

1 **Unsupervised Deep Clustering of Seismic Data:**
2 **Monitoring the Ross Ice Shelf, Antarctica**

3 **William F. Jenkins II¹, Peter Gerstoft¹, Michael J. Bianco¹, Peter D.**
4 **Bromirski¹**

5 ¹Scripps Institution of Oceanography, University of California San Diego, La Jolla, CA, USA

6 **Key Points:**

- 7 • Deep clustering identified classes of seismic signals with similar spectral and tem-
8 poral features.
- 9 • Deep clustering can be adapted to various kinds of data sets, enabling rapid ex-
10 ploration of “big data” in seismology.
- 11 • Paired with environmental data, deep clustering could provide insights into the
12 causes of seismicity.

Abstract

Advances in machine learning (ML) techniques and computational capacity have yielded state-of-the-art methodologies for processing, sorting, and analyzing large seismic data sets. In this work, we consider an application of ML for automatically identifying dominant types of impulsive seismicity contained in observations from a 34-station broadband seismic array deployed on the Ross Ice Shelf (RIS), Antarctica from 2014 to 2017. The RIS seismic data contain signals and noise generated by many glaciological processes that are useful for monitoring the integrity and dynamics of ice shelves. Deep clustering was employed to efficiently investigate these signals. Deep clustering automatically groups signals into hypothetical classes without the need for manual labeling, allowing for comparison of their signal characteristics and spatial and temporal distribution with potential source mechanisms. The method uses spectrograms as input and encodes their salient features into a lower-dimensional latent representation using an autoencoder, a type of deep neural network. For comparison, two clustering methods are applied to the latent data: a Gaussian mixture model (GMM) and deep embedded clustering (DEC). Eight classes of dominant seismic signals were identified and compared with environmental data such as temperature, wind speed, tides, and sea ice concentration. The greatest seismicity levels occurred at the RIS front during the 2016 El Niño summer, and near grounding zones near the front throughout the deployment. We demonstrate the spatial and temporal association of certain classes of seismicity with seasonal changes at the RIS front, and with tidally driven seismicity at Roosevelt Island.

Plain Language Summary

We demonstrate the ability of a machine learning technique called deep clustering to automatically identify different types of seismic signals. A neural network encodes spectrograms into simplified representations. Application of a clustering algorithm separates the representations into distinct clusters of signal types. The deep clustering technique was applied to seismic data recorded by an extensive array of broadband seismometers deployed on the Ross Ice Shelf (RIS), Antarctica from 2014 to 2017. In addition to knowing when and where on the RIS signals are detected, clustering enables users to determine the signal characteristics. Paired with environmental data, deep clustering can be used to identify whether certain environmental factors are associated with particular classes of seismicity.

1 Introduction

Ice sheets and ice shelves in West Antarctica are experiencing rapid change. Between 2003 and 2019, the West Antarctic Ice Sheet (WAIS) experienced a net ice loss of 169 billion tons per year, contributing 7.5 mm to sea level rise (Smith et al., 2020). Warming oceans are enhancing basal melting of ice shelves that reduces the buttressing of grounded ice sheets (De Angelis & Skvarca, 2003; Thoma et al., 2008; Pritchard et al., 2012; Paolo et al., 2015), leading to increased discharge of ice into the ocean and raising sea level (Scambos, 2004; Dupont & Alley, 2005; Rignot et al., 2014; Fürst et al., 2016). With West Antarctica alone containing a sea level rise potential of 5.6 m (Smith et al., 2020), monitoring the loss of ice shelves plays a critical role in anticipating future sea level rise and associated societal impacts on coastlines and the environment. Increased seismic activity, such as icequakes resulting from fracturing, can give indications of changes in iceberg calving rates and the integrity of ice shelves and are observable using glacial seismology methods (Aster & Winberry, 2017). However, the prevalence of extensive, continuously recording seismic observing systems has led to an abundance of data which is becoming increasingly difficult to analyze using conventional signal processing. At the same time, advances in computing capabilities and machine learning algorithms have enabled more efficient, data-driven approaches to study natural processes and phenomena. To analyze large seismic data sets more efficiently, we adapt contemporary machine learning techniques to augment existing signal processing and data analysis techniques.

Seismology is a data-intensive field with well-developed signal processing and analytical methods. The recent introduction of machine learning techniques has led to the development of complementary tools that give seismologists novel approaches to traditional analyses, such as earthquake detection and early warning, phase picking, ground-motion prediction, tomography, and geodesy (Kong et al., 2019; Bianco & Gerstoft, 2018; Bianco et al., 2019; Johnson et al., 2019). In this study we present an implementation of *clustering*, a form of unsupervised machine learning used to discover classes of similar signals within a data set (Bishop, 2006; Holtzman et al., 2018; Johnson et al., 2020), and which is commonly used as an exploratory tool for large, unlabeled data sets.

To test the applicability of clustering groups of similar signals for monitoring ice shelves, we focus specifically on the Ross Ice Shelf (RIS), Antarctica, where a 34-station passive seismic array was deployed from November 2014 to January 2017 to observe the

77 response of the RIS to ocean gravity wave impacts and investigate the structural dynam-
78 ics of the ice shelf (Bromirski et al., 2015). The array, shown in Figure 1, continuously
79 recorded long- and short-period seismic signals that exhibited seasonal and spatial vari-
80 ations related to the shelf’s coupling to the ocean, atmosphere, and crust (Baker et al.,
81 2019). Signals and ambient noise of interest on the RIS include tidally-driven stick-slip
82 seismicity at Whillans Ice Stream (Bindschadler, King, et al., 2003; Bindschadler, Vorn-
83 berger, et al., 2003; D. A. Wiens et al., 2008); basal micro-earthquakes and tremor (Barcheck
84 et al., 2018); tidally and thermally driven rift fractures (Olinger et al., 2019); diurnal seis-
85 micity associated with subsurface melting (MacAyeal et al., 2019); wind-generated res-
86 onance in the ice (Chaput et al., 2018); flexural and plate waves generated by ocean swell,
87 infragravity waves, and tsunami (Bromirski & Stephen, 2012; Bromirski et al., 2017; Chen
88 et al., 2018); regional and teleseismic earthquakes (Baker et al., 2020); and icequakes gen-
89 erated by ocean gravity waves (Chen et al., 2019). Ambient seismic noise, which can be
90 used to estimate the RIS structure (Diez et al., 2016), also contains spectra from ocean
91 gravity waves, whose dispersion can be used to identify their source distance and origin
92 (Bromirski et al., 2015; Hell et al., 2019).

