2 tensor representing the power spectral density through the Discrete Short-Time Fourier Transform (DSTFT) of $\dot{\mathbf{e}}(t, z_i)$. Importantly, to utilize $\mathbf{S}(t, f, z_i)$ for training and classification tasks in the ML model, it is ideal for the dimensions of the training data to match the dimensions of the target data to be classified. As such, the dimensions of the tensors in both the training and analysis were conserved.

2.2 CONSTRUCTING A DAS DATASET

It is of interest to investigate whether microphonic data could be used to train object detection models for DAS data and vice versa. The source data for this study were taken from 50 environmental sound classifications (ESC-50) [22]. A second copy of these data was recorded with the 40 km Treble+ DAS system (Terra15 Technologies Pty Ltd., Perth, Australia) to capture the unique background artifacts

introduced by the system. The DAS dataset was created by playing these files on a speaker (S1 Pro+, BOSE, Framingham, Massachusetts) 1.2 m away from a 1 km fiber spool (Figure 1). All optical fiber used in this study was SMF-28e+ (Corning Inc., Corning, New York, USA) packaged as either ITU-T G.652.D compliant OS2 single-mode fiber with a 3 mm jacket or bare fiber coated with with 200 µm dual acrylate. The DAS system was set to a pulse width of 4.9 m (24 ns) and high-pass filtered at $\geq$ 100 Hz to eliminate the DC component of the signal. Each recording in this combined dataset was resampled to a uniform 32-bit, 16 kHz file and then zero-padded or trimmed to 5 s segments. The time series signals for both the source files and the DAS files were converted to log-mel spectrograms prior to acoustic classification training (Figure 2). The spectrograms were computed using custom code in Python (v3.14) using the Librosa library (v0.11.0) with 128 mel bins, 1,024 samples per STFT, and a hop length of 256. Mel-band power was converted to decibels relative to the maximum value within each 5.0 s window, clipped to a dynamic range of −80 to 0 dB, and mapped to [0,1] for input to the classification model. These data were subject to an 80%/20% split for training and validation.

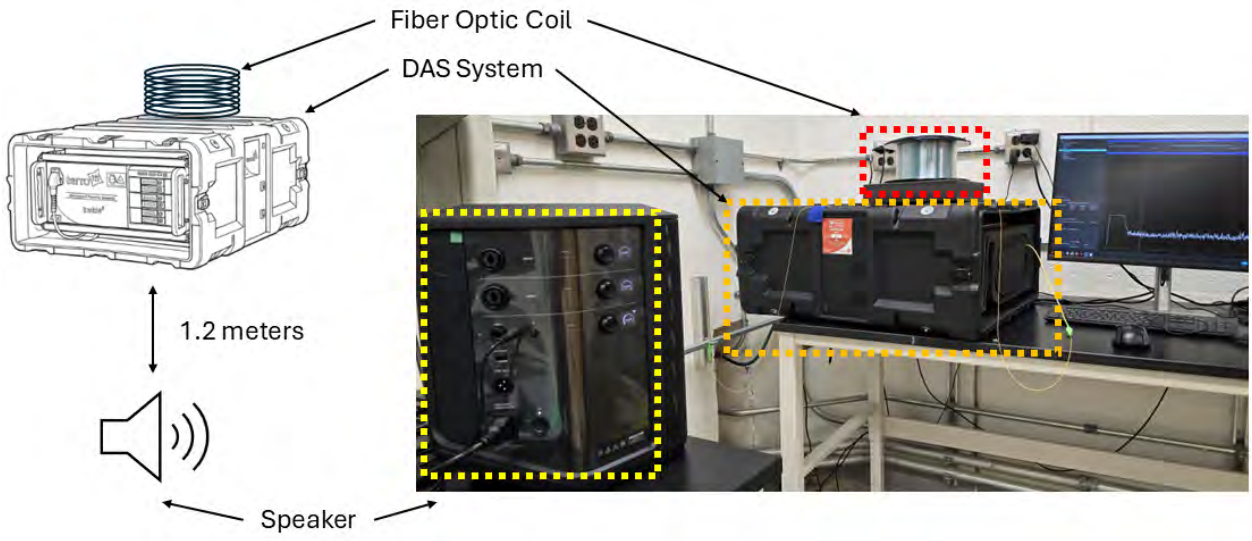


Figure 1. The DAS dataset was created by playing the source files over a speaker 1.2 m away from the fiber-optic spool.

2.3 YOLO26N TRAINING

We trained the head (stage) of a YOLO26n model [23] on the spectrogram dataset. Training used transfer learning from pretrained weights, standard image augmentation, and early stopping on the validation split. A custom trainer was used to train the `YOLO26n-cls` model on single-channel binary spectrogram arrays rather than the image formats typically used with the ML model. The first convolutional layer of the pretrained model was converted from three input channels to one input channel by averaging the pretrained RGB weights across the input channel dimension. The custom dataset loader reads the binary arrays from class folders, converts them to tensors of shape $(1, H, W)$, replaces invalid numeric values, and resizes them to the target input geometry. The baseline model training process used the AdamW algorithm [24] with a weight decay of 0.001 and a fixed seed for reproducibility. Model training settings were 50 epochs (with a patience timeout of 15 epochs), no dropout, and a learning rate of 0.01. Fine-tuning runs used 25 epochs and a reduced learning rate of 0.0001.

To test how the source of training data (DAS-derived versus source data) affects validation accuracy of the ML object detection model (YOLO26n-cls), we conducted eight experiments as displayed in Table 1.

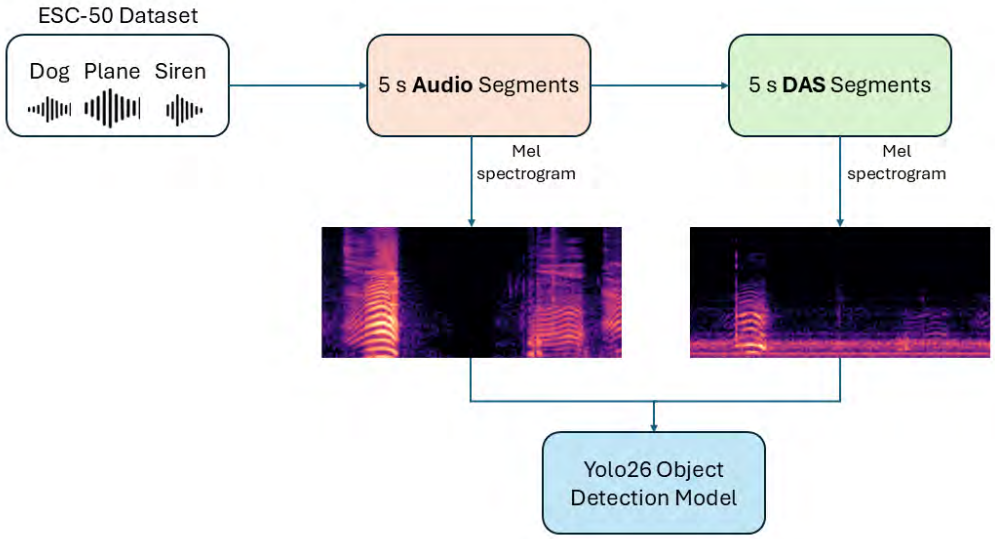


Figure 2. ML-enhanced DAS acoustic classification pipeline using mel spectrograms. 5 s source files were classified and played through a speaker to the DAS system. The DAS- and source-derived spectrograms were used to train an object detection model.

Table 1. Training configurations evaluated, varying training data format, SpecAugment, and fine-tuning data format

Training Data	SpecAugment	Fine-tuning Data
Source	N	—
Source	Y	—
Source	N	DAS
Source	Y	DAS
DAS	N	—
DAS	Y	—
DAS	N	Source
DAS	Y	Source

Spectrogram augmentation improves generalization to noise, missing information, and other real-world variations [25]. The loader applied a frequency mask to 6 mel bins at a time, and a time mask to 16 time frames, and a circular time roll from 0.97 to 1.03 for also 16 time frames, with additive noise contributing a standard deviation of 0.005, as shown in Figure 3. Each of these trained models was validated against held-out data. For each tensor, the validator recorded the true label, the model-predicted label, top-1 probability, and a top-5 class list.

