Key Points:

- Seismic waves from uncontrolled sources, such as natural and induced seismicity, generate waves that contain information related to the fractures through which they pass
- Complex waveforms arising from fracture sets are difficult to analyze when the fracture spacing is small relative to the length of an event wave packet
- A twin neural network identifies changes in signals, generated by uncontrolled sources that indicate associated changes in saturation among fractures in a set

Correspondence to:
L. J. Pyrak-Nolte,
ljpn@purdue.edu

Monitoring Fracture Saturation With Internal Seismic Sources and Twin Neural Networks

D. D. Nolte[1] and **L. J. Pyrak-Nolte[1,2,3]**

[1]Department of Physics and Astronomy, Purdue University, West Lafayette, IN, USA, [2]Department of Earth, Atmospheric and Planetary Sciences, Purdue University, West Lafayette, IN, USA, [3]Lyles School of Civil Engineering, Purdue University, West Lafayette, IN, USA

Abstract Seismic coda-wave analysis is a well-developed method for detecting subtle physical changes in complex media by measuring arrival times in the late-arriving energy from multiply scattered or reflected waves. However, a challenge arises when multiply scattered waves are not sufficiently separated in time from the direct arrivals to provide a clear coda wave train. Additional complications for monitoring changes in fracture systems arise when the signals originate from unsynchronized internal sources, such as natural or induced seismicity, from acoustic emission, or from transportable intra-fracture sources (chattering dust), that generate uncontrolled signals that vary in arrival time, amplitude and frequency content. Here, we use a twin neural network (TNN also known as a Siamese neural network) for dimensionality reduction to analyze signals from chattering dust to classify the fluid saturation state of a synthetic fracture system. The TNN with shared weights generates a low-dimensional representation of the data input by minimizing contrastive loss, serving as the input to a multiclass classifier that accurately classifies whether multiple fractures in a fracture system are fully saturated or partially saturated, or whether a change in saturation has occurred in different fractures in the system. These results show that information buried in unresolved codas from uncontrolled sources can be extracted using machine learning to monitor the evolution of fracture systems caused by physical and chemical processes even when the scattered and direct wave fields overlap.

Plain Language Summary Earthquakes and human activities in the Earth's subsurface generate internal seismic waves that travel from wave sources to receivers. Along their paths, the waves interact with fractures that scatter the waves and cause them to partially interfere, producing complicated signals that carry information on the geometry and the physical conditions of the fractures. Extracting this information could give important insight into environmental safety of underground engineering projects affected by fracture sets. In this work, we use the unique properties of a twin neural network (TNN) to disentangle fracture information buried in complicated signals. The TNN is a dimensionality reduction technique and multiclass classifier that decides whether a fracture system is fully saturated with water or partially saturated, or whether a change in saturation has occurred in different fractures in the system. These results show that information about physical and chemical processes in fractures can be extracted from complex wave fields.

1. Introduction

Seismic coda-wave analysis is a well-developed approach for detecting changes in complex media by taking advantage of the information in multiply scattered signals as they propagate through a material system. Active-source monitoring using coda waves can detect changes in velocity associated with alterations from natural and human activities. Active sources have the advantage of generating signals that have repeatable source characteristics when looking for time shifts to extract changes in velocity of a region that can be linked to time-dependent alterations. For example, a recent study showed that active-source monitoring of coda waves can be used to track the location of a CO_2 plume as it was injected over a 74 hr period into the subsurface (Zhu et al., 2019). Coda wave data were acquired using a continuous active source with known and repeatable source characteristics that identified changes in velocity as small as 0.2%–2% from different components of the waveform.

Alternatively, passive-source monitoring using coda waves from earthquakes or acoustic emissions can also be used to detect such changes. Natural earthquakes and induced seismicity from subsurface activities (such as mining and geothermal energy exploitation) generate elastic waves that have been used to identify changes in the Earth's subsurface located between the source of the seismicity and the receiver location (e.g., Boue et al., 2014;

Olivier et al., 2015; Peng & Ben-Zion, 2006). However, a challenge in such passive-source monitoring is the lack of repeatability of the source characteristics because of variability in the triggering mechanisms (e.g., fluid pressures, stress distributions, tectonic settings, creep, loading rate, etc.), as well as source location and rupture mechanisms (shear vs. tensile, multiple slip events, etc.). Depending on the mechanism of the source of seismicity, sequential events may not have wave form similarity (Uchida & Burgmann, 2019) with consistent frequency content and phase components. In addition, Geller & Mueller (Geller & Mueller, 1980) showed that deviation by as little as a quarter-wavelength in path length reduces waveform similarity which could occur while inducing a fracture. These aspects create a challenge for passive monitoring of the material between the source and receiver when working with non-repeatable or uncontrolled sources that may repeat randomly in time and with varying characteristics (e.g., amplitude, spectral content, and phases).

The challenge of such passive monitoring is made even more difficult when working in fractured media where the strength of coda signals produced from sets of parallel fractures depends on the conditions of the fractures in the sets. Laboratory experiments with controlled ultrasonic sources and numerical simulations have demonstrated that coda patterns can be used to monitor cracking and to identify the orientation of fracture sets, fracture spacing and fracture density (e.g., de Figueiredo et al., 2012; Hadziioannou et al., 2009; Nakagawa et al., 2000b, 2003, 2004; Sang et al., 2020; Sneider et al., 2013; Willis et al., 2006]). However, the presence of reverberations and their amplitudes depend not only on the fracture spacing but also on fluid saturation and fracture specific stiffness.

Fracture specific stiffness is an effective parameter that captures the deformed morphology of a fracture and is intimately related to the contact area and the aperture distribution (Petrovitch et al., 2013, 2014; Pyrak-Nolte & Nolte, 2016). The amount of energy partitioned into transmitted and reflected components, body-mode conversion, and guided waves, as well as the time delays of different scattered components, depend on fracture specific stiffness (Shao & Pyrak-Nolte, 2016; Shao et al., 2015; Xian et al., 2001a). As a fracture is partially closed under stress, the increase in contact area and decrease in fracture aperture increases the fracture specific stiffness (Hopkins, 1990; Hopkins et al., 1990; Pyrak-Nolte, 2019) and in turn increases wave transmission with a decrease in reflections, a decrease in arrival time, the transmission of higher frequency components and a decrease in compressional (P) to shear (S) wave conversions and guided modes (Choi et al., 2014; Lubbe et al., 2008; Lubbe & Worthington, 2006; Nakagawa et al., 2004; Pyrak-Nolte et al., 1990; Schoenberg, 1980; Xian et al., 2001b). Fluid saturation also suppresses the ringing coda-wave tail because saturation decreases the reflections from the fracture (Pyrak-Nolte, 2007; Xian et al., 2001b). In addition, the time delay produced by a fracture depends on the frequency of the signal, on the number of fractures, on fracture spacing and on the fracture specific stiffness of each fracture. Other research (Hildyard & Young, 2002), has shown that non-uniform loading on a set of fractures leads to gradients in stress along and among fractures which in turn creates gradients of stiffness that affect wave scattering.

Gradients in fracture stiffness or differences in fluid saturation among fractures results in complex and noisy signals that evolve over time as stresses and fluid saturations change (e.g.,Oliger et al., 2003; Pyrak-Nolte, 2007). If the originating probe signal has unknown source characteristics, the complexity of the problem increases. Indeed, the limited frequency content and the temporally scattered nature of these signals make their analysis with standard signal processing techniques difficult. All of the aforementioned factors control or influence coda wave generation from fracture systems and present a challenge when trying to determine fracture conditions when using passive monitoring methods.

Therefore, the ability to interpret fracture system evolution from passive monitoring techniques requires a method to identify the key scattered components related to the evolution in physical properties in a waveform independently of the source properties. Here, we begin to unravel the contributions that affect interpretation of fracture condition from coda waves by using laboratory test-bed experiments and machine learning. We focus specifically on two aspects of passive monitoring for a set of parallel fractures: (a) the use of stationary and moving uncontrolled sources, and (b) changing saturation in subsets of the fractures. The machine learning approach we take uses twin-neural-networks (TNN) that create low-dimensional representations among similar and dissimilar classes of seismic behavior.
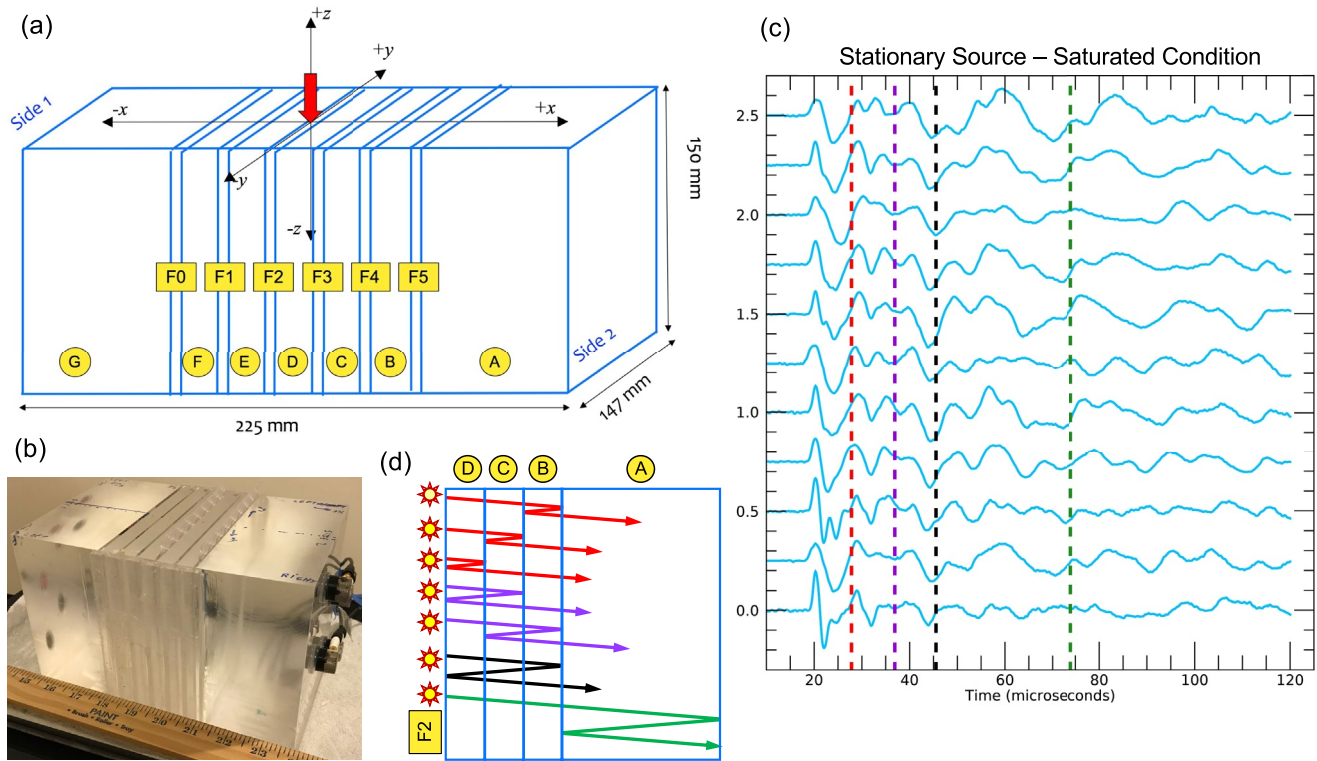
Figure 1. Fracture test bed and representative signals. (a) Sketch of the fracture test bed with the individual matrix blocks labeled A-G (yellow circles) and fractures F0-F5 (numbered yellow rectangles). Dimensions of the individual blocks and fracture apertures are listed in Tables 1 and 2. The x-y-z coordinate system is centered on the top of the sample at the center of F2. (b) Digital image of sample. (c) Acoustic emissions from a stationary source recorded for the saturated condition. The color-coded dashed lines represent the arrival times of internal reflections caused by an emission in F2 in the fracture set shown in (a). (d) Sketch of some of the possible dominant reflected-modes for acoustic emissions from a source in F2 for waves with an initial $+k_x$ wave component received on Side 2. The reflected modes are color-coded to the arrival times in (c).

2. Materials and Methods

2.1. Fracture Set Test Bed

A fracture test bed that contained six synthetic fractures was fabricated from acrylic blocks (Figures 1a & 1b). The dimensions of each block are listed in Table 1. The blocks were separated by 3D printed spacers with spatial relief to create partially open fractures. Silicone rubber was used to attach the spacers and thin acrylic plates to the bottom of the sample to create water-sealed fractures that were only open on the top of the sample from where fluids could be pipetted into or out of a fracture plane. The sample was clamped for 24–48 hr to enable the silicone rubber to solidify. The apertures of each fracture are given in Table 2.