93 The seismic data recorded on the RIS are diverse and encompass numerous source
94 mechanisms with a wide range of spatiotemporal variability. In this study, we apply two
95 unsupervised clustering methodologies to the RIS array seismic data to identify classes
96 of seismic events with similar temporal and spectral characteristics. The occurrences and
97 distributions of these signal classes provide information on glaciological processes affect-
98 ing ice shelf evolution.

99 **2 Background**

100 Grouping seismic signals with similar characteristics (clustering) allows investiga-
101 tion of spatiotemporal variability associated with glaciological processes that result from
102 environmental forcing.

103 **2.1 Clustering**

104 There are numerous methods to cluster data, (Aggarwal & Reddy, 2014), many of
105 which have been adapted for use in seismology and geophysics (Kong et al., 2019). A
106 related approach based on sparse modeling, called dictionary learning, has been applied

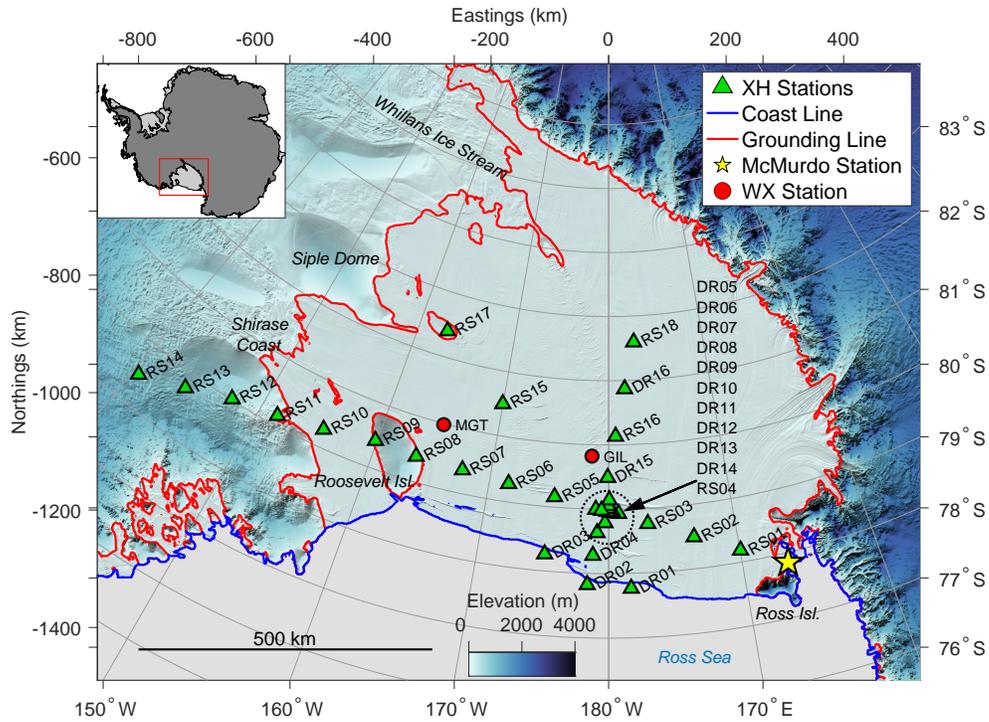


Figure 1. The passive broadband seismic array deployed from November 2014 to January 2017 consisted of 34 seismic stations and was deployed as part of the Ross Ice Shelf Dynamic Response to Wave-Induced Vibrations Project (Bromirski et al., 2015). RIS surface elevation, ice and water layer thicknesses, and grounding and coast lines were obtained from Bedmachine (Morlighem et al., 2017; Greene et al., 2017).

107 to regularizing seismic inverse problems (Bianco & Gerstoft, 2018; Bianco et al., 2019).
 108 Hierarchical clustering has been used by Mousavi et al. (2016) to automatically discrim-
 109 inate between shallow and deep earthquakes, and by Trugman and Shearer (2017) to more
 110 precisely localize earthquakes. Graphical clustering has been used to localize sources in
 111 a dense seismic array by Riahi and Gerstoft (2017), and by Telesca and Chelidze (2018)
 112 to cluster seismic events in time. Distance-based clustering, like the popular k -means al-
 113 gorithm, (MacQueen, 1967; Hartigan & Wong, 1979) has been used by Chamarczuk et
 114 al. (2020) to cluster seismicity based on features extracted from seismic data. Perol et
 115 al. (2018) used k -means to define probabilistic earthquake locations as part of their con-
 116 volutional neural network (CNN) detection and localization technique. Wallet and Hardisty
 117 (2019) used Gaussian mixture model (GMM) clustering, which assumes clusters in the
 118 data exist that can be represented as linearly superimposed Gaussian distributions, en-
 119 abling identification of seismic facies. Seydoux et al. (2020) detected and clustered seis-
 120 mic signals and background noise with the use of a deep scattering neural network and
 121 GMM.