2.4 FIBER DEPLOYMENT AND GAUGE LENGTH

To investigate the effect of gauge length used for classification and single-point versus multipoint sensor regions, a 1 km fiber spool with an additional 100 m of fiber spliced to the end of the spool and loosely

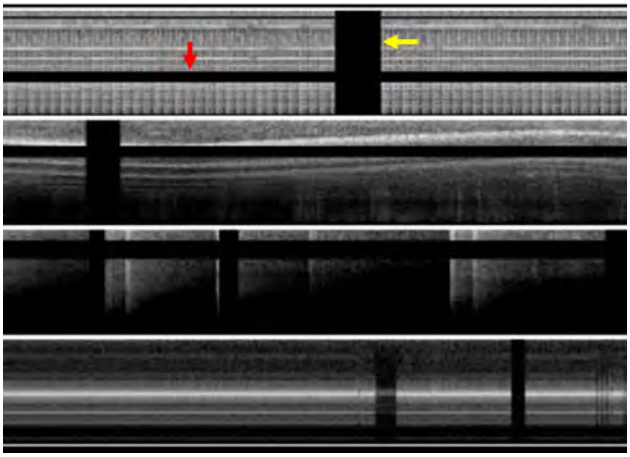


Figure 3. Effects of time and frequency masking implemented as part of SpecAugment on spectrograms. Red arrow points to a horizontal mask. Yellow arrow points to a vertical mask

strung around the lab was deployed. The full 73-class playback session was recorded with this setup, and offline validation was performed using various gauge length and location parameters.

Straight fiber performance was evaluated at 1,047 m for the one-point condition and at 1,032, 1,047, and 1,067 m for the three-point condition. Spooled fiber performance was evaluated at 500 m for the one-point condition and at 250, 500, and 750 m for the three-point condition. Spatial averaging lengths of 1, 2, 5, and 10 m were applied offline to the same physical recording.

The goal of these experiments was to evaluate the effect of the number of discrete sensing regions and how spatially averaging multiple discrete regions contributes to classification results. It was also of interest to investigate how the spatial location of the sensor relative to the acoustic source affects classification.

2.5 MODIFIED Z-SCORE ANOMALY DETECTION ON TIME SERIES, MEL, AND MFCC FEATURES

Because the DAS system continuously streams high-rate, noisy data, developing an event detection pipeline would highlight the regions of interest while condensing data storage needs. To achieve this, established trigger algorithms that continuously monitor input data via statistical methods could be employed [26] but require novel applications to DAS-derived data in an industrial nuclear environment. To this end, we developed a modified z-score event detection algorithm that leverages the MAD statistical method along with a moving RMS window, which were chosen because of robustness to outlier data [27] (Figure 4). The modified z-score was defined with

$$z_i = \frac{\rho_i - \tilde{\rho}_{\text{ref}}}{\text{MAD}_{\text{ref}}}, \tag{5}$$

where z_i is the modified z-score of the i-th RMS amplitude of the signal, ρ_i is the RMS amplitude of the i-th window, and $\tilde{\rho}_{\text{ref}}$ is the median of the reference RMS values. MAD_{ref}, which is the MAD of the reference RMS values, is calculated with

$$\text{MAD}_{\text{ref}} = \text{median}\left(\left|\rho_j - \tilde{\rho}_{\text{ref}}\right|\right), \tag{6}$$

where ρ_j is the j-th RMS amplitude of the reference value. The RMS is calculated by

$$\rho_j = \sqrt{\frac{1}{N}\sum_{k=1}^{N} \dot{\epsilon}_k^2}, \tag{7}$$

where ρ is the RMS amplitude of the window, $\dot{\epsilon}$ is the k-th strain rate sample within the window, N is the number of samples in the window (RMS window size), and k is the sample index, running from 1 to N.

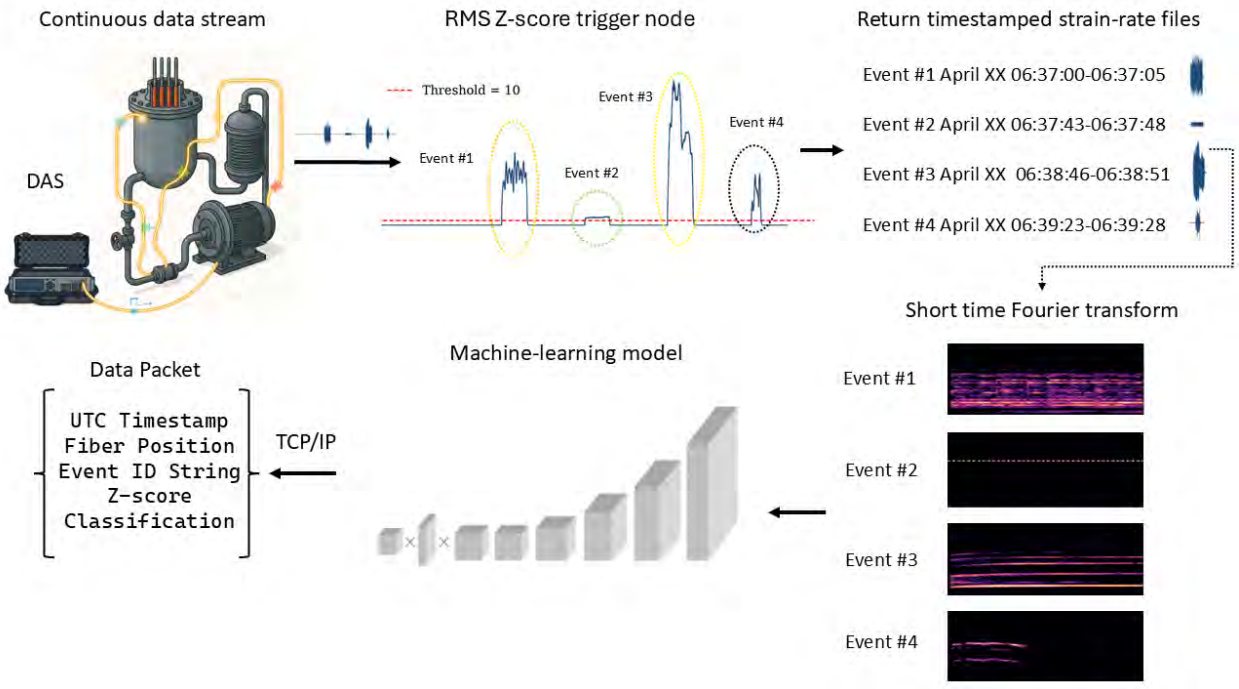


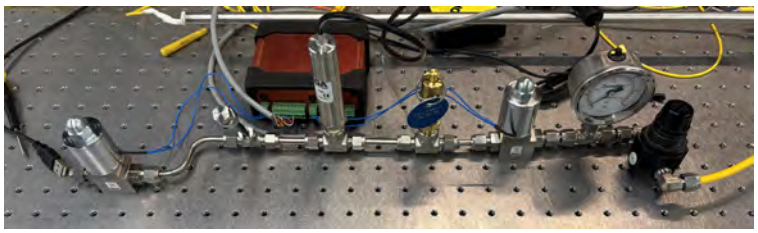
Figure 4. Modified z-score event detection pipeline to condense high-rate DAS stream to usable data packets for streamlined O&M.