Printed 3D plates (Figure 2) were used in fractures F3 and F5 (Figure 1) to simulate partially closed fractures. The plates were printed with a polymethyl methacrylate (PMMA) resin on a FormLabs 2 printer with a minimum layer thickness of 25 µm. The surfaces of the plates contained cylindrical protrusions (Figure 2c) arrayed in a hexagonal packing geometry with a center-to-center spacing of 8.8 mm (Figure 2b) and a height of 0.5 mm (Figure 2d). During desaturation of a fracture, contact was maintained between the plate protrusions and the fracture walls through water menisci. A syringe was used to remove 45–48 cc of water to lower the position of the air-water interface (e.g., Figure 2e) in a fracture. The approximate position of the air-water interface was $z = -46$ mm ± 2 mm below the top surface of the sample.

Four sample saturation conditions were used in the experiments (Table 3) and are referred to as: (a) "Sat" — when all fractures are fully saturated; (b)

Table 1
Dimensions of acrylic blocks shown in Figure 1a

Block	Width -x (mm)	Length (mm)	Height z (mm)
A	75.32	147	150
B	11.88	147	150
C	12.31	147	150
D	12.39	147	150
E	12.25	147	150
F	12.20	147	−56.38
G	77.095	147	150

Table 2
Apertures of fractures shown in Figure 1a

Fracture	Aperture (mm)
F0	2.03
F1	2.07
F2	2.07
F3	2.05
F4	2.06
F5	2.04

"S3" — when fracture F3 is partially saturated with all other fractures fully saturated; (c) "S5" — when fracture F5 is partially saturated with all other fractures fully saturated; and (d) "S3S5" — when fractures F3 & F5 are partially saturated with all other fractures fully saturated.

2.2. Sources

2.2.1. Internal Transportable Sources

Acoustic emissions were produced using chattering dust that are sucrose-based grains that each contain hundreds to thousands of small compressed CO_2 bubbles (Pyrak-Nolte et al., 2020). When the dust is placed in a water-saturated fracture, it dissolves, causing multiple delayed explosive releases of compressed gas from individual bubbles. Pyrak-Nolte et al. (2020) showed that a single grain can produce 100–1000s of emissions in a 3–6 min period with a frequency content that ranged between 50 kHz and 500 kHz. The distribution of signal amplitude was correlated to the distribution of the size of the compressed gas bubbles. 3D X-ray tomography of the grains observed that the gas bubbles ranged in size from 4.5 to 250 μm (Pyrak-Nolte et al., 2020). A single grain as it dissolves provides a multitude of "uncontrollable sources" as opposed to using piezoelectric transducers with known source characteristics. These uncontrolled sources were used to mimic signals more representative of those observed in nature and with the advantage that some location information of the dust is known in this case because the dust was always released in fracture F2 (Figure 1).

The dust was used to simulate: (a) a stationary uncontrolled source, and (b) a moving uncontrolled source. An uncontrolled stationary source mimics repeating sources that can occur in nature, for example, repeating earthquakes within a region (e.g., Abercromnie et al., 2020; Holtzman et al., 2018; Marone et al., 1995; Nadeau et al., 1995; Schaff & Beroza, 2004). The stationary source was achieve by placing a wire at a fixed depth of 30 mm in fracture F2 (Figure 2e). After a dust grain is released into a saturated fracture, the dust grain initially descends and then comes to rest on the wire. For the stationary source, only one grain was used per test and a minimum of 10 tests were performed for each fracture saturation condition. Examples of signals from the dust

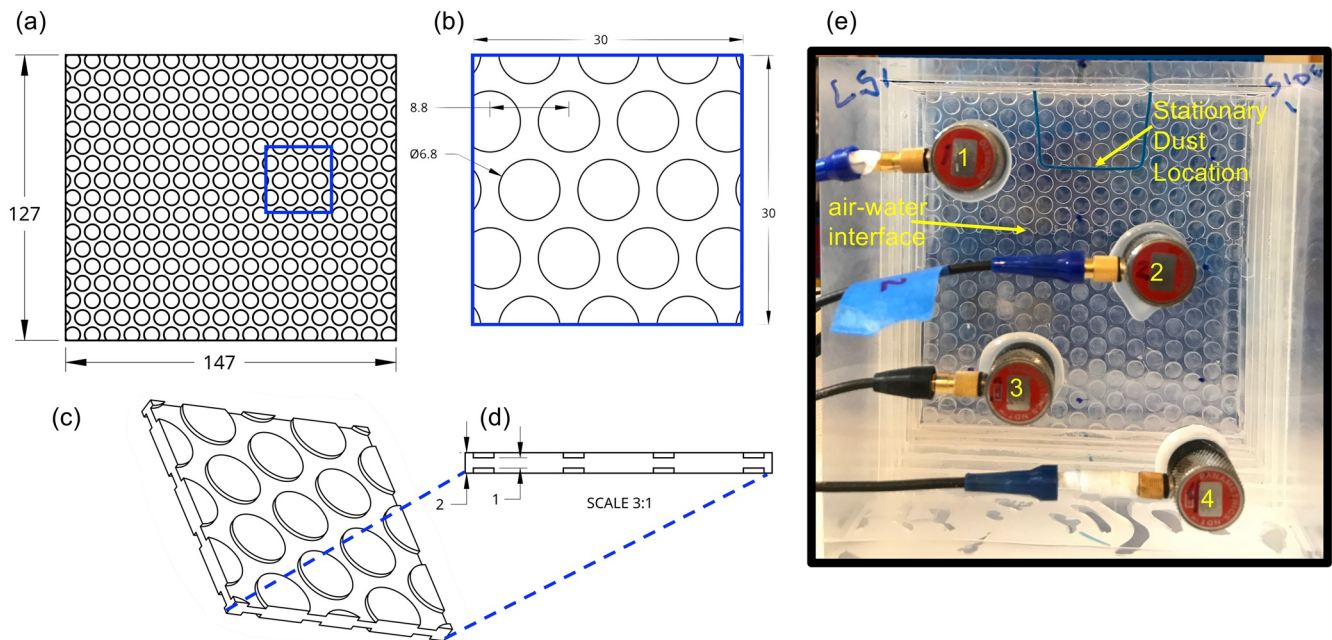**Figure 2.** Mechanical drawings of (a) an entire plate, (b) a sub region of a plate, (c) 3D view and (d) side view of a plate used to create contact between fracture surfaces in F3 and F5. The units are in millimeters. (e) Digital photograph showing a plate in fracture F3 for a partially water-saturated condition, the wire ledge in fracture for the stationary source location (blue wire), and transducers 1–4 on side 1.

Table 3
Saturation conditions for fractures F3 and F5

Condition	F3	F5
Sat	saturated	saturated
S3	partially drained	saturated
S5	saturated	partially drained
S3S5	partially drained	partially drained

released in the fracture test bed with fully saturated fractures are shown in Figure 3a to demonstrate the variability of the received signal amplitudes, frequency components and phases.

Conversely, a moving source was used to simulate induced seismicity along a fracture plane which may occur during hydraulic fracturing or geothermal stimulation (Eaton et al., 2018; Fehler et al., 2001; Holtzman et al., 2018; Lopez-Comino et al., 2021; Rodriguez-Pradilla & Eaton, 2020; Rowe et al., 2002). For the case of the moving source, a dust grain was released in F2 and fell through the fracture under gravity. For each test, 4–5 grains were released sequentially separated in time by 3–4 min to enable the emissions from a previous grain to subside. Again, a minimum of 10 tests were performed per fracture condition. An example of the location of events from a moving source released in fracture F2 is shown in Figures 3b and 3c when all of the fractures were fully saturated (Information on event location and error in location is given in the Supporting Information S1, section S3). Event location as a function of depth is shown for the *y*-direction (Figure 3c) and *x*-direction (Figure 3b). As the dust descends, the path of the dust is confined in the x-direction by the aperture of the fracture but is free to drift in the y-direction parallel to the fracture plane. Furthermore, the explosive emission of trapped gas causes the dust grain to make small sudden jumps (up to 1 mm displacement) in the unconstrained y-z plane.

2.2.2. Acoustic Emission Measurements

Acoustic emissions from chattering dust sources were recorded using eight broadband piezoelectric compressional wave sensors (Olympus V103 with central frequency of 1 MHz) attached to side 1 and side 2 of the sample (Figure 1a) using hot glue. The sensors remained in place for all experiments. The locations of the sensors on the sample are listed in Table 4 for the coordinate system shown in Figure 1a. The piezoelectric sensors were connected to an AE measurement system (8-Channel Mistra Express) through preamplifiers (Mistra 1220–5054, 20/40/60 dB single-ended powered preamplifier) to record signals using Mistra AEWin software. The threshold amplitude for detection was set at 40 dB (with a 60 dB preamplifier with 20 kHz–1200 khz window) which was
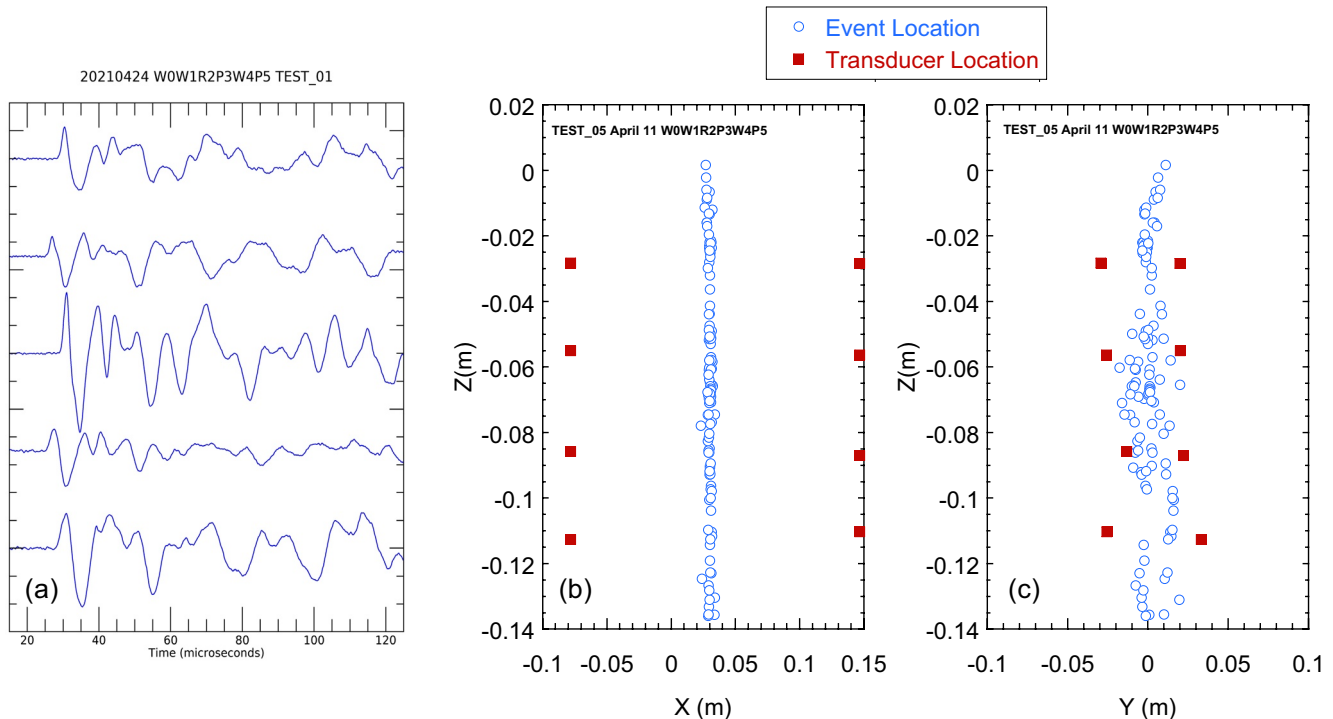

Figure 3. (a) Examples of signals from a fully saturated fracture sample with a stationary source. Locations of acoustic emissions from a moving source in a fully saturated fracture sample for (b) *z-x* and (c) *z-y* views (see Figure 1 for coordinate system). The locations are based on arrival times of events recorded by 5–8 sensors.

Table 4
Positions of AE sensors for the coordinate system shown in Figure 1a. Side 1 of sample: CH1-CH4; Side 2 of sample CH5-CH8.

AE Sensor	x (mm)	y (mm)	z (mm)
CH1	−78.1	−29.19	−28.25
CH2	−78.1	20.21	−55.1
CH3	−78.1	−13.17	−85.74
CH4	−78.1	33.29	−112.64
CH5	146.6	19.9	−28.37
CH6	146.6	−25.79	−56.38
CH7	146.6	22.21	−87.1
CH8	146.6	−25.52	−110.26

determined to be sufficient to eliminate ambient noise for this experimental set-up. The AE system records time to within 0.1 μs. The AEWin software determined the time and amplitude of every hit and saved the information to a summary line file. In addition, the full waveforms were recorded (10 MSPS, 100 μs pre-trigger) and saved to a hard drive. Information on source location methods and errors in location is given in the Supporting Information S1, section S3.