122 Not all clustering methods involve machine learning. Template matching, in which
 123 a matched filter is constructed from a template waveform, is used to scan through con-
 124 tinuous recordings to locate similar signals (Gibbons & Ringdal, 2006; Beaucé et al., 2018;
 125 Chamberlain et al., 2018). Yoon et al. (2015) and Bergen and Beroza (2018) presented
 126 computationally efficient techniques in which locality-sensitive hashing is used to map
 127 seismic signals into a hash table, allowing similar signals to be identified by table entry.
 128 Hotovec-Ellis and Jeffries (2016) developed an approach that uses correlation-based sim-
 129 ilarity search to automatically detect and cluster repeating volcanic seismicity in con-
 130 tinuous data. Cole (2020) adopted the method of Hotovec-Ellis and Jeffries (2016) to clus-
 131 ter RIS array data at stations RS09, RS10, and RS11 in order to characterize tidal forc-
 132 ing of seismicity at these stations.

133 **2.2 Dimensionality**

134 Data are considered high-dimensional when many features are required to repre-
 135 sent or describe the data. Seismic data represented as time series, spectrograms, scalo-
 136 grams, or energy envelopes can contain thousands of features (e.g., discrete samples in
 137 a time series, or bins in a spectrogram). Clustering performed directly on such input data
 138 is vulnerable to the “curse of dimensionality” (Bellman, 1961; Bishop, 2006; Murphy, 2012;

139 Aggarwal & Reddy, 2014), i.e., as the dimensionality of the input data increases, the num-
 140 ber of data points required to maintain sufficient sampling density increases exponen-
 141 tially. A further consideration is that clustering error metrics can give less meaningful
 142 results as dimensionality increases.

143 As high-dimensional data are difficult to cluster (Aggarwal et al., 2001; Steinbach
 144 et al., 2004), dimensionality reduction remains a major focus of development (Yang et
 145 al., 2017). It is often desirable to transform the input data to a lower-dimensional rep-
 146 resentation described by fewer, more salient features. A popular approach is to use prin-
 147 cipal component analysis (PCA), which projects higher dimensional data into lower di-
 148 mensional space (Goodfellow et al., 2016) and was used by Reddy et al. (2012) to com-
 149 press seismic data to maximize feature variance.

150 The approach to reducing dimensionality in this study employs an autoencoder,
 151 a model whose output aims to reproduce its input via a series of non-linear transforma-
 152 tions employing a deep neural network (DNN) (Hinton, 2006; Murphy, 2012; Yang et al.,
 153 2017). These non-linear transformations provide greater capacity in dimension reduc-
 154 tion, and can better model data with low-dimensional representations than, for exam-
 155 ple, PCA. The autoencoder first encodes input data such as an image—in our case, a
 156 spectrogram—into a latent feature vector. Next, the autoencoder decodes the latent fea-
 157 tures and reconstructs the original image. Since the autoencoder provides a non-linear
 158 transformation of the data, it must be trained using gradient descent. In this iterative
 159 training, the error between the input and output is minimized. In doing so, the salient
 160 features of the data are learned by the network weights. With the dimensionality of the
 161 input data reduced in the latent feature space, clustering algorithms can be applied to
 162 the data’s latent feature space.

163 **2.3 Deep Embedded Clustering**

164 In deep clustering, a DNN such as an autoencoder is used to reduce the dimension-
 165 ality of the data. A recent deep clustering method that has shown improvement over tra-
 166 ditional clustering techniques was developed by Xie et al. (2016), whose *deep embedded*
 167 *clustering* (DEC) consists of two processes: (1) An autoencoder is trained to represent
 168 the data’s salient features; and (2) the encoding layers and clustering layer are jointly
 169 optimized. Yang et al. (2017) extended the approach in DEC by jointly optimizing the

170 clustering step with training the entire autoencoder, not just the encoder layers. Addi-
171 tional variations of DEC have been proposed: Xie et al. (2016) used a stacked denois-
172 ing autoencoder (Vincent et al., 2010) in their original implementation, but Min et al.
173 (2018) employed autoencoders composed of CNN layers and other architectures. More
174 recently, Chazan et al. (2019) developed an approach in which joint clustering is performed
175 with a mixture of autoencoders, each representing a cluster, and Boubekki et al. (2021)
176 demonstrated improved performance using a clustering algorithm that is jointly optimized
177 with the embeddings of the autoencoder.

178 Mousavi et al. (2019) used DEC to predict whether seismic detections were local
179 or teleseismic, and Snover et al. (2021) demonstrated the ability of DEC to cluster an-
180 thropogenically generated seismic noise. In a similar signal processing and clustering work-
181 flow to ours, Ozanich et al. (2021) compared DEC and GMM on spectrograms of acous-
182 tic data collected on a coral reef, but in their case found GMM performed better than
183 DEC.

184 In this study, we implement GMM clustering in the latent feature space and com-
185 pare its performance with DEC. Using RIS seismic data from December 2014 to Novem-
186 ber 2016, we identify several different classes of signals, and further demonstrate the util-
187 ity of deep clustering as an exploratory tool for large, real-world seismic data sets by as-
188 sociating the clustering results with observed environmental factors.

189 **3 Ross Ice Shelf (RIS) Seismic Array and Data**

190 Each station in the RIS seismic array consisted of 3-component Nanometrics Tril-
191 lium 120 PHQ seismometers emplaced 1 m below the surface of the ice, powered by solar
192 panels during the austral summers, and lithium-ion batteries during the austral win-
193 ters. Two subarrays comprised the array. The larger subarray consisted of 18 stations
194 spaced approximately 80 km apart (prefix RS), primarily oriented parallel to the RIS
195 front. The RS stations sampled short-period orthogonal components of ground veloc-
196 ity at a sampling rate of 100 Hz, except for two stations that sampled at 200 Hz. The
197 smaller subarray consisted of 16 stations (prefix DR) arranged approximately orthog-
198 onal to the ice shelf front along the international date line, sampling ground velocity with
199 a sampling rate of 200 Hz. For this study, we were primarily interested in the detection
200 and classification of icequakes and local/regional earthquakes, using only vertical com-