To investigate the appropriate parameters in the context of DAS-derived acoustic data, we studied how varying the RMS window size, z-score threshold, and choice of reference frame impacted our false positives (FPs), false negatives (FNs), and F1 error score. Detector performance was quantified by matching detections against a labeled manifest of event start times, using a ±1 s tolerance window around each labeled start. A manifest event was counted as a true positive (TP) if at least one detection fell within its tolerance window and as an FN otherwise. A detection that did not fall within the tolerance window of any manifest event was counted as an FP. A 6 s lockout period was enforced following each trigger so that a single sustained event could not register as multiple detections. From these counts, we computed the F1 score, which is an established metric for evaluating anomaly detection algorithms for system health monitoring [28]. This metric is especially useful because it depends only on the events and the detections, penalizing missed events and spurious triggers symmetrically. Initially, we investigated z-score triggering on time series data with raw strain rate values. Then, we studied how data representation in the frequency domain could be used for event detection algorithms by introducing an artificial 100 Hz noise to our clean data and evaluating time and frequency sourced triggers.

2.6 INDUSTRIAL ENVIRONMENT EVENT CLASSIFICATION

Events such as cracking of pipes or pressure boundaries, pressure release valve (PRV) actuation, turbulence in working fluids, and motor whine or frequency shifts in pumps can each indicate facility condition. To

build a working dataset of relevant industrial sounds, the DAS system sampled a custom-built PRV (Figure 5a), a vacuum oven (Symphony Vacuum Oven, 414004-578, VWR, Radnor, Pennsylvania, Figure 5b), a laser engraver, and a leak checker. These classes were incorporated into the existing ESC-50 dataset and subjected to training and validation.



(a) PRV.



(b) Vacuum oven.

Figure 5. Fiber-optic sensing deployed on industrial valves to capture industrially relevant acoustic sources: (a) a PRV and (b) a vacuum oven.

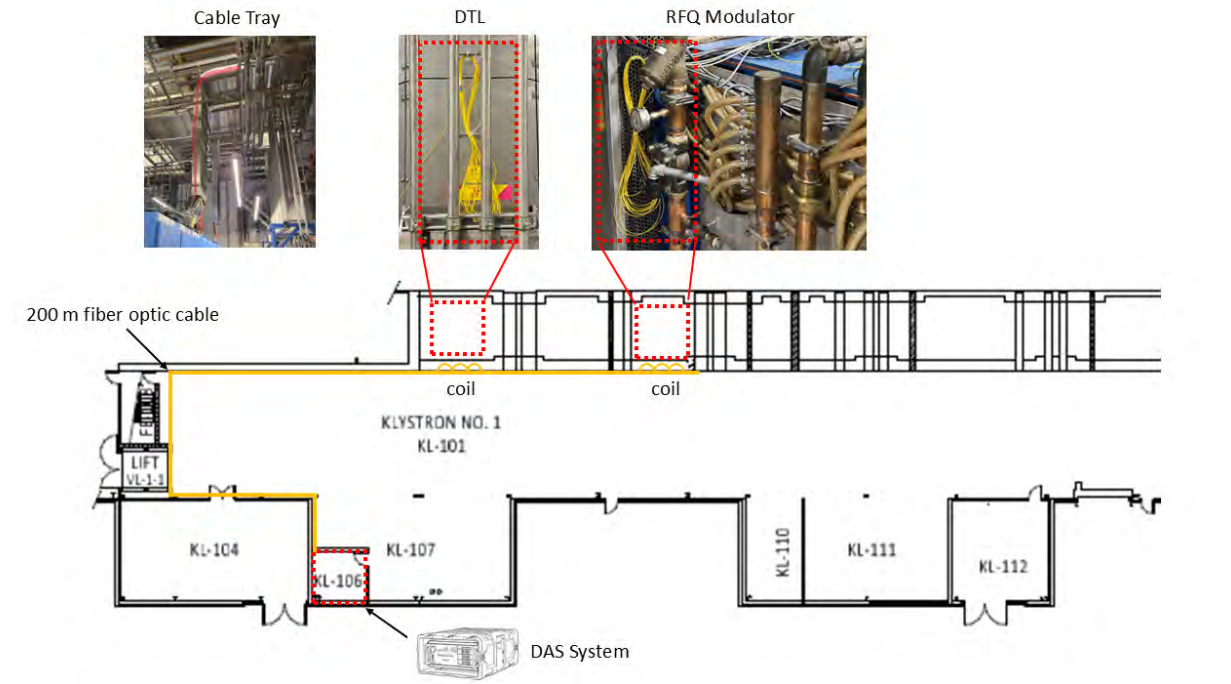
2.7 DEPLOYMENT AT THE SPALLATION NEUTRON SOURCE

To test the event classification pipeline in a relevant industrial environment, the DAS system was deployed in the SNS facility at Oak Ridge National Laboratory (ORNL). Three 20 m segments of a 200 m fiber were coiled and placed against equipment of interest in the klystron gallery at SNS, as shown in Figure 6. The DAS system remotely collected 7 days of acoustic data at 54.253 kHz (4.5 TB). A custom Python script was implemented to segment the data to generate spectrograms from 3 h slices. Changes to the acoustic signatures of the spectrograms are qualitatively assessed to demonstrate feasibility.

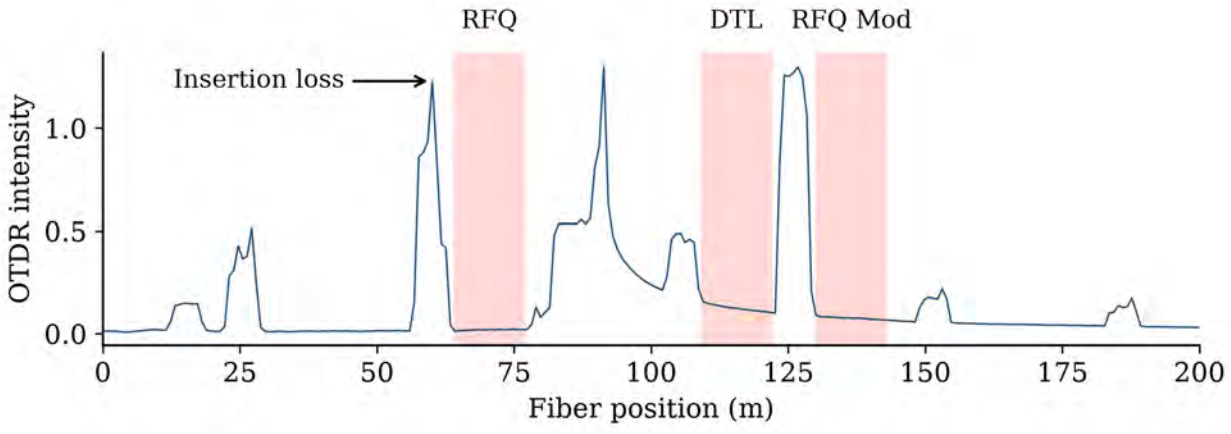
2.8 SNS FREQUENCY-BAND ENERGY ANALYSIS

Frequency band energy (FBE) was extracted from the spectrograms of the SNS facility by isolating the power spectrum in the region of interest (ROI), integrating across frequency, and plotting against our

previously established time frames. The ROI was determined by qualitative assessment of the spectrogram images. The power spectrum was plotted in decibels with a reference from each piece of equipment: the radio-frequency quadrupole (RFQ), the RFQ modulator, and the drift tube linac (DTL).



(a) Fiber-optic cable route.



(b) OTDR trace of SNS.

Figure 6. (a) DAS fiber-optic cable along equipment in the klystron building at SNS at ORNL. (b) Instantaneous OTDR trace of the fiber run.

3. RESULTS

3.1 EFFECTS OF CROSS-DOMAIN MODEL TRAINING ON VALIDATION ACCURACY

Table 2 summarizes the development and live model results. The initial deployment baseline, the replicated combined recipe, and the selected final model were all validated against the fixed 1,095-clip source-group partition, with the final model reaching 79.54% top-1 (true class predicted correctly) and 95.89% top-5 accuracy (true class predicted within top 5 guesses). The source-reference evaluation was performed offline on a 73-source subset with manual alignment to eliminate any offset introduced by triggering, yielding 84.93% top-1 accuracy. In the live testing, every scheduled class event was detected by the trigger function (100% recall), while classification accuracy on the live sync marker and live trigger sessions reached 71.23% and 73.97% top-1, respectively.