2.2.3. Sample Characterization With Controlled Source

Controlled-source transmitted signals were recorded for each of the four sample conditions (Figure 4) with a controlled source to compare with the uncontrolled sources. Transducers for CH1, CH2, CH3, and Ch4 on side 1 of the sample (Figure 1a) were used as sources and excited in the AST (auto-sensing testing) mode of the Mistra AE system. The pulses were transmitted through the sample and recorded by the transducers on side 2 of the sample (CH5, CH6, CH7, and CH8). For transducer pairs, CH3-CH8 and CH4-CH7, the conditions of the fractures were always the same (i.e., fully saturated) because both transducers were located (Table 3) below the air-water interface even for the partially saturated conditions. These signals exhibit the same amplitudes, phases and arrival times (Figures 4c & 4d). For the transducer pairs that are located within the partially saturated region of the fracture system, CH1-CH6 & CH2-CH5, the signal amplitude is largest for the fully saturated condition (Figures 4a & 4b) and smallest when the signal is transmitted through both partially saturated fractures, F3 and F5 (S3S5 condition). Similar amplitudes and arrival times are observed when either F3 or F5 are partially saturated. Signal amplitude alone is sufficient to distinguish the fully saturated case from the partially saturated cases of S3, S5, and S3S5. In addition, amplitude information from the 1st wave packet (119–130 microseconds) is not sufficient to distinguish differences in the signal when only a single fracture is partially saturated (either condition S3 or S5). The information that indicates whether it is fracture F3 or fracture F5 that is partially saturated is contained in the subtle changes in the waveforms that arises from the coda wave. Depending on whether F3 or F5 is the partially saturated fracture changes where the internal reflections occur from the parallel set of fractures and changes their amplitudes (Figure 1d). These changes tend to be subtle differences in the data
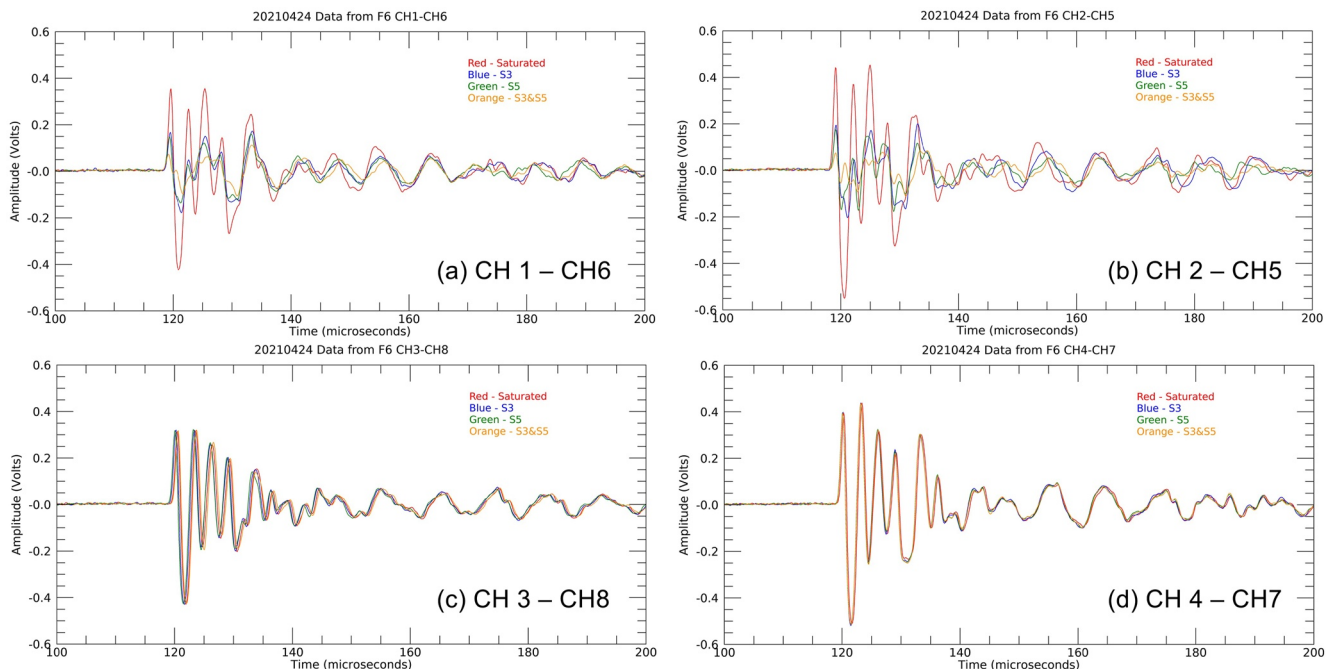

Figure 4. Recorded transmitted signals using controlled, repeatable sources for transducer pairs (a) CH1–CH6, (b) CH2–CH5, (c) CH3–CH8 and (d) CH4–CH7 for 4 samples conditions. The change in saturation is detectable in the CH1–CH6 and CH2–CH5 transducer pairs.

generated even with an active controlled source, hence an extra challenge is faced to distinguish such changes when the source parameters are unknown or unrepeatable when using the chattering dust as the seismic source.

2.3. Twin Neural Network

The challenge posed by uncontrolled chattering dust sources is the loss of three key information elements used by active monitoring: arrival time, amplitude and frequency content. These are the most information-rich aspects of seismic wave analysis of fracture systems in rocks (Nakagawa et al., 2000a, 2002; Pyrak-Nolte & Nolte, 1992, 1995; Pyrak-Nolte et al., 1990). When a chattering dust grain launches an acoustic wave, the event time is unknown, and the subsequent analysis aligns the arrival times of all the signals. In addition, the amplitudes and frequency contents of individual events vary (within a range, see (Pyrak-Nolte et al., 2020)), preventing direct assessment of fracture properties from individual signals. On the other hand, when the system under study is composed of multiple fractures, a rich coda-wave signal is produced from each event which contains information on the fracture geometry as well as the state of saturation. While well-separated codas can be analyzed directly, the codas are unresolved when the fracture spacing is small relative to the length of the event wave packet, creating complex waveforms that are difficult to analyze. For example, Figure 1c shows the waveforms of late-arriving multiple reflections. The reflections can occur from one or more slabs which can result in interference, or the reflections can have multiple alternative paths sharing the same travel time, illustrated in Figure 1d, but the amplitudes of the various overlapping alternatives depend on the fracture stiffness or saturation. This information is buried among the multiple overlapping waveforms, creating a challenge to extract fracture saturation or stiffness information from the signals.

In addition, the uncontrolled sources have high variability. For each of the saturation states, stacks of 1,400 signals are shown in Figure 5. The signals have all been aligned to have the same first arrival times, and they have been normalized by the square-root of the energy (as described in the Supporting Information S1, section S2.2). The signals under the different saturation conditions share general features that make it difficult to distinguish one saturation case from another. The averages are shown as the red curves, but the signals within each class vary by nearly 100%. In the absence of arrival time and amplitude information (which have been compensated and normalized), matched filter approaches fail to accurately classify single events, and alternative analysis methods must be used.

Neural networks are well-suited to analyzing complex data with many channels of information with high covariance among the channels and conditional associations that pose challenges to linear filtering approaches (Murphy, 2012). Deep learning in the past decade has moved beyond multi-layer perceptrons (MLP) to apply fully connected layers of high numbers of nodes to generate internal representations of data. When the internal representations capture the complex character of the data in a few degrees of freedom, it is much easier to visualize data relationships in lower dimensions (3D and lower), and classifiers are easier to train. Stacked autoencoders (Marsland, 2015) are one method used in dimensionality reduction, identifying common signal features within highly varying background. However, autoencoders can lose subtle differences among multiple classes, finding the common features rather than the differences.

A recent deep network architecture that solves this problem is the TNN (also known as a Siamese neural network) (He et al., 2018; Zheng et al., 2016). A TNN is a deep neural net, shown in Figure 6a that uses two identical deep networks with identical weights that are exposed to a pair of data inputs that are known to be from the same or from dissimilar classes. This network has an advantage in that distinct class labels are not needed for training: only whether the members of the pair are from the same class or different classes. Furthermore, any number of different classes are allowed. The outputs from a pair of data sets is used to calculate a contrastive loss function that is minimized, thereby training the network to recognize when the two inputs are from similar or dissimilar classes. For data analysis problems with small signals in the presence of high backgrounds, this twin network can learn to disregard the background and to recognize subtle signal differences (Berlemont et al., 2018). A TNN performs dimensionality reduction through successive bipartite (fully connected) layers and ReLU layers. A large number of input channels, for instance from a time-series seismic signal, are sent through multiple layers, and in this case, the output dimensionality is three-dimensional (the maximum for conventional data visualization).
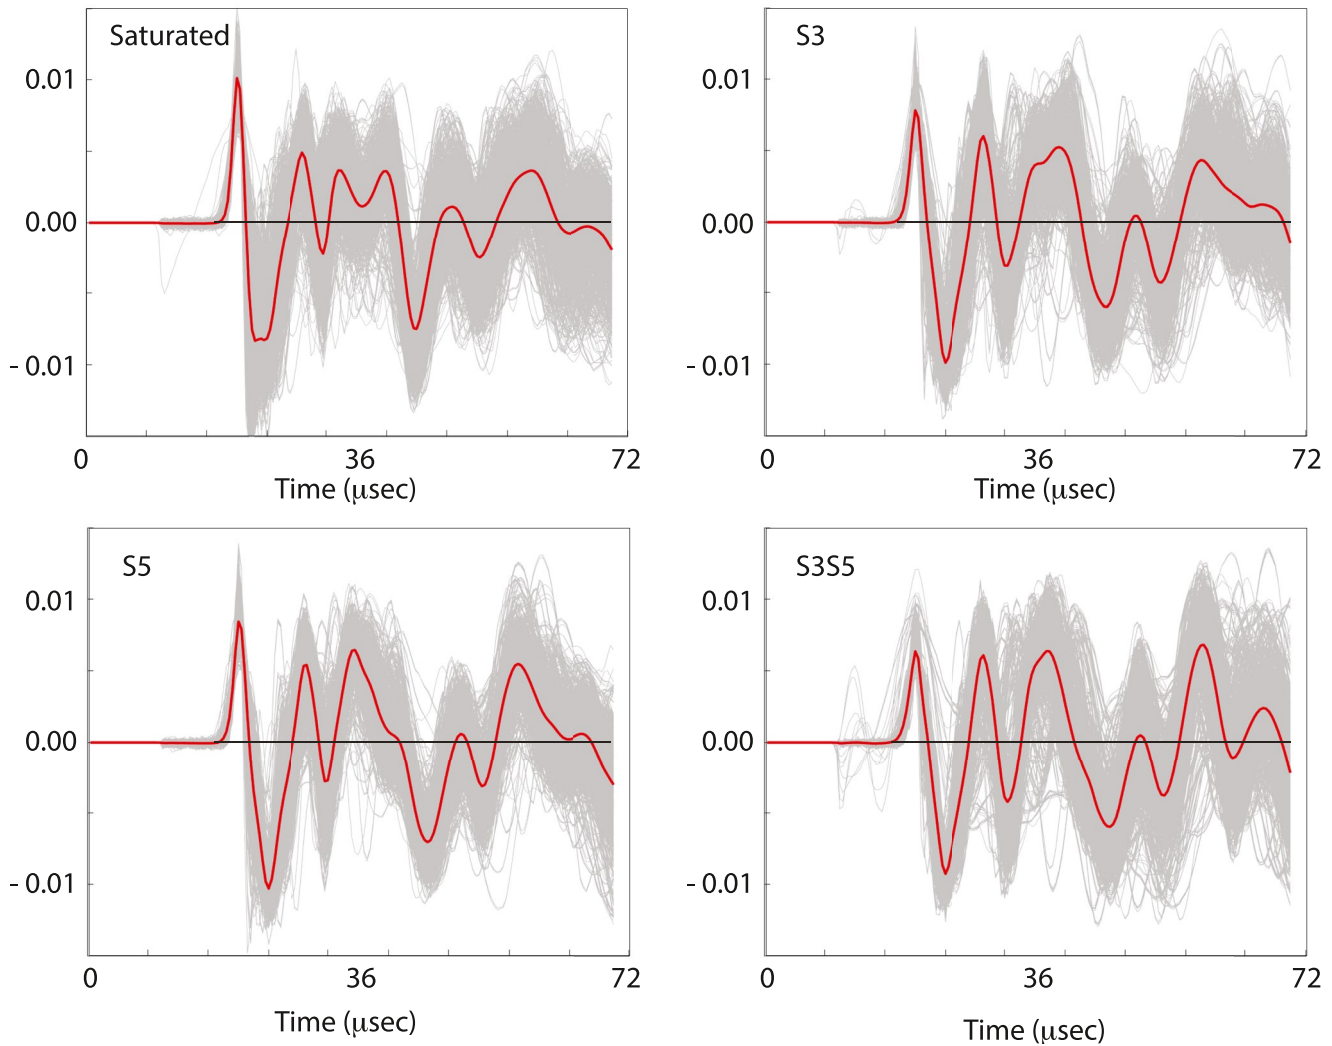
Figure 5. Stacks of all training signals (energy normalized) for the four saturation conditions at constant depth. Each set has 1,400 signals that have been shifted to the same arrival time and normalized by received energy. The signal-to-signal variability caused by the uncontrolled sources within each class is around 100%.