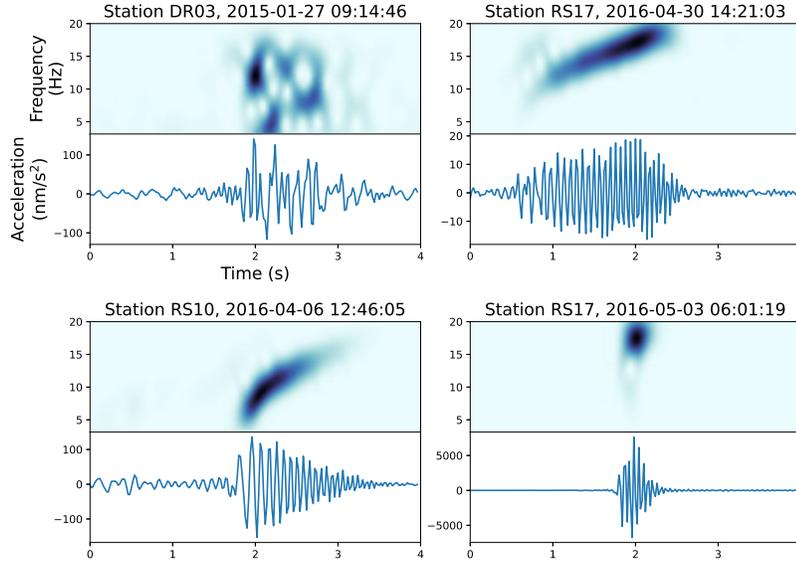


Figure 2. Seismic signals detected on the Ross Ice Shelf exhibited diverse characteristics with variation in time, space, and source mechanism. Shown are examples of acceleration response seismograms and their respective normalized spectrograms spanning the 3-20 Hz band that were typical for the data set. The normalized spectrograms were used as input to the deep clustering analysis.

201 ponent observations with frequencies of interest occurring between 3 and 20 Hz. This
 202 passband was selected to preserve impulsive signals, eliminate high-energy noise preva-
 203 lent at low frequencies, and exclude resonances generated by wind at frequencies above
 204 20 Hz. Representative types of signals detected are shown in Figure 2.

205 Seismic data from each station were processed in 24-hour segments as follows: 1) Data
 206 were linearly de-trended and tapered with a Hann window. 2) Instrument responses for
 207 all stations were removed, giving acceleration in m/s^2 . 3) Since the bandwidth of inter-
 208 est was from 3 to 20 Hz, data were decimated to 50 Hz, using low-pass filtering followed-
 209 by downsampling. 4) A band-pass filter with cutoff frequencies at 3 and 20 Hz was ap-
 210 plied to remove long-period signals originating from tides, tsunamis, infragravity waves,
 211 ocean swell, and teleseisms. 5) A short-term average/long-term average (STA/LTA) de-
 212 tection algorithm (Allen, 1982) was used to detect impulsive signals, particularly icequakes
 213 and local earthquakes, employing an STA window of 0.5 s, LTA window of 30 s, trigger
 214 threshold of 15, and de-trigger threshold of 10. The detector was applied to data from

215 each station from 3 December 2014 to 21 November 2016 for a total of 719 days of ar-
 216 ray data, yielding 531,407 detections.

217 Upon detection, a 4 s trace centered on the spectral peak of each triggered event
 218 was saved for processing. Centering the trace at the spectral peak yielded more unique
 219 clusters by preventing the clustering algorithm from labeling similar signals as different
 220 classes based only on their relation to the trigger time. For each seismic trace saved, a
 221 spectrogram was computed using the short-time Fourier transform with a 0.4 s Kaiser
 222 window, NFFT=256, and 90% overlap. Spectrograms (samples) contained one channel
 223 of amplitude information, 87 frequency bins, and 100 time bins for a total of 8,700 fea-
 224 tures per spectrogram. To improve DNN learning, sample-wise normalization was per-
 225 formed by dividing each spectrogram by its vector norm (LeCun et al., 2012).

226 4 Deep Clustering Implementation

227 The objective of deep clustering models is to first encode the input data—in this
 228 case, spectrograms of seismic signals—into a layer containing latent (lower-dimensional)
 229 features, called the *embedded* layer, and to then apply a clustering algorithm in this la-
 230 tent feature space. In the implementation that follows, the 8,700 features of an input spec-
 231 trogram are reduced to a latent feature space of just 9 embedded features with the use
 232 of a convolutional autoencoder, a type of DNN composed of convolutional and transposed
 233 convolutional layers. We then describe the GMM and DEC clustering algorithms that
 234 are used in the clustering analysis.

235 4.1 Dimensionality Reduction with a Convolutional Autoencoder

236 Autoencoders provide a useful means of data approximation using a lower-dimensional
 237 representation via a sequence of non-linear transformations. The autoencoder model con-
 238 sists of three components: an *encoder*, a *bottleneck*, and a *decoder* (Murphy, 2012). First,
 239 the encoder maps input data from a data space X into a latent feature space Z , which
 240 is contained within the bottleneck of the model. Next, the decoder attempts to recon-
 241 struct X from Z . This process is performed iteratively with the objective of minimiz-
 242 ing the error between X and the decoder output, X' . In minimizing the error, the au-
 243 toencoder learns the salient features of X and accurately encodes them in Z , thus re-
 244 ducing the dimensionality of the clustering task.