Table 2. Summary of model development and live testing results. Source-reference accuracy was measured before physical playback, live sync marker accuracy used marker-aligned DAS windows, and live trigger accuracy used windows produced by the live trigger. Event recall measures whether each scheduled class event was detected, independently of class-label accuracy.

Experiment	Evaluation axis	Result
Initial deployment baseline	Source-group validation	797/1,095 (72.79%) top-1
Replicated combined recipe	Source-group validation	78.17%, 78.54%, and 78.36% top-1
Selected final model	Source-group validation	871/1,095 (79.54%) top-1; 1,050/1,095 (95.89%) top-5
Source reference	Subset of source-group validation; validated after recording with manual alignment	62/73 (84.93%) top-1
Live event detection	Scheduled class events detected with triggering	73/73 detected; 100% recall
Live testing using sync marker	Validation sources in a new live session	52/73 (71.23%) top-1
Live testing using triggering	Validation sources in a new live session	54/73 (73.97%) top-1

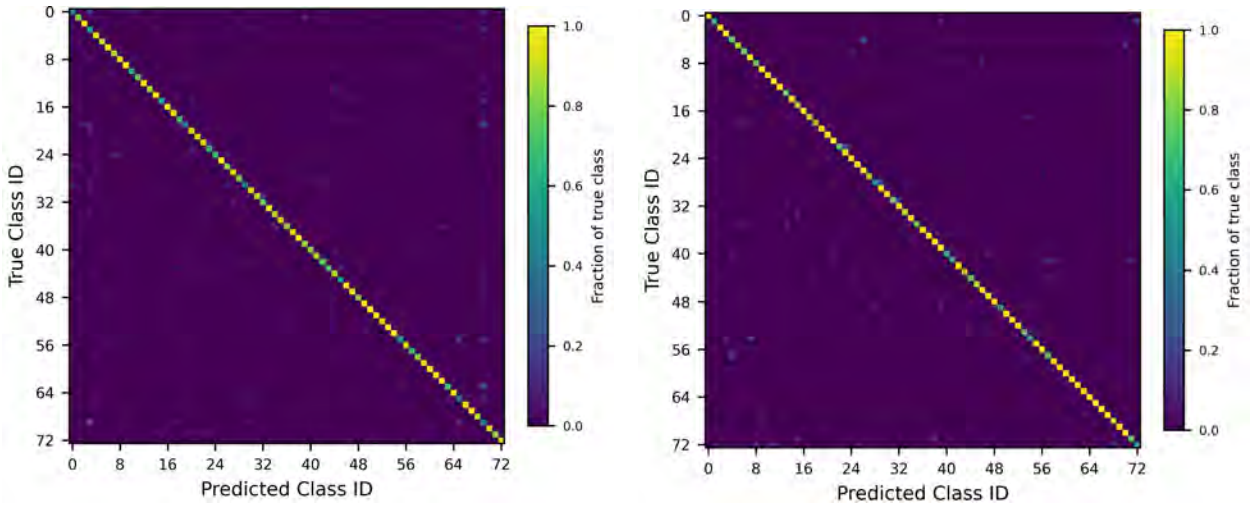
Table 3 reports validation accuracy for models trained on either source files or DAS recordings under different combinations of spectrogram augmentation and fine-tuning, with each model evaluated on both the source and DAS validation partitions to assess cross-domain transfer. Across all conditions, models perform strongly on the partition matching their training format but poorly on the mismatched format, indicating that acoustic features do not transfer directly between the two domains.

Table 3. Validation accuracy results for models trained with source files or DAS recordings under different SpecAugment and fine-tuning conditions. Domain refers to the medium used to produce the data on which the classifier was trained on.

Domain	SpecAugment	Fine-tuning	Source Val. Accuracy (%)	DAS Validation Accuracy (%)
Source	N	N	97.7	6.9
DAS	N	N	5.3	78.0
Source	Y	N	98.3	8.4
DAS	Y	N	10.5	84.8
Source	N	Y	25.1	76.1
DAS	N	Y	97.5	8.9
Source	Y	Y	31.7	84.4
DAS	Y	Y	98.1	16.8

Spectrogram augmentation improves performance for validation accuracy for matching format cases (Figure 7). For source-trained models without fine-tuning, adding spectrogram augmentation raised source validation accuracy from 97.7% to 98.3%, and for DAS-trained models, it raised DAS validation accuracy from 78.0% to 84.8%. The different sourced training format accuracy also improved, with source-trained models going from 6.9% to 8.4% on DAS, and DAS-trained models from 5.3% to 10.5% on source.

Fine-tuning shifts accuracy toward the target format at the expense of the original training format. For source-trained models without augmentation, fine-tuning raised DAS accuracy from 6.9% to 76.1% while dropping source accuracy from 97.7% to 25.1%. The corresponding DAS-trained models saw source accuracy rise from 5.3% to 97.5% and DAS accuracy fall from 78.0% to 8.9%. The source-trained model with spectrogram augmentation and DAS fine-tuning reached 84.4% on DAS while falling to 31.7% on source validation. The DAS-trained model with augmentation and source fine-tuning reached 98.1% on source while falling to 16.8% on DAS-derived spectrograms.



(a) **(b)**

Figure 7. Confusion matrices with spectrogram augmentation from same-domain validation of the DAS file trained model validated with (a) DAS recorded data and (b) the source recording trained model validated against source files.

3.2 FIBER DEPLOYMENT AND GAUGE LENGTH IMPACT

A same-session comparison was performed using a single recording session on a combined length of fiber and evaluated offline across four gauge lengths (1, 2, 5, and 10 m) 4. The spooled condition produced identical accuracy at every gauge length, indicating that gauge length had no measurable effect on classification for that topology. The straight condition varied only marginally across gauge lengths (within roughly 1.4% top-1).

The spooled topology outperformed the straight topology on top-1 accuracy. In the one-point condition, spooled fiber achieved 58.90% top-1 versus 43.84%–45.21% for straight fiber. The three-point condition showed the same trend, with spooled fiber at 52.97% top-1 compared with 43.84%–44.75% for straight fiber.

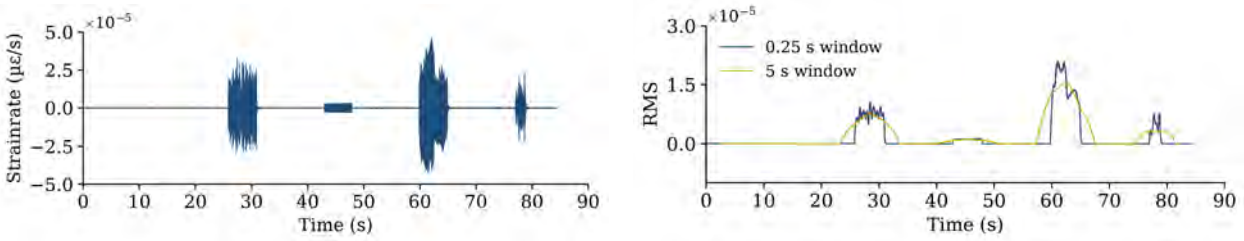
Top-5 accuracy was substantially less sensitive to topology. In the one-point condition, spooled fiber reached 82.19% top-5 versus 80.82%–82.19% for straight fiber, and in the three-point condition, 79.45% versus 79.91%–81.28%. The principal benefit of the spooled fiber therefore appears to be improved top-1 ranking rather than increased inclusion of the correct class among the top-5 candidates; in the three-point case, the straight topology actually edged out the spooled topology on top-5.