The contrastive loss function L_{contr} is defined as

$$L_{contr} = \frac{1}{2}yd^2 + \frac{1}{2}(1-y)|\max(m_0 - d, 0)|^2 \tag{1}$$

where y is the class-agreement label ($y = 1$ for similar and $y = 0$ for dissimilar), and d is the Euclidean distance between the two output features of the two subnetworks

$$d = \|F_1 - F_2\|_2 \tag{2}$$

The margin parameter m_0 is used as a constraint: if two images in a pair are dissimilar, then their distance should be at least equal to the margin, or a loss will be incurred. Only one term of the loss function is active for a given pair. If they are similar ($y = 1$) then the algorithm minimizes their distance d. If they are dissimilar, then the algorithm maximizes their distance to be at least a distance m_0 away from each other.

An additional function can be added to the TNN as a tool to identify which parts of the waveforms carry the most significant information. This is accomplished by adding a regularization term to the loss function to minimize the input weights from channels that carry little distinguishing information. The regularization loss function L_{reg} is
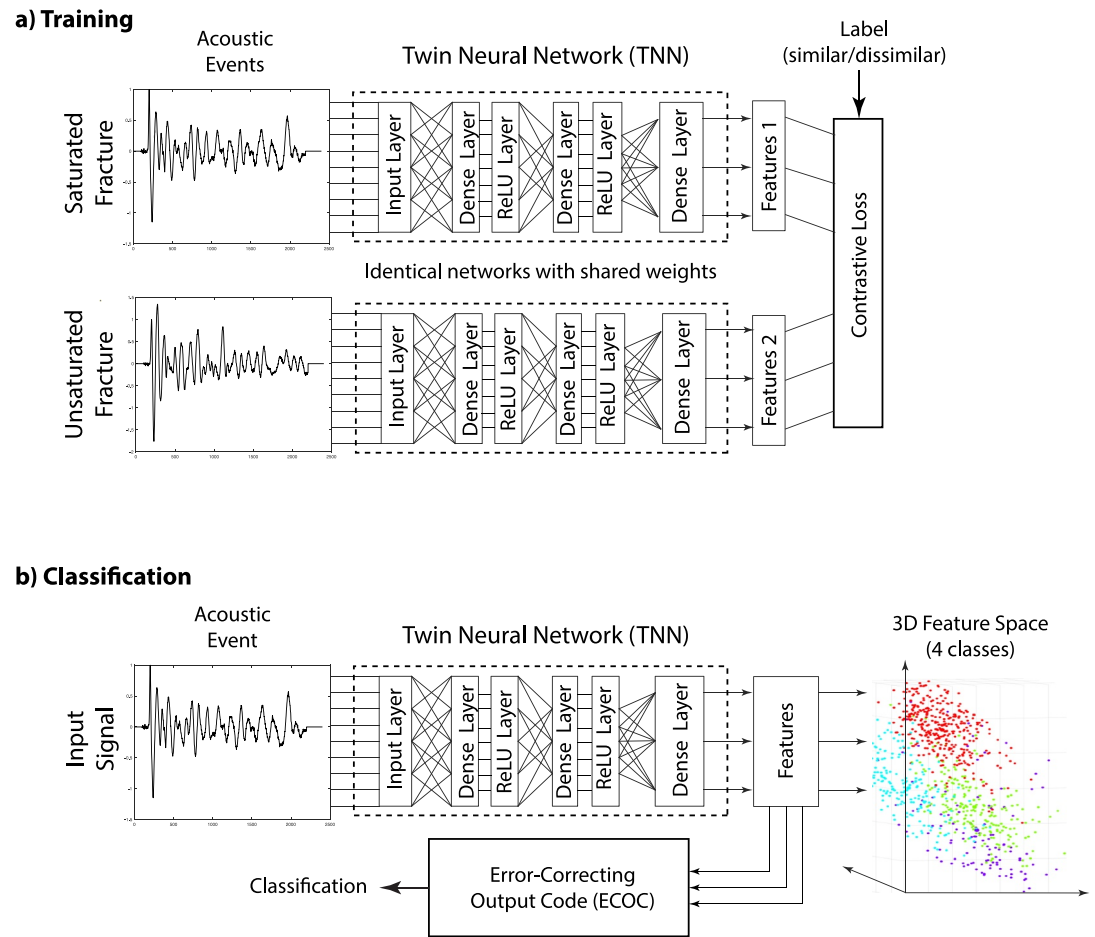
Figure 6. The "twin" network for dimensionality reduction consists of identical twin networks with shared weights. (a) In training, data structures are applied in pairs, then the contrastive loss is used to update the shared parameters. The network in this figure has two layers of fully connected ReLUs that are connected to the three output neurons. (b) A single signal is input into the trained twin neural network (TNN), and the 3-dimensional feature output is the input to a multi-class classifier that is trained separately. The three-dimensional latent space can be visualized independently.

$$L_{\text{reg}} = \lambda \|w_1^a\|_1 \tag{3}$$

where the norm is *L1*, w_1^a is the matrix of input weights into the first hidden layer, and the hyperparameter λ is tuned so that the regularization loss is approximately a third of the total loss value at the end of training. Optimization of the total loss function with the *L1* regularization loss causes the weights of unimportant channels to be down-weighted during training after being initialized with random weight values. By analyzing the input weights in the trained network, it is possible to identify the most information-rich parts of the waveform.

The deep twin network used for this analysis had three fully connected (dense) layers with 128 nodes in the first layer, 64 nodes in the second layer and 3 nodes in the final layer. The decreasing dimensionality was chosen to encourage the system to funnel information from the high-dimensional input to the low-dimensional output. Choosing a smaller number of nodes in the second layer below this value causes the data to form a one-dimensional manifold in the three-dimensional output space with poor classification performance. Therefore, the middle layer was the controlling factor in the selection of the number of nodes in the multiple layers. The activation function for the first two layers is ReLU, while the final-layer neurons perform a simple sum with bias and linear output. The margin in the contrastive loss function is set comparable to the variabilities caused by the uncontrolled sources, balancing intra-class variance against inter-class separations. The hyperparameters for the TNN were tuned for speed of analysis and accuracy of classification. The training was performed using the Matlab version

R2019b Deep Learning toolbox with the adamupdate optimizer. The minibatch size was 128 with a learning rate of 1×10^{-4}. The number of training steps was 1,500, which allowed convergence to a stable trained network while preventing overfitting. The data in the output space is used as the input to the standard Matlab error-correcting output code (ECOC) for multiclass analysis. The ECOC is trained separately after dimensionality reduction. It is a standard Matlab toolbox multiclass support vector machine classifier based on multiple one-versus-other maximum-margin classifiers.

3. Classification of Fracture States From a Deep Twin Network

The goal of the TNN is to identify the saturation states of multiple fractures when probed by stationary or falling sources. The network is trained on acoustic signals passing through sections of the fractures that are either saturated or have been partially drained. In the drained condition there remains a water meniscus around the protrusions in contact with the acrylic blocks. The meniscus properties and the saturation vary with depth, providing an additional test of the sensitivity of the trained network to partial saturation.

3.1. Stationary Source Probing Fracture Saturations

A first series of experiments held the chattering dust sources at a constant depth in the launching fracture (to mimic a stationary source) to probe two fractures under the four saturation conditions described in Section 2.1 and listed in Table 3. The set of 6,600 waveforms is shown in Figure 7a from which 70% (4,620) and 30% (1,980) are used for training and validation, respectively. Additional information on the partitioning of the data into training, validation and testing sets is given in the Supporting Information S1, section S2.3. The TNN was trained on the training set to generate five models each with random initialization of the neural weights. The best model of the five was selected based on the performance of the validation set. This model was then applied to the test set. The test set (1,192 events) is shown in Figure 7b. In the figure, the color represents the amplitude of the signal and the time base is 0.4 microseconds per channel. The training and validation set predictions at the end of training are shown in Figure 7c, and the predictions for the test set are shown in Figure 7d. A key aspect of this approach is to acquire and analyze each event separately without averaging. Because the source is uncontrolled, amplitudes and frequency content can vary from one event to another, and averaging can wash out significant information contained in the wave form.

The confusion matrix for the test set is shown in Figure 8a. The classes 1–4 correspond to the four saturation conditions: Sat, S3, S5, and S3S5. The diagonal terms (0.91, 0.88, 0.88, and 0.84 for classes 1 through 4, respectively) are much larger than the off-diagonal (largest value of 0.11 for class 4 misidentified as class 2) for all conditions. By testing a large number of events under a given condition (see Supporting Information S1, Table S5 for the number of signals for each condition), the classification of the fracture-set condition is taken from majority rule (mode-pooling) for the most likely classification within the data. Majority rule (as seen in the confusion matrix) provides perfect classification of the four saturation conditions.

The reduced-dimension feature space for all of the test events is shown in Figure 8b color-coded on the fracture saturation condition. The TNN dimensionality-reduction down to three features has separated the events among the four conditions. There is some overlap among events in the different classes, but the majority of the points have large Euclidean distances among the separate point clouds. The saturated events (red) are well separated from the others. The two single-fracture-syringed cases do have some overlap, as would be expected because the direct waves would be similar, but there is still good separation in the classification because of differences in the codas. This is a key finding of this paper, because it demonstrates that the coda is sensitive to the difference between syringing water from fracture F3 versus fracture F5 because of the change in the reflection conditions of the acoustic waves from the fractures. The double-syringed condition has some overlap with each of the single-syringed conditions, which also may be expected because the condition S3S5 involves the same fractures as the conditions S3 and S5. Nonetheless, the majority-rule correctly classifies each saturation conditions.