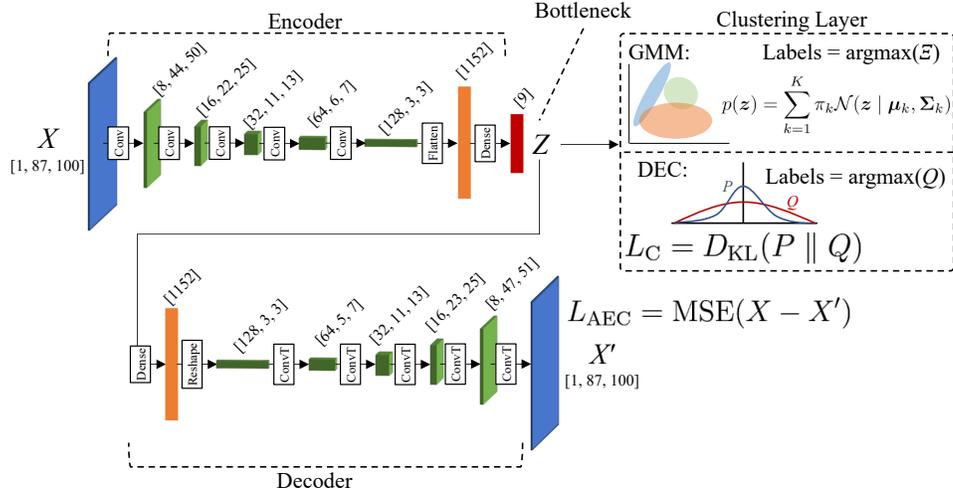


Figure 3. The deep clustering framework in this study uses a convolutional autoencoder that encodes the data space X into the latent feature space Z , and a decoder that recovers the original input X from Z . The mean squared error (MSE) between the input X and the reconstruction X' is used as the autoencoder loss function. The latent feature space Z lies at the bottleneck between the encoder and decoder, providing the input to the clustering layer. Gaussian mixture model (GMM) clustering labels each data sample according to its most likely cluster membership using an expectation-maximization algorithm. Deep embedded clustering (DEC) provides label assignments, and also outputs a clustering loss function that is combined with the MSE to further train the parameters that map $X \rightarrow Z \rightarrow X'$.

245 Consider a data set of spectrograms $\mathcal{D} = \{\mathbf{x}_n \in X^M\}_{n=1}^N$, where \mathbf{x}_n is a vector rep-
 246 resentation of the n^{th} spectrogram in a data set containing N spectrograms, and the num-
 247 ber of features in \mathbf{x}_n , M , is the spectrogram size (the product of the number of frequency
 248 bins and time bins). In the encoder stage, the mapping of X to Z is described by $f_\theta : X \rightarrow Z$,
 249 where θ are parameters that are learned through iterative model training. The decoder
 250 stage is a mirror operation of the encoder and seeks to map the latent feature space Z
 251 to the reconstruction X' by $g_\theta : Z \rightarrow X'$. The overall mapping of the autoencoder can
 252 be described as $F_\theta : X \rightarrow Z \rightarrow X'$, where $F_\theta = g_\theta \circ f_\theta$. Input spectrograms \mathbf{x}_n map to
 253 their corresponding latent feature vectors by $\mathbf{z}_n = f_\theta(\mathbf{x}_n) \in Z^D$, where D is the num-
 254 ber of embedded features, and to their reconstructions by $\mathbf{x}'_n = F_\theta(\mathbf{x}_n) \in X'$.

255 As the autoencoder is composed of convolutional and transposed convolutional lay-
 256 ers, F_θ is a nonlinear mapping that must be appropriately parameterized. This is accom-
 257 plished by iteratively learning the parameters θ in order to minimize the error between
 258 the input and reconstructed data. The mean squared error (MSE) between an input spec-
 259 trogram with M features and its reconstruction, defined as

$$260 \quad \ell(\mathbf{x}, \mathbf{x}') = \frac{1}{M} \sum_{m=1}^M (x_m - x'_m)^2, \quad (1)$$

261 is averaged over the N samples in the data set to obtain the autoencoder loss function:

$$262 \quad L_{\text{AEC}} = \frac{1}{N} \sum_{n=1}^N \ell(\mathbf{x}_n, \mathbf{x}'_n). \quad (2)$$

263 Performing this calculation over the entire data set at once is computationally expen-
 264 sive, memory intensive, and can lead to poor convergence. Instead, the loss is calculated
 265 in mini-batch subsets of the data space. For each mini-batch loss, stochastic gradient de-
 266 scent (Goodfellow et al., 2016) is used to update the weights. When all mini-batches have
 267 been processed, the next training epoch begins and the process is repeated. After each
 268 epoch, a subset of the data separate from the training data is used to validate the model’s
 269 performance without updating the weights, yielding a validation MSE. Training is per-
 270 formed until a specified maximum number of epochs is reached, or stopped early if the
 271 validation MSE fails to decrease below its minimum value after ten epochs. The early
 272 stopping criterion prevents the autoencoder from overfitting the training data.

273 The design choice of autoencoder architecture can be informed by prior knowledge
 274 of a data set and its features, as well as practical considerations such as computational
 275 resources available. Our DNN architecture, detailed in Table 1, is designed to be com-
 276 putationally efficient, simple to construct, and robust enough to learn salient features

Table 1. *Convolutional Autoencoder Architecture*

Layer Name	Type	Input Shape	Filters	Activation	Output Shape	Trainable Parameters
Input	-	-	-	-	[1, 87, 100]	-
Conv1	Convolution	[1, 87, 100]	8	ReLU	[8, 44, 50]	80
Conv2	Convolution	[8, 44, 50]	16	ReLU	[16, 22, 25]	1,168
Conv3	Convolution	[16, 22, 25]	32	ReLU	[32, 11, 13]	4,640
Conv4	Convolution	[32, 11, 13]	64	ReLU	[64, 6, 7]	18,496
Conv5	Convolution	[64, 6, 7]	128	ReLU	[128, 3, 3]	73,856
Flat	Flatten	[128, 3, 3]	-	-	[1152]	0
Encoded	Fully Connected	[1152]	-	ReLU	[9]	10,377
FC	Fully Connected	[9]	-	ReLU	[1152]	11,520
Reshape	Reshape	[1,152]	-	-	[128, 3, 3]	0
ConvT1	Transposed Conv	[128, 3, 3]	64	ReLU	[64, 5, 7]	73,792
ConvT2	Transposed Conv	[64, 5, 7]	32	ReLU	[32, 11, 13]	18,464
ConvT3	Transposed Conv	[32, 11, 13]	16	ReLU	[16, 23, 25]	4,624
ConvT4	Transposed Conv	[16, 23, 25]	8	ReLU	[8, 47, 51]	1,160
Decoded	Transposed Conv	[8, 47, 51]	1	Linear	[1, 95, 101]	73
Output	Crop	[1, 95, 101]	-	-	[1, 87, 100]	-
Total						218,250