Table 4. Comparison of classification model accuracy based on gauge length and fiber layout topology

Topology	Source Count	Gauge (m)	Top 1 (%)	Top 5 (%)
spooled	1	1	58.90	82.19
		2	58.90	82.19
		5	58.90	82.19
		10	58.90	82.19
	3	1	52.97	79.45
		2	52.97	79.45
		5	52.97	79.45
		10	52.97	79.45
straight	1	1	45.21	80.82
		2	43.84	82.19
		5	43.84	82.19
		10	45.21	80.82
	3	1	44.29	79.91
		2	43.84	80.37
		5	44.29	80.37
		10	44.75	81.28

3.3 DEVELOPMENT OF AN ANOMALY DETECTION PIPELINE FOR DAS MONITORING

Windowing the high-rate data stream with an RMS function is useful to minimize the influence of transient spikes, or outliers, on our data processing. As shown in Figure 8(a), the raw strain rate series contains sharp, high-amplitude transients against a low-amplitude baseline. Figure 8(b) shows that a 0.25 s RMS window preserves these peaks with tight temporal localization, whereas a 5 s window suppresses their amplitude and spreads their energy across a wider span in time. This behavior implies that the window should be matched to the temporal scale of the acoustic signature of interest. For nonstationary noise that persists over longer durations, a longer RMS window is acceptable and can improve robustness by averaging over the event. For brief, stationary noise, a shorter RMS window is preferable, as it retains



(a) **(b)**

Figure 8. Effect of RMS window length on the strain rate envelope. (a) Raw strain rate time series. (b) Running RMS for a 0.25 s (blue) and 5 s (orange) window, showing peak suppression and temporal spreading at the longer window.

the peak amplitude and temporal localization needed for the event to remain distinguishable above the baseline.

Figure 9 compares two strategies for calibrating the modified z-score reference values on stitched DAS RMS data. When the median and MAD are calculated from a fixed 20 s reference window, the MAD calculation reflects only the artificially quiet background noise of the lab environment (MAD = 1.83×10^{-9}), so subsequent events are inflated. When the reference is instead computed over the entire stitched signal, the transient events are themselves included in the statistics that define the baseline. The median is essentially unchanged (7.14×10^{-8} versus 7.79×10^{-8}), but the global MAD rises to 4.38×10^{-9} (compared to 1.83×10^{-9} for the artificially quiet value). The MAD value with the global reference lowers the peak z-values recorded for the same physical events and narrows their separation from the baseline. Consequently, a fixed background reference yields higher-amplitude z-score peaks. The advantage of the moving global reference is that it absorbs nonstationary events into its own calibration, systematically correcting for gradual changes.

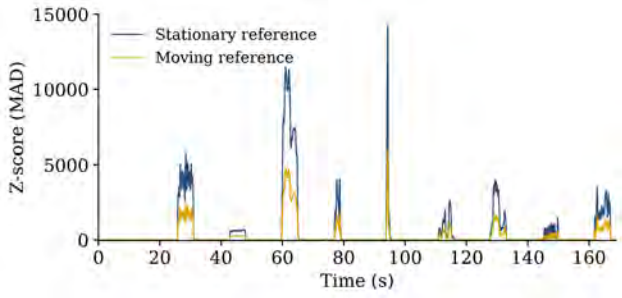


Figure 9. MAD-based modified z-score computed on two stitched DAS clips with a 0.25 s RMS window, comparing two reference-frame choices on a single axis. The fixed 20 s reference (blue) calibrates the median and MAD on a background window at the start of the record, while the moving reference (orange) calibrates them over the entire stitched signal.

When calibrating the median and MAD over the entire stitched record rather than a fixed background window, the trigger achieved a best F1 of 0.679 at a z-score threshold of 69.8, recovering 19 of the 20 manifest events (1 FN) while producing 17 FPs among 36 total detections. Figure 10 shows that at low thresholds, the detector recovers nearly all events at the cost of numerous spurious detections, while raising the threshold suppresses FPs but eventually begins to miss true events.

It is of interest to study how low-frequency background impacts our event detection method, as it is

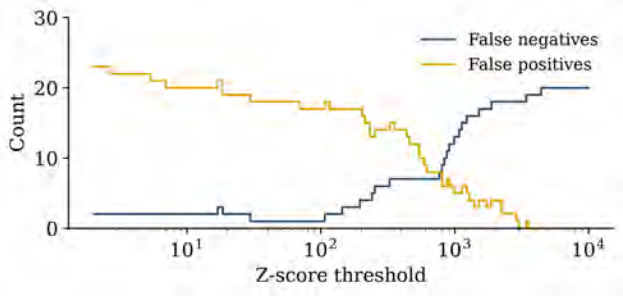
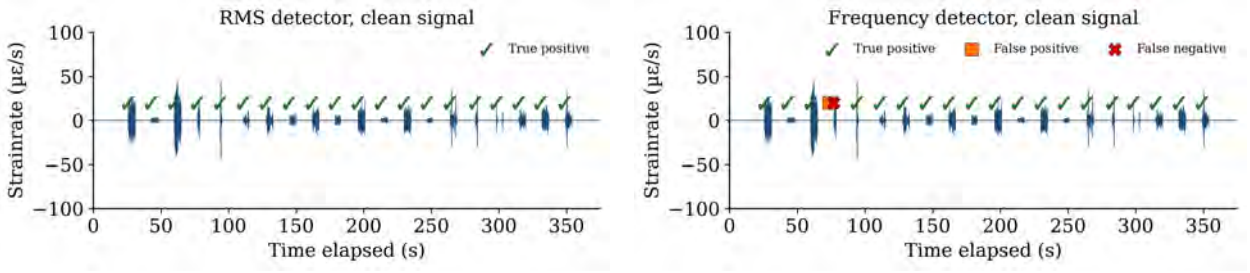


Figure 10. FNs (blue) and FPs (orange) versus z-score threshold for the DAS trigger detector, scored against the manifest with a ± 1 s match tolerance. FNs count missed manifest events; FPs count detections not matching any event.



(a) **(b)**

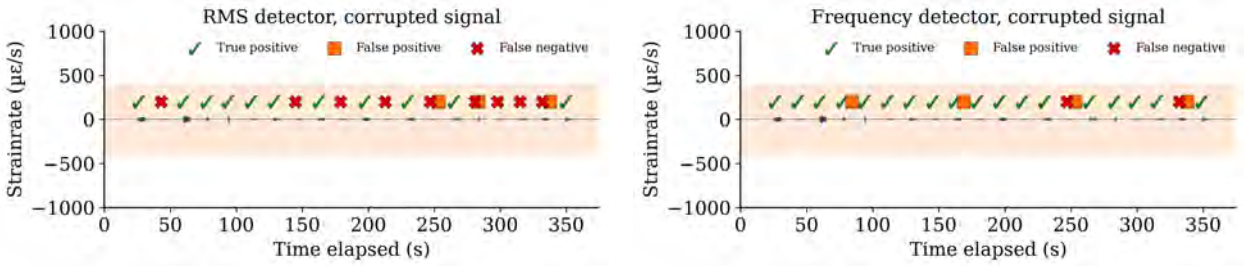
Figure 11. Event detection on the clean DAS record. (a) Time-amplitude detector: F1 = 1.000 (20 TPs, 0 FPs, 0 FNs). (b) Mel spectrogram detector: F1 = 0.950 (19 TPs, 1 FP, 1 FN). Green checks mark TPs, orange squares mark FPs, and red crosses mark FNs.

likely that any industrial environment will have noise. We expanded our trigger algorithm to include the frequency domain by transforming the strain rate signals to mel spectrogram representations. Each mel bin is subject to the MAD algorithm, and if any bin is beyond the threshold, then it is considered a trigger. On the clean signal, the two detectors performed comparably; the time-amplitude detector achieved an F1 of 1.00 (11a), while the mel spectrogram detector achieved an F1 of 0.950, missing one true event (11b). The 100 Hz tone was introduced to highlight how a low-frequency noise signal, relevant in an industrial setting, will skew the RMS amplitude based trigger detector. In this context, the time-amplitude detector was not able to trigger on relevant events, with its F1 falling to 0.647, missing 9 events (12). The mel spectrogram detector also degraded, but less substantially (12b), retaining an F1 of 0.857 and recovering 18 of the 20 events (6 FNs).