One of the challenges of deep learning is the physical interpretation of the internal weights and states of the network. Deep neural networks suffer from the "black box" problem in which an extremely large number of latent variables form an internal representation that is an encrypted form of the external information input into the network. As an example, all of the input weights for the first layer of the TNN are shown in Figure 9a for the 128 first-layer neurons and 130 input channels. There is no recognizable pattern in the matrix of neural weights, al-
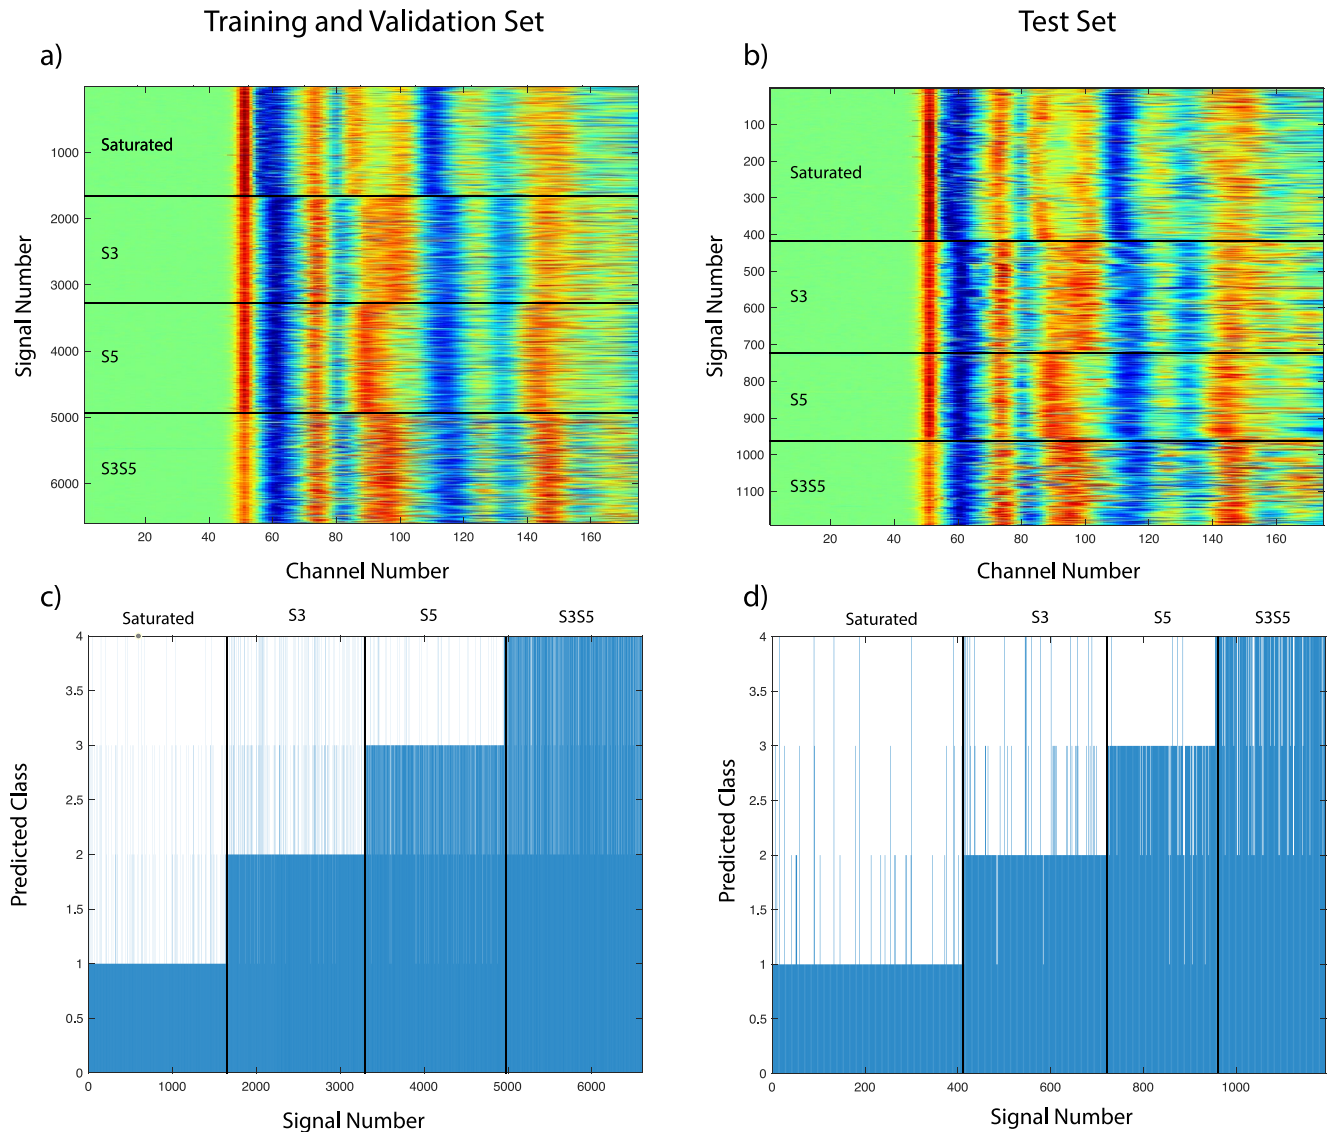
Figure 7. Training, validation and test signals. (a) Training and validation set under four conditions. (b) Independent test set under the four conditions. (c) Training and validation classifications. (d) Test set classifications. The test set is composed of 20% of the acquired signals that were not used for training and validation.

though the L1 regularization has down-weighted the least informative channels. However, creating the cross-correlation matrix among the neural weights and clustering on similarity produces the square correlation matrix in Figure 9b that shows clear groupings in the neural weights. In this case the correlation matrix is approximately block diagonal with three similar groups of neural weights. This demonstrates that there is a large amount of covariance among the weights with only a few orthogonal behaviors. This is expected because of the deterministic time series behavior of the input signals. The oscillatory character of the acoustic time-series induces strong co-dependencies along the time axis which is reflected in the clustering of the input weights. This can also be seen in the correlation matrix among the input channels without clustering in Figure 9c. The correlation matrix has roughly a checker-board pattern that reflects the oscillatory character of the input signals. These correlations in the time-series signal might also be extracted using the properties of convolutional neural networks, but this will be explored in a later study. In addition to the internal weights, it is also possible to extract the internal states of the hidden neurons as shown in the bottom row of Figure 9 that show the internal states for the training events sorted by class number. Clear class divisions are observable in the internal states for all three hidden layers.
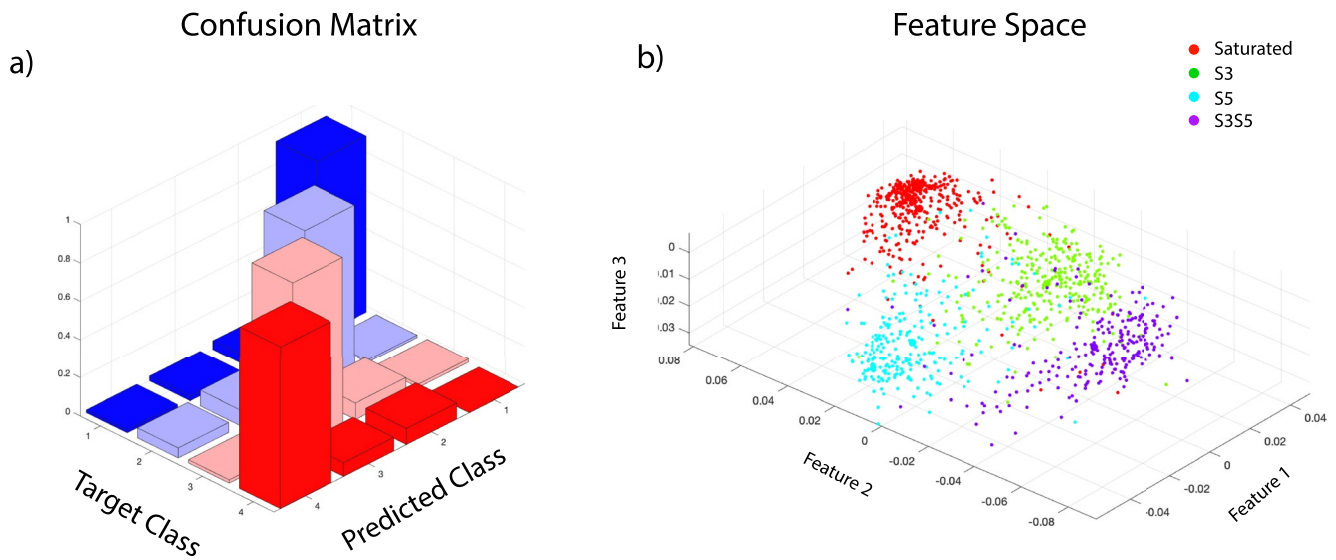
Figure 8. Stationary-source classification results for the test set with four saturation conditions. (a) Confusion matrix of predicted class against target class. (b) Feature space of the events of the four conditions for the test set.

By interrogating the input weights to the first hidden neuron layer it is possible to identify which input channels carry the dominant information. This is made especially clear when *L1* regularization is constructed to depend only on the input weights and is added to the loss function for optimization. The *L1* regularization suppresses weights from channels that do not provide significant information for the low-dimensional clustering. The results of this analysis are shown in Figure 10 that identifies the channels that had the highest input weights. It is clear that the portions of the waveforms where there are distinct differences among the four conditions contribute the most weight to the twin network inputs. In addition to the peak of the first arrival, there is considerable weight from the region of the second coda around 36 microseconds. The third coda also contributes weights on the input, but the first coda contributes little. The signal packet of an unscattered source wave has about one and a half cycles, causing all of the codas to overlap. In this analysis, the deep twin network is not only convenient for identifying and classifying different fracture saturations, it is also a tool to identify the most information-rich portions of the waveform.

Another challenge for neural networks is their ability to generalize trained behavior to be able to predict the properties of never-before-seen data. During training, the upper bound to the accuracy of the predictive model is established by the validation set. However, validation data sets are subsets of the original data. To test the predictive power of a neural network model, it is necessary to have independent data that are acquired separately from the training and validation data sets. We performed a repeat of the constant-depth data acquisition (Section 3.1) using the same synthetic fracture network but under an additional cycle of saturation and drainage. Approximately 8,268 events (see Table S5 in Supporting Information S1) were acquired for each saturation condition (saturated, S3, S5, and S3S5) that were split 70% and 30% (5,787 and 2,481) for the training and validation sets, respectively (see Supporting Information S1, section S2.3). The system saturation was physically cycled through the four saturation conditions and an additional 7,559 test data events were captured for each condition. The TNN was trained on the training set to generate five models each with random initialization of the neural weights. The best model of the five was selected based on the performance of the validation set. This model was then applied to the test set. The resulting confusion matrix of the test set is shown in Figure 11. The four saturation conditions for the independent test set are predicted with accuracies ranging from 67% to 85% (compared to random prediction that would yield 25% accuracy for each condition under four-class classification and compared to Figure 8 for the non-prospective test set that had accuracies ranging from 82% to 92%). The maximum confusion values (off-diagonal) are 26% of the double pipetted case S3S5 confused with two single saturations S3 and S5. This first may be anticipated because the first coda is strongly affected by the F3 reflection for both of these cases. The next may be anticipated because of the similarity of the transmitted signal for these two single-fracture saturation conditions. The accuracies for the predictions are on a per-event basis, meaning that the probability of correctly
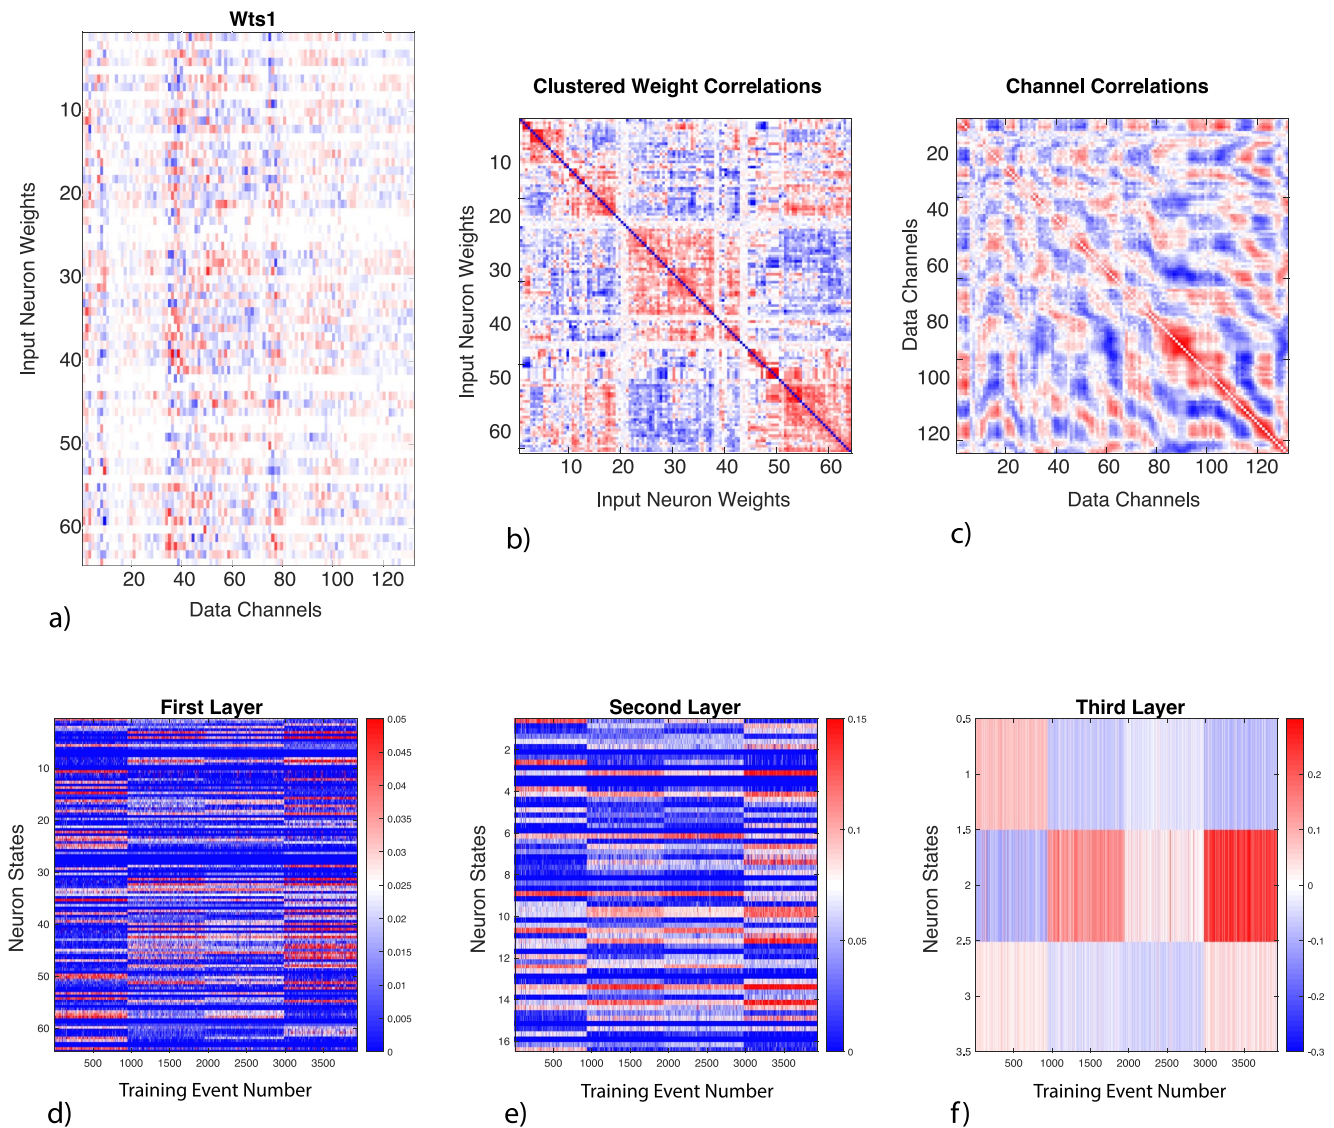
Figure 9. Internal weights and states of the twin neural network. (a) The input weights between 130 input channels and 128 first-layer neurons. (b) Correlation matrix of the input weights correlated across the data channels and clustered showing a small number of similarities among the input weights. (c) Correlation matrix of the input channels correlated across the input weights. The checkerboard pattern reflects the oscillations in the input signals. (d) Output states of the first neuron layer sorted by class. (e) Output states of the second neuron layer. (f) Output states of the the third and last neuron layer.

classifying a single event is approximately three times higher than at random. In applications, characterizing the saturation of a fracture uses multiple events of multiple chattering dust grains to characterize a single condition. Therefore, this approach achieves perfect accuracy for assessing fracture saturation based on mode-pooling the predictions of the individual events.