Table 2. *Sample Sizes and Hyperparameters used to Train the Autoencoder and Deep Embedded Clustering Model*

Samples			Hyperparameters				
Total (N)	Training (N_{train})	Validation (N_{val})	Initial learning rate	Mini-batch size	Classes (K)	Clustering loss factor (λ)	Updates per epoch
531,407	40,000	10,000	10^{-3}	64	8	10^{-4}	10

277 from a noisy seismic data set. In total, θ contains 218,250 trainable parameters under
 278 this DNN architecture.

279 Autoencoder training is implemented using 50,000 spectrograms randomly selected
 280 without replacement from the 531,407 detections. Of the selected spectrograms, 80% are
 281 used for training and 20% for validation. The trainable parameters are optimized using
 282 the Adaptive Moment Estimation (Adam) algorithm (Kingma & Ba, 2017). In training,
 283 there are two principal hyperparameters to address. First is the initial learning rate, which
 284 controls the initial step size used by Adam to step down the gradient of the loss. The
 285 second hyperparameter is the mini-batch size, which sets the number of spectrograms
 286 to be passed through the model at one time. The optimal configuration is found through
 287 a grid search of the hyperparameters. A summary of the optimal hyperparameters and
 288 the number of spectrograms used are listed in Table 2. As seen in Figure 4a, training
 289 and validation losses fall off exponentially with each training epoch until the early stop-
 290 ping criterion is met; in this case, at 48 epochs. The effectiveness of the autoencoder’s
 291 ability to reconstruct the input spectrogram is illustrated in Figure 5. Though some loss
 292 of resolution in time and frequency is expected due to the convolutional and transposed
 293 convolutional layers, the structure of the spectrogram is largely preserved, with the salient
 294 information of the input encoded to the latent feature space. To test that the autoen-
 295 coder adequately generalized the entire data set, all spectrograms were fed through the
 296 model, yielding an average MSE of 5.9381×10^{-6} , which is consistent with the valida-
 297 tion MSE at the early stopping point.

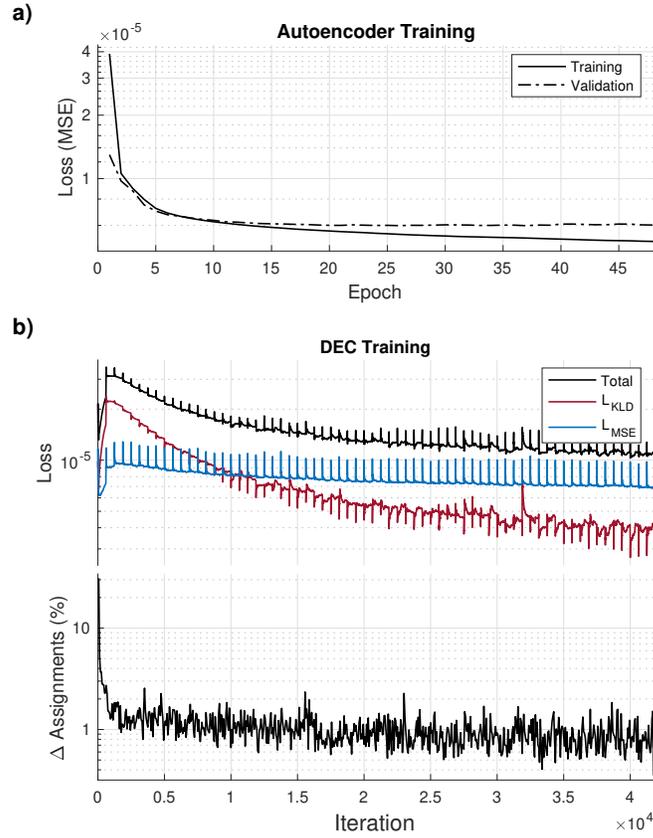


Figure 4. (a) Training and validation losses during autoencoder training. To avoid over-fitting the model, training is stopped when the early stopping criterion is met (in this case, at 48 epochs). (b) In the upper plot, loss curves are shown for deep embedded clustering (DEC). In the lower plot, the percentage of samples which undergo class reassignment at each update interval is shown; training is stopped once the change is less than 0.4%

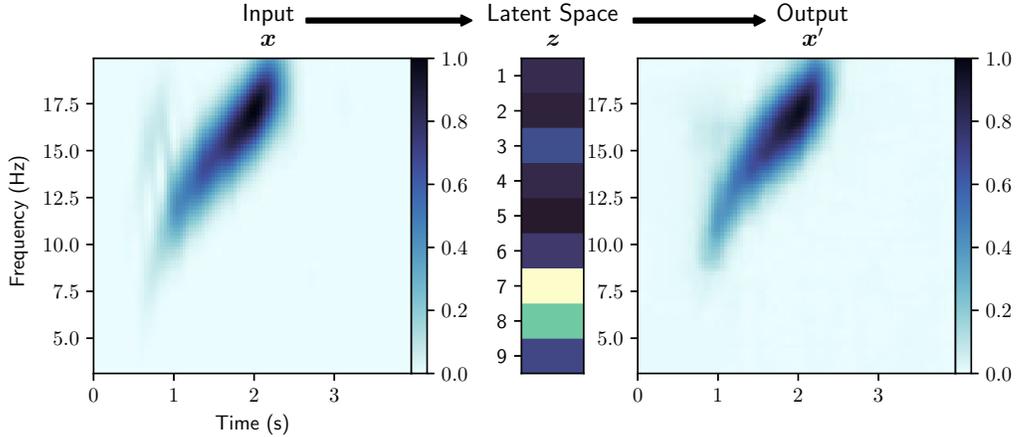


Figure 5. A trained autoencoder takes an input spectrogram \mathbf{x} , encodes it to a 9-dimensional latent feature vector \mathbf{z} , then reconstructs the input as \mathbf{x}' . The autoencoder preserves features correlated within a given cluster and discards the remaining signal, which can help with signal identification.