3.4 EXTENDED DAS DATASET CLASSIFICATION WITH INDUSTRIAL SOUNDS

Figure 13 shows the mel spectrograms of the furnace acoustic signature across the six operating states. Qualitatively, with the pump off, the low-frequency band carries little energy, whereas engaging the pump introduces pronounced low-frequency content that persists across the record. These qualitative differences indicate that DAS can capture acoustic features that vary with operational state, suggesting it as a promising basis for automated state classification.

The acoustic signatures of a vacuum oven, a laser engraver, and a leak checker were incorporated into the existing ESC-50 dataset and subjected to training and validation, yielding a final validation accuracy of 77.3%, comparable to that reported for the original dataset.

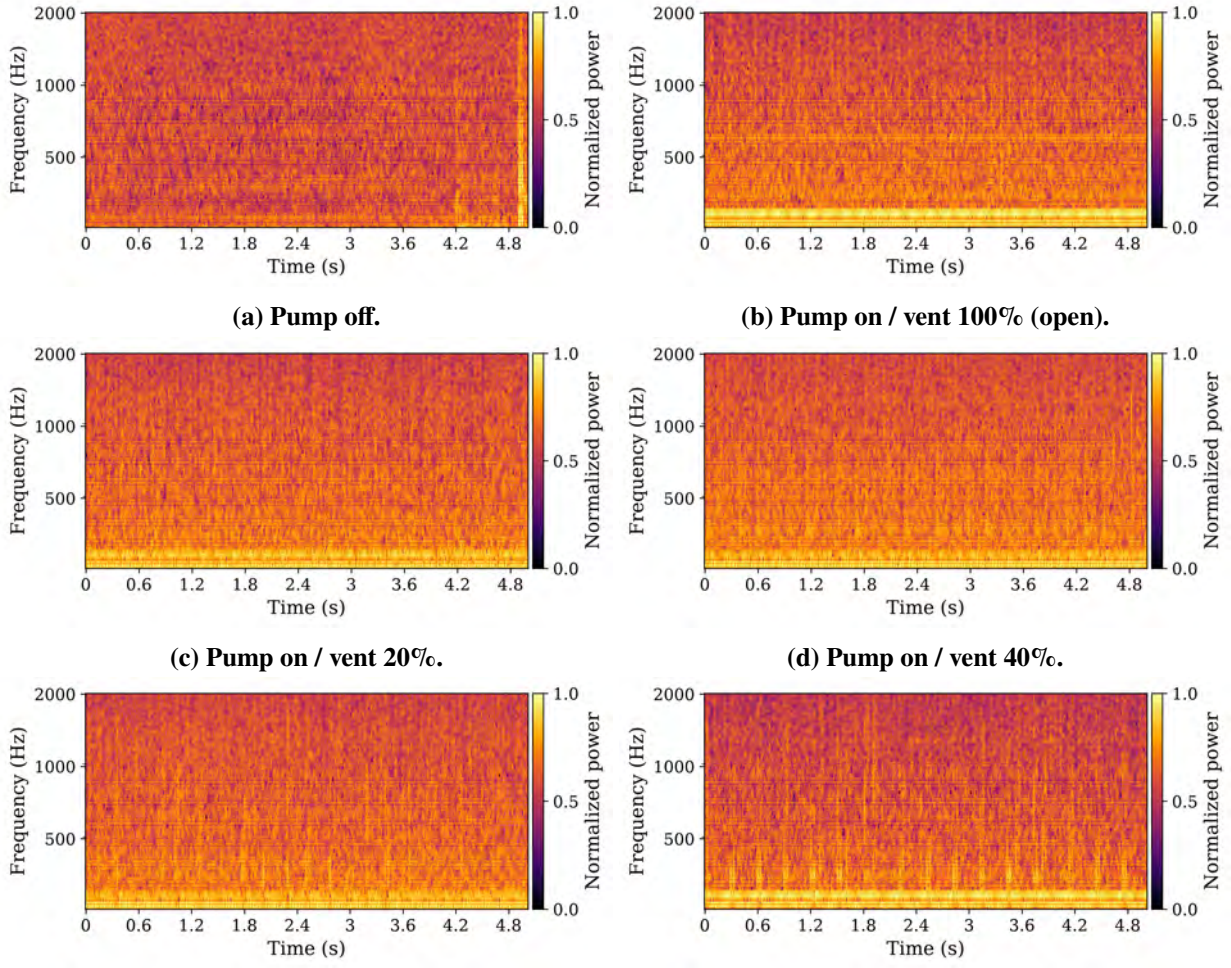


(a) **(b)**

Figure 12. Event detection on the corrupted DAS **record (100 Hz tone added).** (a) Time-amplitude detector: F1 = 0.647 (11 TPs, 3 FPs, 9 FNs). (b) Mel spectrogram detector: F1 = 0.857 (18 TPs, 4 FPs, 2 FNs). Green checks mark TPs, orange squares mark FPs, and red crosses mark FNs. The noisy signal is in pale color and the original signal is overlaid in blue for reference.

3.5 PROOF-OF-CONCEPT DEPLOYMENT AT THE SPALLATION NEUTRON SOURCE

The deployment of the DAS system at the SNS facility demonstrated proof of concept of our pipeline. The cyclic patterns on the spectrogram could indicate various on/off cycles of supplementary equipment, such as fans (14). These data could be implemented in an object detection model if ground truth state data were collected and used to train the model. The SNS deployment validates several properties critical to a remote MR monitoring scenario, such as unattended operation over a multiday period, tolerance to the electromagnetic environment of an accelerator facility, and the ability to resolve operationally relevant acoustic signatures. For the frequency band energy, there is a 180° shift between equipment (15), indicating a correlation between operational states (i.e., one turns off when the other turns on). For MR applications, where physical access is infrequent and operator staffing is minimal, these characteristics are particularly valuable. Future work should focus on collecting labeled ground truth data in a representative nuclear environment (i.e., a research reactor). Integration of the classification pipeline with existing control and safety systems would represent a step toward a deployable remote monitoring solution.



(a) Pump off.

(b) Pump on / vent 100% (open).

(c) Pump on / vent 20%.

(d) Pump on / vent 40%.

(e) Pump on / vent 60%.

(f) Pump on / vent 0% (closed).

Figure 13. Mel spectrograms of acoustic signature across six operating states of the pressure vacuum.

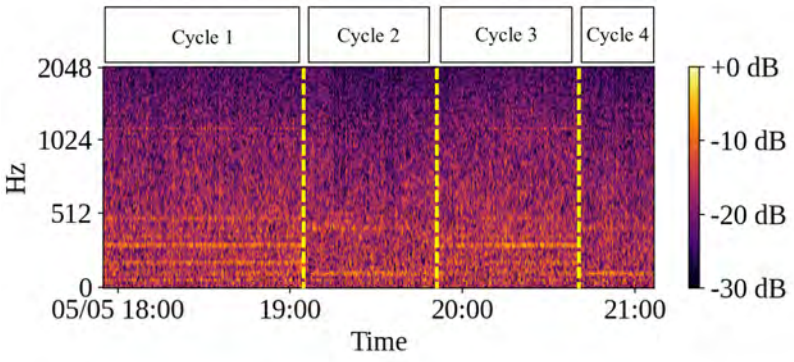


Figure 14. DAS-derived spectrogram of auxiliary machinery.

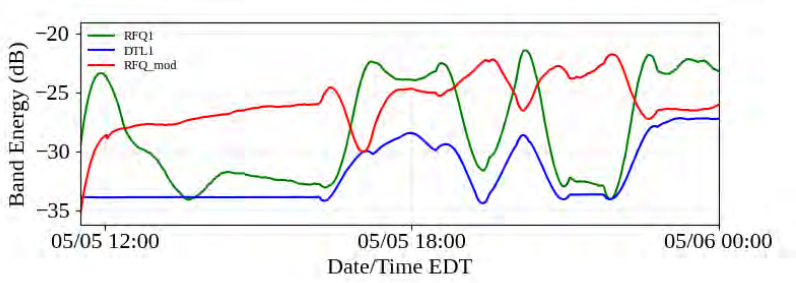


Figure 15. FBE from spectrograms measured in the klystron gallery of SNS.