3.2. Moving Source Probing Different Saturations

One of the valuable assets of the chattering dust is the ability for the dust grains to move through a fracture system while being tracked and located by the receivers. The changing position of a dust grain as it falls under gravity, or as it is transported by internal fluid currents, allows it to probe spatial variations in system properties. For instance, in our experiments that alter fracture saturation, the saturation is a function of depth because of the residual meniscus held by capillary forces around the protrusions of the fracture inserts. As a dust grain falls, the wave path probes different portions of the fractures.
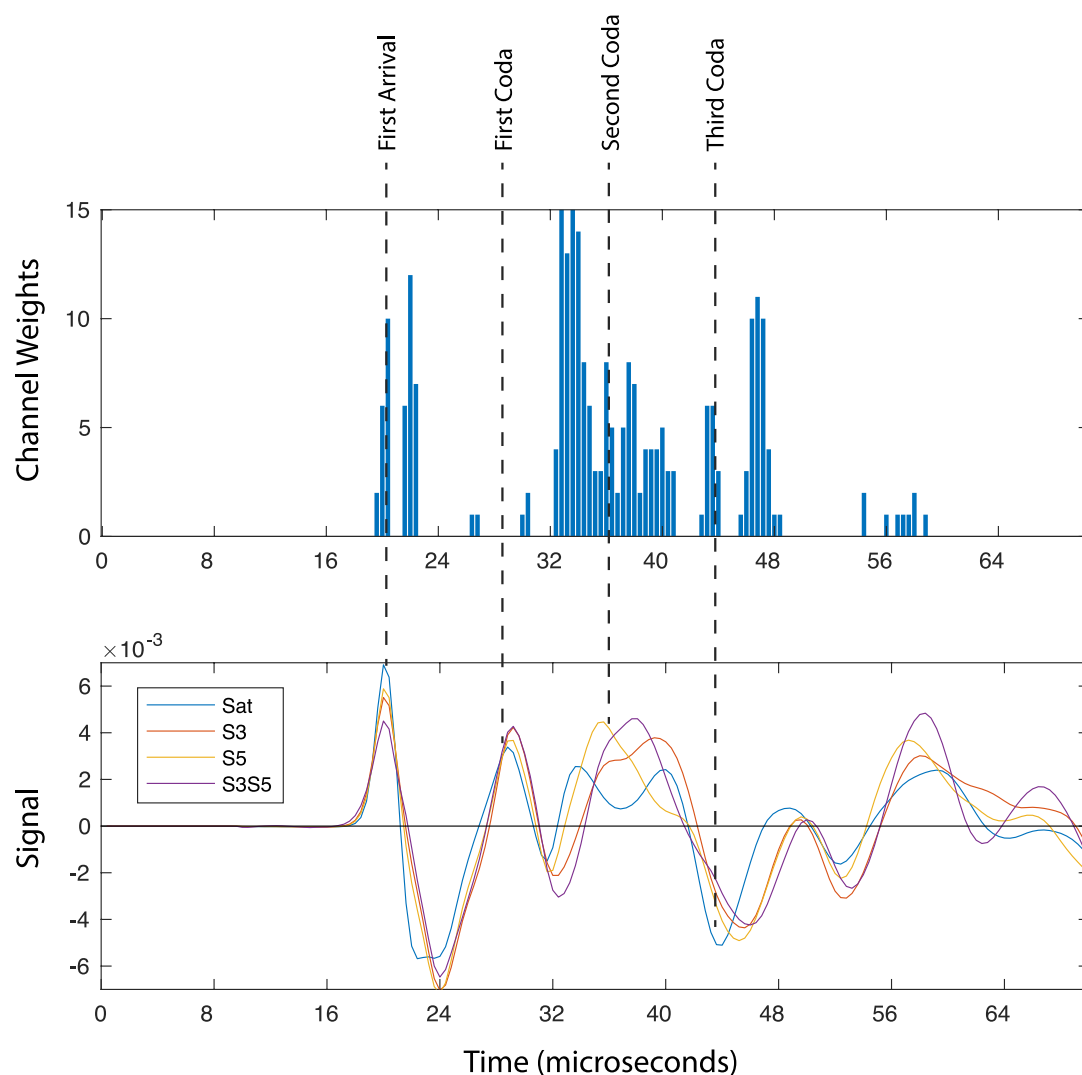
Figure 10. (top) Channel weights from L1 regularization. (bottom) Average signals for the four saturation conditions. The locations of the first arrival and codas are shown.
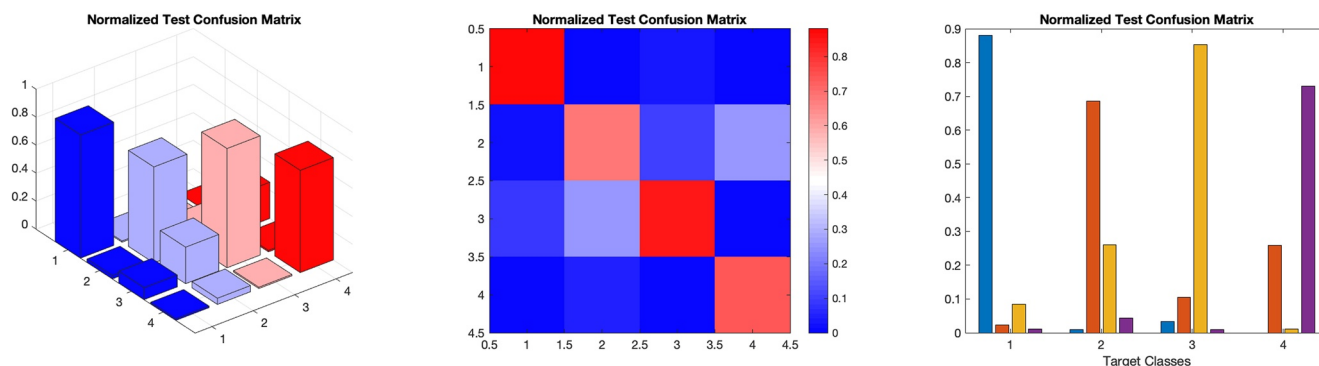

Figure 11. Three representations of the confusion matrix for prospective 4-class data at constant depth. Approximately 8,000 events were used for training with an additional prospective set (after saturation cycling) of 8,000 signals for prediction.
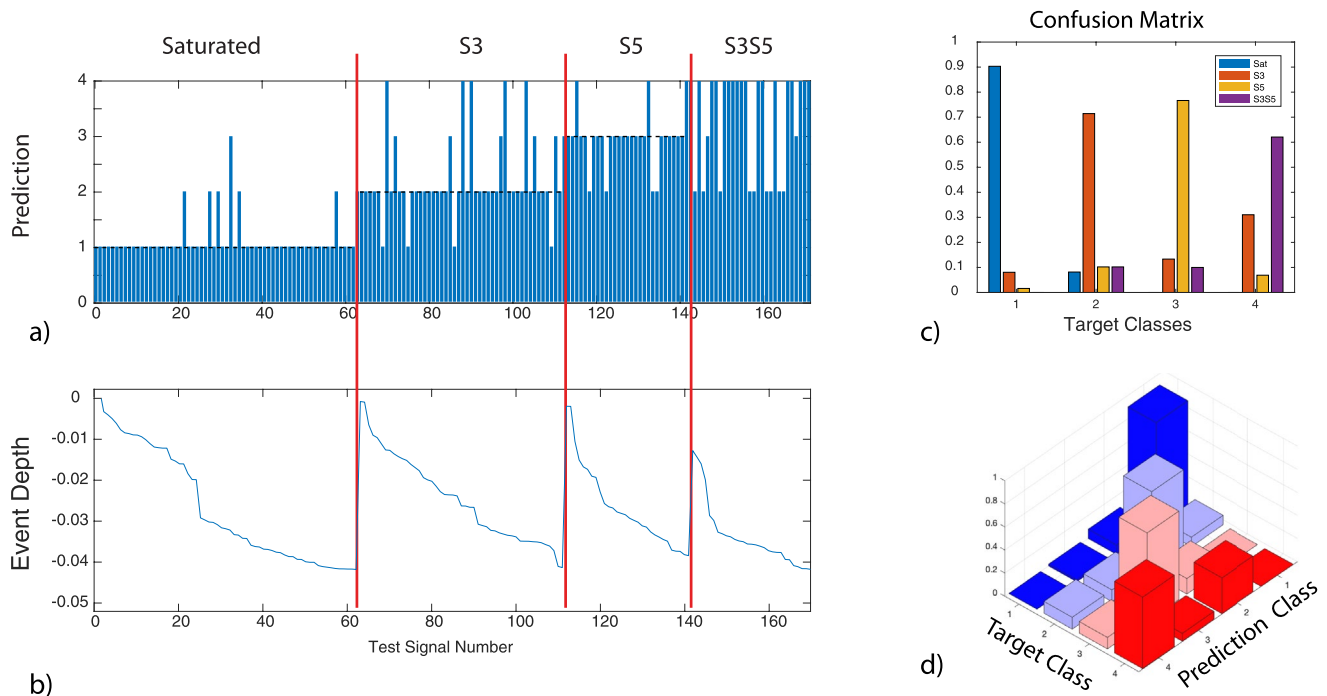
Figure 12. Network performance on falling dust grains. (a) Class predictions for the four saturation conditions. (b) The location of the grain at the moment of emission. (c and d) Confusion matrix for four-class classification.

A series of experiments were performed under the same experimental conditions as in the previous section, except that the dust grains fall under gravity after being released at the top of the fracture F2. A similar data set was acquired but with the position of the dust grain tracked by triangulating the signals received at the multiple acoustic receivers. The TNN was trained on the new data with a 20% (170 test events) reserved for testing relative to 691 training events and 297 validation events. The TNN was trained on the training set to generate five models each with random initialization of the neural weights. The best model of the five was selected based on the performance of the validation set. This model was then applied to the test set. The classifications of the test events are shown in Figure 12a that can be compared to the depth of the dust grain at the moment of emission in Figure 12b. The first set of signals in Figure 12 are for the fully saturated condition. Even as the dust grain falls under gravity there is little mis-classification. The second condition S3 has fracture F3 partially saturated. In this case the predictions are seen to be distributed across all three possible partial-saturation conditions with accuracies of 90%, 65%, 73%, and 62% for conditions Sat, S3, S5, and S3S5, respectively. The changing depths of the falling grains produces depth-dependent changes in the waveforms that make the distinction among the conditions more difficult to distinguish. The confusion matrix, shown in Figures 12c and 12d has large on-diagonal values. The maximum confusion is 31% for class S3S5 mistaken for S3. This may be anticipated because of the similar coda structure for the S3 and S3S5 saturation conditions. Despite the misclassification of some individual events, the majority rule (mode-pooling) gives accurate classification in this case with the transporting internal seismic source.