298

4.2 Clustering Methodologies

299

300

301

302

303

304

In our deep clustering framework, clustering is performed in the latent feature space, Z , to find K distinct classes of signals within the data. We assume that the data form clusters which are separable in Z space, and that these clusters coalesce around unique locations $\{\boldsymbol{\mu}_k \in Z\}_{k=1}^K$, i.e., centroids around which other similar signals may be found. We use Euclidean distance between a centroid and a latent feature vector to measure similarity:

305

$$d_{n,k} = \|\mathbf{z}_n - \boldsymbol{\mu}_k\|_2. \quad (3)$$

306

$d_{n,k}$ is a measure of the similarity between features indexed by n and k .

307

4.2.1 Gaussian Mixture Model (GMM)

308

309

310

311

312

In GMM clustering, the latent feature vectors \mathbf{z} are described by a mixture of K Gaussian distributions that are linearly superimposed in the latent space Z , where each Gaussian model has its own centroid $\boldsymbol{\mu}_k$ and covariance $\boldsymbol{\Sigma}_k$. We follow the methods of Bishop (2006, p. 430) and Murphy (2012, p. 339). The overall distribution of the mixture model is given by the convex combination of their distributions,

313

$$p(\mathbf{z}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{z} \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k). \quad (4)$$

314 Consider the latent feature vectors \mathbf{z}_n as rows of a matrix $\mathbf{Z} \in \mathbb{R}^{N \times D}$ with N sam-
 315 ples and D features. To estimate the parameters of each Gaussian distribution, an expectation-
 316 maximization (EM) algorithm is used to maximize the Gaussian mixture model's like-
 317 lihood function of \mathbf{Z} with respect to the parameters $\boldsymbol{\mu}_k$, $\boldsymbol{\Sigma}_k$, and π_k (Bishop, 2006, p. 433):

$$318 \quad \ln p(\mathbf{Z} \mid \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_K\}, \{\boldsymbol{\Sigma}_1, \dots, \boldsymbol{\Sigma}_K\}, \{\pi_1, \dots, \pi_K\}) = \sum_{n=1}^N \ln \left[\sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{z}_n \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right]. \quad (5)$$

319 For every sample \mathbf{z}_n , a binary K -dimensional random variable $\xi_k \in \{0, 1\}$ is intro-
 320 duced that has one element equal to one and all others to zero. The marginal distribu-
 321 tion over $\boldsymbol{\xi}$ is $p(\xi_k = 1) = \pi_k$, where the mixing coefficients π_k satisfy $0 \leq \pi_k \leq 1$ and $\sum_{k=1}^K \pi_k = 1$
 322 in order to be valid probabilities. Since $\boldsymbol{\xi}$ is a 1-of- K (categorical) representation, this
 323 distribution is written as

$$324 \quad p(\boldsymbol{\xi}) = \prod_{k=1}^K \pi_k^{\xi_k}, \quad (6)$$

325 and the conditional distribution of \mathbf{z}_n given $\boldsymbol{\xi}$ as

$$326 \quad p(\mathbf{z}_n \mid \boldsymbol{\xi}) = \prod_{k=1}^K \mathcal{N}(\mathbf{z}_n \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)^{\xi_k}. \quad (7)$$

327 Equation (4) is then rewritten in terms of the factored joint distribution $p(\mathbf{z}_n, \boldsymbol{\xi}) = p(\boldsymbol{\xi})p(\mathbf{z}_n \mid \boldsymbol{\xi})$:

$$328 \quad p(\mathbf{z}_n) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{z}_n \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) = \sum_{\boldsymbol{\xi}} p(\boldsymbol{\xi})p(\mathbf{z}_n \mid \boldsymbol{\xi}). \quad (8)$$

329 Using Bayes' theorem and equations (4) and (8), the conditional probability of $\boldsymbol{\xi}$ given
 330 \mathbf{z}_n is:

$$331 \quad \gamma(\xi_k) \equiv p(\xi_k = 1 \mid \mathbf{z}_n) = \frac{p(\xi_k = 1)p(\mathbf{z}_n \mid \xi_k = 1)}{\sum_{j=1}^K p(\xi_j = 1)p(\mathbf{z}_n \mid \xi_j = 1)} = \frac{\pi_k \mathcal{N}(\mathbf{z}_n \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{z}_n \mid \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j)}, \quad (9)$$

332 where π_k is the prior probability of $\xi_k = 1$, and $\gamma(\xi_k)$ is the posterior probability hav-
 333 ing observed \mathbf{z}_n . As with \mathbf{Z} , we construct a matrix $\boldsymbol{\Xi} \in \mathbb{R}^{N \times K}$ whose rows consist of
 334 the binary random variables ξ_n for each sample \mathbf{z}_n . Thus indexed, $\gamma(\xi_{nk})$ is defined as
 335 the *responsibility* that distribution k has for *explaining* sample \mathbf{z}_n , and is analogous to
 336 soft clustering, where the probability that sample \mathbf{z}_n belongs to distribution k is deter-
 337 mined for each of the K distributions. In practice, each latent feature vector \mathbf{z}_n is as-
 338 signed to one of K Gaussian distributions by $\arg \max_{\xi} [\gamma(\xi_{nk})]$.