4. DISCUSSION

MRs are attractive because they are modular and small. To make them useful, maintenance and control must be minimized. Combining continuous anomaly detection with acoustic classification offers a route to flagging (and identifying) off-normal conditions early and remotely, which could reduce O&M costs and thereby ultimately advance the widespread adoption of MRs.

Audio classification ML models and anomaly detection algorithms are well researched in the context of system health monitoring [17]. Wireless sensors that continuously monitor industrial conditions and report off-normal events are commercially available [18]. This real-time, remote, autonomous monitoring is advantageous because it moves toward decreasing staffed O&M responsibility and costs. Furthermore, reactor downtime is a critical limiting factor in the economic feasibility of MRs. In unanticipated downtime, the lack of predictive maintenance results in surprise faults and unknown time to restoration. However, even with preventative maintenance and scheduled downtime, this is not an optimized economic system, as there is too often no necessary maintenance required by the system. Thus, a remote sensing system is crucial for optimized predictive and preventative maintenance for nuclear systems. Because of electromagnetic noise, radiation [29], and temperature [11], fiber-optic sensing is an advantageous hardware instrumentation option compared to traditional sensors. Thus, the key advancement of this work was to establish the feasibility of applying well-established ML models, such as the YOLO object recognition model [15], to DAS-derived vibroacoustic data. To the best of our knowledge, there is no existing model that classifies DAS data in the context of industrial system health monitoring, specifically for MRs. Because DAS data have a unique noise signature compared to traditional microphonic data, we had to develop a novel, end-to-end pipeline for our event classification model.

As there is not a readily available dataset for industrial sound signatures recorded by DAS devices, it was of interest to investigate whether a model trained on microphonic data could be used to classify DAS data and vice versa. The results presented here indicate that without advanced signal processing techniques, it is not viable to use variable source data for event classification. However, models trained on DAS data performed well when validated against unseen DAS data (Table 3). Thus, we also attempted to augment our training to identify routes for cross-source classification. Spectrogram augmentation did not prove to be an effective way to make microphonic sourced data usable for DAS validation. Fine-tuning our source-derived weights with DAS spectrograms was an effective approach to accurately classify unseen DAS data. The results presented here indicate that established ML models can be used for classification of DAS data. Based on the experiments conducted in this study, it is not feasible to classify DAS data without a DAS training dataset. This is an important future direction to investigate, as there are more open-source industrial sound signature datasets available recorded by microphone compared to DAS hardware. The results in Table 3 show a mild improvement in cross-domain validation accuracy when spectrogram augmentation is applied. This indicates there may be advanced signal processing and ML techniques that could be eventually employed to integrate microphonic data with DAS data.

Accuracy depended on where an event sat on the fiber and in particular differed between straight and spooled sections because spooling changes how strain couples into the fiber and how adjacent channels overlap spatially. The gauge length sets the spatial averaging of each DAS channel and trades signal-to-noise against spatial resolution. These results should not necessarily be interpreted solely as topology-caused. The straight and spooled measurements occurred at different absolute distances along the fiber, so factors such as attenuation and coupling could contribute as well. Optimizing gauge length for the deployment geometry, and accounting for spooled versus straight runs when laying out the fiber, could improve both classification and anomaly detection.

In the context of O&M, the FN and FP error modes are not equally costly. An FN, where a true event is missed, represents a permanent loss of information about the monitored system, potentially masking the onset of a developing fault. An FP, by contrast, merely records an additional data packet that later can be ignored or deleted. The trigger purpose is to preserve events for later inspection, and the cost of overrecording is storage, while the cost of underrecording is a missed fault. A straightforward way to exploit this asymmetry is to have the trigger return the underlying strain rate data packet alongside each detection so that a human operator, or a downstream classifier, can resolve whether the recorded excursion was a true event or an erroneous trigger. The z-score threshold is most useful as a user-facing sensitivity control rather than a fixed parameter. Lowering the threshold slides the detector toward high sensitivity, while raising it will only capture major events. Because the appropriate balance depends on the tolerance for missed events at a given facility, giving the threshold as an adjustable sensitivity setting lets the operator position the detector along this trade-off curve to suit their context, rather than committing to the single F1-optimal point identified here.

An important extension of the temporal anomaly detection is the spatial location of the source event along the fiber. A single physical event may couple into many channels at once, so a simple per-channel trigger would fire redundantly across the array. If the event signature is similar in its frequency content across channels but differs in amplitude, the spatial origin can be localized by identifying the channel of maximum response rather than triggering everywhere the signature appears. Examining the frequency components channel by channel would allow the detector to confirm that a distributed set of responses shares a common spectral response, while using the amplitude distribution along the fiber to pinpoint the location of highest coupling. This would suppress redundant triggers and add a spatial dimension to the anomaly report.

Using a time-amplitude RMS detector is effective when stationary events occur with low background noise. However, because the detector sums energy across the spectrum, it cannot distinguish changes in frequency. To artificially create a noisy industrial background environment, a 100 Hz tone was introduced. This caused the RMS amplitude z-score detector to fail. A spectral representation such as the mel spectrogram addresses this directly by preserving the frequency axis that the RMS detector discards. Distributing energy across many mel bins and triggering on any bin that exceeds the threshold localizes the interference more effectively than the time-amplitude trigger. In practice, this suggests a layered strategy. The RMS detector suffices for events greater than the background noise, while the spectrogram-based detector is warranted wherever persistent narrowband noise is expected, which, in most operating industrial settings, is common.

The SNS deployment demonstrated that the full chain functions in an operational accelerator environment. As a proof of concept, the deployment validated the approach end to end; however, SNS is not a reactor, and the results should be read as evidence of feasibility rather than as a completed facility monitoring benchmark.

5. CONCLUSIONS

This work establishes the feasibility of applying off-the-shelf ML models and statistical anomaly detection methods to DAS-derived vibroacoustic data for the remote, autonomous health monitoring of MRs. Using a novel end-to-end pipeline based on the ML model, we showed that classifiers trained on DAS data accurately classify unseen DAS data. Cross-source classification between microphonic and DAS data was not viable without adaptation. Spectrogram augmentation proved ineffective, whereas fine-tuning source-derived weights with DAS spectrograms was an effective route to accurate classification. This indicates that a DAS training dataset is necessary for accurately classifying DAS spectrograms. Developing advanced signal processing techniques to make open-source microphone-recorded industrial sound data usable for DAS classification would be an ambitious and well worthwhile avenue to explore. Alternatively, building such open-source datasets with DAS data is a more straightforward but labor intensive method. For anomaly detection, a time-amplitude RMS z-score trigger is effective for stationary events above a quiet background but fails in the presence of persistent narrowband interference. A mel spectrogram detector addresses this by preserving the frequency axis. The z-score threshold is best used as an operator-facing sensitivity control, allowing each facility to position the detector according to its tolerance for missed events. Classification accuracy depended on event location and on spooled versus straight fiber sections, identifying gauge length as a key tunable parameter for future deployments. Additional natural extensions include spatially localizing anomalies along the fiber by exploiting shared spectral signatures and channel-wise amplitude distributions to suppress redundant triggers. Finally, the SNS deployment demonstrated that the full chain functions in an operational accelerator environment. Although SNS is not a reactor and these results should be read as evidence of feasibility rather than a completed facility monitoring benchmark, they collectively establish DAS-based acoustic monitoring as a promising path toward reduced O&M costs and the wider adoption of MRs.