3.3. Probing Saturation With Depth

When the water level is drained (pipetted) from a fracture containing the artificial asperities, capillary forces trap a meniscus around each of the protrusions of the 3D printed fracture inserts. Therefore, the saturation is a function of height, transitioning from fully saturated at low values of height to unsaturated at high values of height. This z-dependent saturation can be probed by a series of falling dust grains. The TNN in this case was trained on a two-class system for either a fully saturated sample or a partially saturated F5. The TNN was trained on 9,100 events and validated on 3,900 events to select the best model that was then applied to 2,600 test events. The classification of the test events are shown in Figure 13a with excellent accuracy identifying the saturated case. However, it is
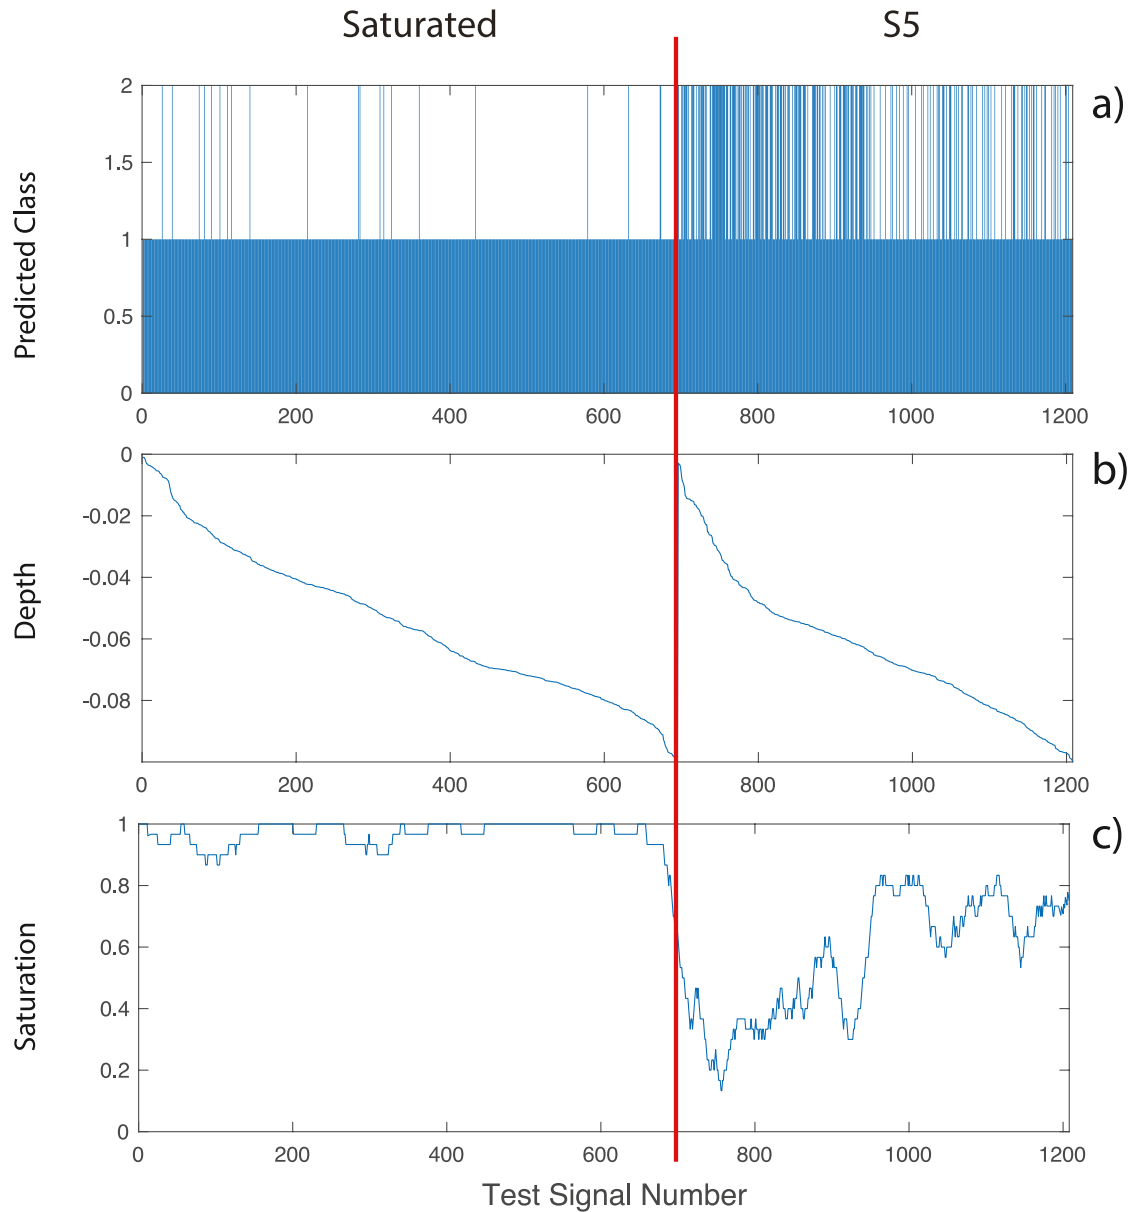
Figure 13. Two class classification on saturated versus partially saturated conditions. (a) Classification from the twin neural network. (b) Depth of the emission event. (c) Average saturation based on the percent classified correctly.

important to note that the partially saturated case is only accurately identified for high z values of the emitting dust grains. When the dust has a high z, the direct raypath through the sample passes through the dry portion of fracture F5. As the z becomes more negative, more of the wave samples more of the water meniscus closer to the water level. Therefore for lower values of z (i.e., increasing depth) for the partially saturated condition, the classification shifts back to the saturated class. By averaging the classifications using a windowed average, a pseudo-saturation graph can be obtained, shown in Figure 13c. The saturation is high for the first class, shifts to low saturation initially for the second class, but then increases again for deeper emission depths as the wave probes more of the water meniscus. This demonstrates the potential use of the TNN to map out the saturation of a partially saturated fracture.

4. Discussion and Conclusions

Networks of fractures in the subsurface are inaccessible to direct interrogation, unless accessed destructively through boreholes or mining. Therefore, studies of fracture networks have relied either on remote detection using active seismic sources, or on passive internal sources such as induced seismicity, acoustic emission or ambient noise. A goal of the current study is to establish a middle ground between active probing and passive sensing in which transportable sources infiltrate fractures to illuminate fracture networks acoustically from the inside. A principle question is whether such sources of acoustic emission internal to a set of fractures can be used to characterize the conditions of other fractures within the set. The challenges to this approach are two-fold: (a) internal sources are uncontrolled (in terms of the wave attributes such as frequency and amplitude), and (b) internal reflection modes that sample the fractures within the set often overlap with portions of the direct wave and with other combinations of reflections, complicating wave analysis. In the work presented here, by using a laboratory-based fracture test bed, the expected arrival times of different reflected modes (Figure 1d) were known. By examining the signals (Figure 1e), these reflection modes (codas) were confirmed to occur within the first few cycles of the wave, obscuring identification of the specific reflected modes that contain information from the waves interacting with the fractures in the set. In addition, the acoustic emissions from the internal sources varied in magnitude, leading to differences in amplitude and arrivals amongst thousands of recorded signals (Figure 5).

To handle the large amount of signal information with large signal-to-signal variations, we employed a machine learning approach known as the TNN that uses supervised learning on a training set to analyze the waves to classify the saturation conditions of the fractures within the set. The approach used the normalized time-series data (i.e., signals) without relying on arrival times, amplitudes or spectral content. Despite this elimination of the primary information content (that is a fundamental part of active sensing), the signal from the fully saturated condition was shown to be distinctly different from signals from the partially saturated conditions as observed in the controlled source data (Figure 4) and the uncontrolled source data (Figures 7a & 7b). The differences among the S3, S5, and S3S5 conditions were more subtle but still recognizable by the neural network. Based on a mode-pooling decision rule, the classifications of individual events within a cohort of events led to perfect classification among the four saturation conditions.

A 4-class prospective study demonstrated the robustness of the method by training on data acquired under an initial set of saturation conditions, then physically cycling the saturation and testing data collected at a later time. This again achieved perfect classification using the mode-pooling rule. This indicates that information on the conditions of fractures can be extracted from coda wave signals that are not strongly separated in time from direct arrivals or other modes. On the other hand, all of the TNN studies reported here used a training set that is needed to capture the subtle differences among the saturation conditions. Practical extension of this work to the field, or even to other laboratory studies, is currently limited to situations in which fracture conditions are cycled so that conditions can be recognized from known training cases. A future extension of this study would use unsupervised training to find clusters of behavior that would seek to identify several different saturation conditions. The TNN requires class-similarity labels for training and so may not be suited to unsupervised learning. On the other hand, the TNN would be amenable to iterative semi-supervised learning in which class-similarity labels are initialized randomly and then updated according to the separations observed within the latent representation. However, such a semi-supervised approach is beyond the scope of this current study.

An important extension of the use of the multi-class TNN classifier was demonstrated to provide probabilities that can be continuous valued. While the class labels from the TNN are discrete, the latent representation is used in a direct multiclass classifier on an event-by-event basis. In the 2-class study, the low-dimensional representation from the TNN enabled estimation of changes in saturation with increasing depth (Figure 13) using signals from an internal moving source based simply on the ratio of class-1 to class-2 classifications on the events. The increasing saturation with depth was reflected by an increasing proportion of class-1 predictions relative to class-2. This continuous-valued ratio becomes a surrogate for changing saturation with depth.

The void geometry of a fracture, defined both by the spatial distribution of apertures as well as the distribution of aperture sizes, affects the fracture specific stiffness, κ, and affects the partitioning of energy into reflected and transmitted modes. There are two limiting cases in which stiffness goes to infinity or to zero. For the first limit, a fracture behaves as a welded contact and no coda waves are generated. Based on the displacement discontinuity theory, welded contact behavior occurs for $\kappa > 30 \, \omega Z$, where ω is the frequency of the signal and Z is the

seismic impedance (density × phase velocity). Therefore, for a given ω, the fracture can still be open but most of the energy is transmitted, and reflections are not significant. For the second case, a fracture will behave as a free surface when $\kappa < 0.03\ \omega Z$ and all of the energy is reflected from the fracture. However, when the frequency of the source enables optimal detection, that is, for $0.03 < \kappa < 30\omega Z$ (Pyrak-Nolte, 2019), then the ability of the fracture set to generate interfering coda waves from reflections between different fractures within a fracture set will depend on how the energy is partitioned into reflected and transmitted modes at each fracture. In an extreme case for which the aperture is uniform and large, that is, similar to parallel plates, additional reverberations might arise from internal reflections within a single fracture, acting as an additional source of tightly spaced codas if the fractures were saturated with water. However, if this parallel-plate fracture were unsaturated, then no energy would be transmitted.

Another question is how the results presented here can scale to field applications where varying signal-to-noise ratios can occur as a fracture evolves, or distances of the fracture from the sensors change, or the signal quality can vary from location to location. The advantage of our technique for potential field applications is three-fold. First, only data from one sensor was needed to classify the signal, eliminating the need to consider sensor-to-sensor variation. Second, the additional sensors were only used to locate the position of the source to enable the analyst to know whether to consider a stationary versus a moving source. If it is a moving source, then the location of the source and comparison of signals from different source locations can be included to predict the expected codas from a fracture network. Third, as long as a fracture set or fracture network produces codas, changes in the condition of fracture system can be identified with this technique. The relative fracture spacings control the timing of the coda signals caused by reverberations among a fracture set, and changing saturation does not affect the spacings, only relative coda amplitudes.

Challenges that may be faced while developing the technique for the field could depend on whether multiple aspects of the system between the source and sensors are changing. For example, both the matrix and fractures may exhibit changes in saturation that affect wave scattering. Furthermore, the optimal detection frequency range may limit the range of fracture stiffnesses that can be detected for any given frequency, that is, fractures are present, but the source frequency content is not optimal for scattering or partitioning into multiple internal reflections. Numerical modeling would be needed to simulate coda wave generation from fracture sets and fracture networks to compare codas for a system when only the fracture saturation is changing, when only the matrix saturation is changing and when both are undergoing changes in saturation. The simulations would help identify how the coda in the received signals would vary for these saturation conditions and could also simulate how signals from a stationary versus moving source would affect the coda wave interpretation. These simulations could help develop further training sets.