REFERENCES

[1] R. Testoni, A. Bersano, and S. Segantin, "Review of nuclear microreactors: Status, potentialities and challenges," *Progress in Nuclear Energy*, vol. 138, p. 103 822, Aug. 1, 2021. DOI: 10.1016/j.pnucene.2021.103822.

[2] H. C. Bryan, K. W. Jesse, C. A. Miller, and J. M. Browning, "Remote nuclear microreactors: A preliminary economic evaluation of digital twins and centralized offsite control," *Frontiers in Nuclear Engineering*, vol. 2, Dec. 21, 2023. DOI: 10.3389/fnuen.2023.1293908.

[3] K. Stevens, J. Oncken, R. Boring, *et al.*, "Opportunities, challenges, and research needs for remote microreactor operations," *Nuclear Technology*, vol. 210, no. 12, pp. 2257–2273, 2024. DOI: 10.1080/00295450.2024.2344903.

[4] I. Naranjo de Candido, A. Al Rashdan, A. Abou Jaoude, and J. Buongiorno, "Assessment of technoeconomic opportunities in automation for nuclear microreactors," *Nuclear Science and Engineering*, 2024. DOI: 10.1080/00295639.2024.2372511.

[5] F. Peng, N. Duan, Y.-J. Rao, and J. Li, "Real-time position and speed monitoring of trains using phase-sensitive OTDR," *IEEE Photonics Technology Letters*, vol. 26, no. 20, pp. 2055–2057, Oct. 2014. DOI: 10.1109/LPT.2014.2346760.

[6] F. Tanimola and D. Hill, "Distributed fibre optic sensors for pipeline protection," *Journal of Natural Gas Science and Engineering*, vol. 1, no. 4, pp. 134–143, Nov. 1, 2009. DOI: 10.1016/j.jngse.2009.08.002.

[7] H. Liu, J. Ma, W. Yan, W. Liu, X. Zhang, and C. Li, "Traffic flow detection using distributed fiber optic acoustic sensing," *IEEE Access*, vol. 6, pp. 68 968–68 980, 2018. DOI: 10.1109/ACCESS.2018.2868418.

[8] T. Xie, B. Shi, C.-C. Zhang, *et al.*, "Distributed acoustic sensing (DAS) for geomechanics characterization: A concise review," *IOP Conference Series: Earth and Environmental Science*, vol. 861, no. 4, p. 042 033, 2021. DOI: 10.1088/1755-1315/861/4/042033.

[9] A. Xenaki, P. Gerstoft, E. Williams, and S. Abadi, "Overview of distributed acoustic sensing: Theory and ocean applications," *The Journal of the Acoustical Society of America*, vol. 158, no. 1, pp. 801–825, Jul. 2025. DOI: 10.1121/10.0037218.

[10] D. C. Sweeney and C. M. Petrie, "Minimization of measurement uncertainty in optical frequency domain reflectometry," *Journal of Lightwave Technology*, vol. 43, no. 6, pp. 2972–2981, Mar. 2025. DOI: 10.1109/JLT.2024.3507004.

[11] A. Birri, D. C. Sweeney, H. C. Hyer, B. Schreiber, E. Cakmak, and C. M. Petrie, "A miniaturized, high-bandwidth optical fiber fabry–perot cavity vibration sensor demonstrated up to 800 °c," *IEEE Sensors Journal*, vol. 25, no. 7, pp. 11 082–11 091, Apr. 2025. DOI: 10.1109/JSEN.2025.3541557.

[12] I. McLoughlin, L. Pham, Y. Song, *et al.*, "Spectrogram features for audio and speech analysis," *Applied Sciences*, vol. 16, no. 2, 2026. DOI: 10.3390/app16020572.

[13] T. Zhang, G. Feng, J. Liang, and T. An, "Acoustic scene classification based on mel spectrogram decomposition and model merging," *Applied Acoustics*, vol. 182, p. 108 258, 2021. DOI: https://doi.org/10.1016/j.apacoust.2021.108258.

[14] K. Zaman, M. Sah, C. Direkoglu, and M. Unoki, "A survey of audio classification using deep learning," *IEEE Access*, vol. 11, pp. 106 620–106 649, 2023. DOI: 10.1109/ACCESS.2023.3318015.

[15] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788. DOI: 10.1109/CVPR.2016.91.

[16] A. Patil, G. Soni, and A. Prakash, "Data-driven approaches for impending fault detection of industrial systems: A review," *International Journal of System Assurance Engineering and Management*, vol. 15, no. 4, pp. 1326–1344, Apr. 2024. DOI: 10.1007/s13198-022-01841-9.

[17] A. Mesaros, R. Serizel, T. Heittola, T. Virtanen, and M. D. Plumbley, "A decade of dcase: Achievements, practices, evaluations and future challenges," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5. DOI: 10.1109/ICASSP49660.2025.10887673.

[18] R. Giri, S. V. Tenneti, F. Cheng, K. Helwani, U. Isik, and A. Krishnaswamy, "Self-supervised classification for detecting anomalous sounds," in *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2020 Workshop (DCASE2020)*, Tokyo, Japan, Nov. 2020, pp. 46–50.

[19] X. Yu, Y. Fu, W. Shen, *et al.*, "Edge intelligence-based acoustic trigger sensing for real-time monitoring of train-crossing events," *Measurement*, vol. 258, p. 119 308, 2026. DOI: https://doi.org/10.1016/j.measurement.2025.119308.

[20] N. Issa, M. A. F. ROELENS, and S. J. Frisken, *Distributed optical sensing systems and methods*, US Patent 11,815,373, Nov. 2023.

[21] A. Alekseev, V. Vdovenko, B. Gorshkov, V. Potapov, and D. Simikin, "A phase-sensitive optical time-domain reflectometer with dual-pulse phase modulated probe signal," *Laser Physics*, vol. 24, no. 11, p. 115 106, 2014.

[22] K. J. Piczak, "ESC: Dataset for environmental sound classification," in *Proceedings of the 23rd ACM international conference on Multimedia*, ser. MM '15, New York, NY, USA: Association for Computing Machinery, Oct. 13, 2015, pp. 1015–1018. DOI: 10.1145/2733373.2806390.

[23] G. Jocher, J. Qiu, M. Liu, S. Lyu, F. C. Akyon, and M. E. Kalfaoglu, *Ultralytics YOLO26: Unified real-time end-to-end vision models*, Version Number: 1, 2026. DOI: 10.48550/ARXIV.2606.03748.

[24] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2019.

[25] M. Hira, *Audio feature augmentation*, https://docs.pytorch.org/audio/main/tutorials/audio_feature_augmentation_tutorial.html, TorchAudio 2.8.0 documentation, accessed July 28, 2026, PyTorch, 2026.

[26] M. Withers, R. Aster, C. Young, *et al.*, "A comparison of select trigger algorithms for automated global seismic phase and event detection," *Bulletin of the Seismological Society of America*, vol. 88, no. 1, pp. 95–106, 1998. DOI: 10.1785/BSSA0880010095.

[27] H. Azimi, M. Bouchard, R. A. Goubran, and F. Knoefel, "Apnea event detection methodology using pressure sensors," in *2019 IEEE International Symposium on Medical Measurements and Applications (MeMeA)*, 2019, pp. 1–6. DOI: 10.1109/MeMeA.2019.8802214.

[28] G. Princz, M. Shaloo, and S. Erol, "Anomaly detection in binary time series data: An unsupervised machine learning approach for condition monitoring," *Procedia Computer Science*, vol. 232, pp. 1065–1078, 2024, 5th International Conference on Industry 4.0 and Smart Manufacturing (ISM 2023). DOI: https://doi.org/10.1016/j.procs.2024.01.105.

[29] B. Conry, Y.-R. Lin, J. M. Kurley, *et al.*, "Testing and characterization to develop a mechanistic explanation for unsaturated drift of fiber optic sensors during high-dose irradiation," Oak Ridge National Laboratory (ORNL), Technical Report, Sep. 2024, OSTI ID: 2589799. DOI: 10.2172/2589799.