The approach investigated here using the TNN could be applicable in the future to hydraulic fracturing sites. The induced seismicity from one induced fracture can be used to monitor the state of other previously induced fractures. As a hydraulic fracture propagates, the induced-seismicity is emitted from the propagating front as well as from other regions of the fracture. The uncontrolled nature of these sources could be managed using the TNN approach demonstrated here to analyze the properties of fracture networks. In addition, by increasing the strength of the chattering dust emissions, transportable acoustic sources may in the future be able to illuminate fractures from inside a subsurface fracture network.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Data Availability Statement

References

Abercrombie, R. E., Chen, X., & Zhang, J. (2020). Repeating earthquakes with remarkably repeatable ruptures on the San Andreas fault at Parkfield. *Geophysical Research Letters*, *47*, e2020GL089820. https://doi.org/10.1029/2020GL089820

Berlemont, S., Lefebvre, G., Duffner, S., & Garcia, C. (2018). Class-balanced siamese neural networks. *Neurocomputing*, *273*, 47–56. https://doi.org/10.1016/j.neucom.2017.07.060

Boue, P., Poli, P., Campillo, M., & Roux, P. (2014). Reverberations, coda waves and ambient noise: Correlations at the global scale and retrieval of the deep phases. *Earth and Planetary Science Letters*, *391*, 137–145. https://doi.org/10.1016/j.epsl.2014.01.047

Choi, M.-K., Bobet, A., & Pyrak-Nolte, L. J. (2014). The effect of surface roughness and mixed-mode loading on the stiffness ratio, Kx/Kz for fractures. *Geophysics*, *79*(5), D319–D331. https://doi.org/10.1190/geo2013-0438.1

de Figueiredo, J. J. S., Schleicher, J., Stewart, R. R., & Dyaur, N. (2012). Estimating fracture orientation from elastic-wave propagation: An ultrasonic experimental approach. *Journal of Geophysical Research*, *117*, B08304. https://doi.org/10.1029/2012JB009215

Eaton, D. W., Igonin, N., Poulin, A., Weir, R., Zhang, H., Pellegrino, S., & Rodriguez, G. (2018). Induced seismicity characterization during hydraulic-fracture monitoring with a shallow-wellbore geophone array and broadband sensors. *Seismological Research Letters*, *89*(5), 1641–1651. https://doi.org/10.1785/0220180055

Fehler, M. C., Asanuma, H., & Jupe, A. (2001). More than cloud: New techniques for characterizing reservoir structure using induced seismicity. *The Leading Edge*, *20*, 324–328. https://doi.org/10.1190/1.1438942

Geller, R. J., & Mueller, C. S. (1980). Four similar earthquakes in central California. *Geophysical Research Letters*, *7*(10), 821–824. https://doi.org/10.1029/gl007i010p00821

Hadziioannou, C., Larose, E., Coutant, O., Roux, P., & Campillo, M. (2009). Stability of monitoring weak changes in multiply scattering media with ambient noise correlation: Laboratory experiments. *Journal of the Acoustical Society of America*, *125*, 3688–3695. https://doi.org/10.1121/1.3125345

He, H. Q., Chen, M., Chen, T., & Li, D. J. (2018). Matching of remote sensing images with complex background variations via siamese convolutional neural network. *Remote Sensing*, *10*(2), 355. https://doi.org/10.3390/rs10020355

Hildyard, M. W., & Young, R. P. (2002). Modelling seismic waves around underground openings in fractured rock. *Pure and Applied Geophysics*, *159*(1–3), 247–276. https://doi.org/10.1007/978-3-0348-8179-1_12

Holtzman, B. K., Paté, A., Paisley, J., Waldhauser, F., & Repetto, D. (2018). Machine learning reveals cyclic changes in seismic source spectra in Geysers geothermal field. *Scientific Advances*, *4*(5), eaao2929. https://doi.org/10.1126/sciadv.aao2929

Hopkins, D. L. (1990). Ph.D. thesis*The effect of surface roughness on joint stiffness, aperture and acoustic wave propagation*. Berkeley: University of California.

Hopkins, D. L., Cook, N. G. W., & Myer, L. R. (1990). Normal joint stiffness as a function of spatial geometry and surface roughnes. paper presented at *International Symposium on Rock Joints*. Loen, Norway.

Lopez-Comino, J. A., Cesca, S., Niemz, P., Dahm, T., & Zang, A. (2021). Rupture directivity in 3D inferred from acoustic emissions events in a mine-scale hydraulic fracturing experiment. *Frontiers of Earth Science*, *9*, 670757. https://doi.org/10.3389/feart.2021.670757

Lubbe, R., Sothcott, J., Worthington, M. H., & McCann, C. (2008). Laboratory estimates of normal and shear fracture compliance. *Geophysical Prospecting*, *56*(2), 239–247. https://doi.org/10.1111/j.1365-2478.2007.00688.x

Lubbe, R., & Worthington, M. H. (2006). A field investigation of fracture compliance. *Geophysical Prospecting*, *54*(3), 319–331. https://doi.org/10.1111/j.1365-2478.2006.00530.x

Marone, C., Vidale, J. E., & Ellsworth, W. (1995). Fault healing inferred from time dependent variations in soruce properties of repeating earthquakes. *Geophysical Research Letters*, *22*(22), 3095–3098. https://doi.org/10.1029/95gl03076

Marsland, S. (2015). *Machine learning, an algorithmic perspective* (2nd ed.). London: Chapman and Hall/CRC.

Murphy, K. P. (2012). *Machine learning: A probabilistic prespective*. Cambridge, Massachusetts: The MIT Press.

Nadeau, R. M., Foxall, W., & McEvilly, T. V. (1995). Clustering and periodic recurrence of microearthquakes on the san Andreas fault at Parkfield, California. *Science*, *267*(5197), 503–507. https://doi.org/10.1126/science.267.5197.503

Nakagawa, S., Nihei, K. T., & Myer, L. R. (2000a). Shear-induced conversion of seismic waves across single fractures. *International Journal of Rock Mechanics and Mining Sciences*, *37*(1–2), 203–218. https://doi.org/10.1016/s1365-1609(99)00101-x

Nakagawa, S., Nihei, K. T., & Myer, L. R. (2000b). Stop-pass behavior of acoustic waves in a 1D fractured system. *Journal of the Acoustical Society of America*, *107*(1), 40–50. https://doi.org/10.1121/1.428292

Nakagawa, S., Nihei, K. T., & Myer, L. R. (2002). Elastic wave propagation along a set of parallel fractures. *Geophysical Research Letters*, *29*(16), 571–590. https://doi.org/10.1029/2002gl014925

Nakagawa, S., Nihei, K. T., & Myer, L. R. (2004). Plane wave solution for elastic wave scattering by a heterogeneous fracture. *Journal of the Acoustical Society of America*, *115*(6), 2761–2772. https://doi.org/10.1121/1.1739483

Nakagawa, S., Nihei, K. T., Myer, L. R., & Majer, E. L. (2003). Three-dimensional elastic wave scattering by a layer containing vertical periodic fractures. *Journal of the Acoustical Society of America*, *113*(6), 3012–3023. https://doi.org/10.1121/1.1572139

Oliger, A., Nolte, D. D., & Pyrak-Nolte, L. J. (2003). Seismic focusing by a single planar fracture. *Geophysical Research Letters*, *30*(5), 1203. https://doi.org/10.1029/2002gl016264

Olivier, G., Brenguier, F., Campillo, M., Lynch, R., & Roux, P. (2015). Body-wave reconstruction from ambient seismic noise correlations in an underground mine. *Geophysics*, *80*(3), KS11–KS25. https://doi.org/10.1190/geo2014-0299.1

Peng, Z., & Ben-Zion, Y. (2006). Temporal changes of shallow seismic velocty around the Karadere-Duzce branch of the North Anatolian faulr and strong ground motion. *Pure and Applied Geophysics*, *163*, 567–600. https://doi.org/10.1007/s00024-005-0034-6

Petrovitch, C. L., Nolte, D. D., & Pyrak-Nolte, L. J. (2013). Scaling of fluid flow versus fracture stiffness. *Geophysical Research Letters*, *40*, 2076–2080. https://doi.org/10.1002/grl.50479

Petrovitch, C. L., Pyrak-Nolte, L. J., & Nolte, D. D. (2014). Combined scaling of fluid flow and seismic stiffness in single fractures. *Rock Mechanics and Rock Engineering*, *47*, 1613–1623. https://doi.org/10.1007/s00603-014-0591-z

Pyrak-Nolte, L. J. (2007). Fracture anisotropy: The role of fracture-stiffness gradients. *The Leading Edge*, *26*(9), 1124–1127. https://doi.org/10.1190/1.2780781

Pyrak-Nolte, L. J. (2019). Chapter 14. Fracture specific stiffness: The critical link between the scaling behavior of hydro-mechanical coupling in fractures and seismic monitoring. In P. Newell, & A. Ilgen (Eds.), *Science of carbon storage in deep saline formations: Process coupling across time and spatial scales* (pp. 233–251). Elsevier.

Pyrak-Nolte, L. J., Braverman, W., Nolte, N. J., Wright, A., & Nolte, D. D. (2020). Probing complex geophysical geometries with chattering dust. *Nature Communications*, *11*, 5282. https://doi.org/10.1038/s41467-020-19087-z

Pyrak-Nolte, L. J., Myer, L. R., & Cook, N. G. W. (1990). Transmission of seismic waves across single natural fractures. *Journal of Geophysical Research*, *95*(B6), 8617–8638. https://doi.org/10.1029/jb095ib06p08617

Pyrak-Nolte, L. J., & Nolte, D. D. (1992). Frequency-dependence of fracture stiffness. *Geophysical Research Letters*, *19*(3), 325–328. https://doi.org/10.1029/91gl03179

Pyrak-Nolte, L. J., & Nolte, D. D. (1995). Wavelet analysis of velocity dispersion of elastic interface waves propagating along a fracture. *Geophysical Research Letters*, *22*(11), 1329–1332. https://doi.org/10.1029/95gl01323

Pyrak-Nolte, L. J., & Nolte, D. D. (2016). Approaching a universal scaling relationship between fracture stiffness and fluid flow. *Nature Communications*, *7*, 10663. https://doi.org/10.1038/ncomms10663

Rodriguez-Pradilla, G., & Eaton, D. W. (2020). Automated microseismic processing and integrated interpretation of induced seismicity during a multistage hydraulic-fracturing stimulation. Alberta, Canada: Bulletin of Seismologicl Society of America, 110, 2018–2030. https://doi.org/10.1785/0120200082

Rowe, C. A., Aster, R. C., Phillips, W. S., Jones, R. H., Borchers, B., & Fehler, M. C. (2002). Using automated, high-precision repicking to improve delineation of microseismic structures at the Soultz geothermal reservoir. *Pure and Applied Geophysics*, *159*, 563–596. https://doi.org/10.1007/pl00001265

Sang, G., Liu, S., & Elsworth, D. (2020). Quantifying fatigue-damage and failure-precursors using ultrasonic coda wave interferometry. *International Journal of Rock Mechanics and Mining Sciences*, *131*, 104366. https://doi.org/10.1016/j.ijrmms.2020.104366

Schaff, D. P., & Beroza, G. C. (2004). Coseismic and postseismic velocity changes measured by repeating earthquakes. *Journal of Geophysical Research*, *109*, B10302. https://doi.org/10.1029/2004JB003011

Schoenberg, M. (1980). Elastic wave behavior across linear slip interfaces. *Journal of the Acoustical Society of America*, *5*(68), 1516–1521. https://doi.org/10.1121/1.385077

Shao, S., Petrovitch, C. L., & Pyrak-Nolte, L. J. (2015). Wave guiding in fractured layered media. In S. M. Agar & S. Geiger (Eds.), (Eds.,) *SP406 fundamental controls on fluid flow in carbonates: Current workflows to emerging technologies* (in press). Geological Society of London. https://doi.org/10.1144/SP406.8

Shao, S., & Pyrak-Nolte, L. J. (2016). Elastic wave propagation in isotropic media with orthogonal fracture sets. *Rock Mechanics and Rock Engineering*, *49*, 16. https://doi.org/10.1007/s00603-016-1084-z

Sneider, R., Gret, A., Douma, H., & Scales, J. (2013). Coda wave interferometry for estimating nonlinear behavior in seismic velocity. *Science*, *295*, 2253–2255.

Uchida, N., & Burgmann, R. (2019). Repeating earthquakes. *Annual Review of Earth and Planetary Sciences*, *47*, 305–332. https://doi.org/10.1146/annurev-earth-053018-060119

Willis, M. E., Burns, D. R., Rao, R., Minsley, B., Toksoz, M. N., & Verti, L. (2006). Spatial orientation and distribution of reservoir fractures from scattered seismic energy. *Geophysics*, *71*(5), P043–P051. https://doi.org/10.1190/1.2235977

Xian, C. J., Nolte, D. D., & Pyrak-Nolte, L. J. (2001a). Compressional waves guided between parallel fractures. *International Journal of Rock Mechanics and Mining Sciences*, *38*(6), 765–776. https://doi.org/10.1016/s1365-1609(01)00041-7

Xian, C. J., Nolte, D. D., & Pyrak-Nolte, L. J. (2001b). Evanescent compressional-mode guided waves propagating between parallel fractures. *Rock Mechanics in the National Interest*, *1 and 2*, 837–842.

Zheng, L. L., Duffner, S., Idrissi, K., Garcia, C., & Baskurt, A. (2016). Siamese multi-layer perceptrons for dimensionality reduction and face identification. *Multimedia Tools and Applications*, *75*(9), 5055–5073. https://doi.org/10.1007/s11042-015-2847-3

Zhu, T., Ajo-Franklin, J., Daley, T. M., & Marone, C. (2019). Dynamics of geologic CO_2 storage and plime motion revealed by seismic coda waves. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(7), 2464–2469. https://doi.org/10.1073/pnas.1810903116