339 Using superscript t to denote the iteration index, the EM algorithm for a Gaus-
 340 sian mixture is:

- 341 1. Initialization of parameters $\boldsymbol{\mu}_k^{t-1}$, $\boldsymbol{\Sigma}_k^{t-1}$, and π_k^{t-1} .

342 2. Expectation step. This step encodes the samples’ probability of assignment to each
 343 Gaussian distribution by evaluating responsibilities $\gamma(\xi_{nk})$ using $\boldsymbol{\mu}_k^{t-1}$, $\boldsymbol{\Sigma}_k^{t-1}$, and π_k^{t-1}
 344 (equation (9)).

345 3. Maximization step. Using the responsibilities $\gamma(\xi_{nk})$, this step updates the centroid
 346 location ($\boldsymbol{\mu}_k^t$), shape ($\boldsymbol{\Sigma}_k^t$), and normalization (π_k^t) of each distribution in the latent space
 347 Z by:

$$\begin{aligned}
 \boldsymbol{\mu}_k^t &= \frac{1}{N_k} \sum_{n=1}^N \gamma(\xi_{nk}) \mathbf{z}_n \\
 \boldsymbol{\Sigma}_k^t &= \frac{1}{N_k} \sum_{n=1}^N \gamma(\xi_{nk}) (\mathbf{z}_n - \boldsymbol{\mu}_k^t)(\mathbf{z}_n - \boldsymbol{\mu}_k^t)^T \\
 \pi_k^t &= \frac{N_k}{N}
 \end{aligned} \tag{10}$$

where

$$N_k = \sum_{n=1}^N \gamma(\xi_{nk}).$$

349 4. Convergence check. The log likelihood of \mathbf{Z} is evaluated with respect to the param-
 350 eters $\boldsymbol{\mu}_k^t$, $\boldsymbol{\Sigma}_k^t$, and π_k^t (equation 5). If convergence occurs in the log likelihood or in the
 351 parameters $\boldsymbol{\mu}_k^t$, $\boldsymbol{\Sigma}_k^t$, and π_k^t , the EM algorithm has reached a local maximum and ter-
 352 minates; otherwise, the algorithm returns to step 2.

353 To accelerate EM convergence, k -means clustering is used to initialize the GMM
 354 clustering algorithm (Bishop, 2006, p. 438). EM stops after 1,000 iterations have elapsed
 355 or when the change in log likelihood from equation (5) is less than 0.001. To avoid con-
 356 verging on local maxima, the initialization is run 100 times and the initialization with
 357 the best log likelihood is retained.

358 **4.2.2 Deep Embedded Clustering (DEC)**

359 In DEC, clustering is performed in conjunction with continued training of the au-
 360 toencoder, with the clustering layer attached to the bottleneck providing an additional
 361 loss function that is backpropagated through the autoencoder layers (Figure 3). The DEC
 362 model DNN parameters are initialized using the parameters of the trained autoencoder,
 363 and clustering layer parameters are initialized using the centroids from GMM cluster-
 364 ing. DEC seeks to improve the GMM clustering by using the Euclidean distance between
 365 embedded spectrograms and cluster centroids (equation (3)) as an additional loss func-
 366 tion for updating model parameters. Because the input data is unlabeled, a self-supervised

method is required. We implement the method developed by Xie et al. (2016), who, drawing from the t-distributed stochastic neighbor embedding (t-SNE) algorithm (van der Maaten & Hinton, 2008), propose measuring the difference between a Student’s t-distribution kernel of the latent feature vectors \mathbf{z} and an auxiliary target distribution. A simplified Student’s t-distribution is used to measure the similarity between embedded spectrograms \mathbf{z}_n and the cluster centroids $\boldsymbol{\mu}_k$:

$$q_{nk} = \frac{(1 + \|\mathbf{z}_n - \boldsymbol{\mu}_k\|^2)^{-1}}{\sum_k (1 + \|\mathbf{z}_n - \boldsymbol{\mu}_k\|^2)^{-1}}. \quad (11)$$

Equation (11) results in a set of soft class assignments, i.e., the probability that embedded spectrogram n will be assigned to class k . Latent feature vectors \mathbf{z}_n are assigned to one of K classes by $\arg \max_q [q_{nk}]$. The soft class assignments q_{nk} are then used to compute the auxiliary target distribution, p , whose form is designed to improve clustering performance, emphasize embeddings with high-confidence assignments, and normalize each cluster centroid’s contribution to the loss function so that large clusters minimally distort Z (Xie et al., 2016):

$$p_{nk} = \frac{q_{nk}^2 / \sum_n q_{nk}}{\sum_k (q_{nk}^2 / \sum_n q_{nk})}. \quad (12)$$

The dissimilarity between the distributions given by equations (11) and (12) is measured using the Kullback-Leibler divergence (Kullback & Leibler, 1951). From the divergence the clustering layer’s loss function is obtained:

$$L_C = D_{\text{KL}}(P \parallel Q) = \sum_n \sum_k p_{nk} \log \frac{p_{nk}}{q_{nk}}. \quad (13)$$

In DEC, the clustering layer is attached to the trained autoencoder’s bottleneck. During training of the DEC model, the loss functions from equations (2) and (13) are combined into a total loss function,

$$L = L_{\text{AEC}} + \lambda L_C, \quad (14)$$

where λ is a hyperparameter that balances the contributions of the two losses, since they are of differing magnitudes. λ must be tuned: if it is too large, the clustering loss will cause model instability and lead to distortion of the latent space, in which case the latent space will no longer represent the salient features of the data. If λ is too small, the effect on clustering performance will be minimal. We found that $\lambda = 10^{-4}$ yielded optimal performance for model training and clustering.

Two constituent processes occur simultaneously during DEC model training. First, the full loss from equation (14) is backpropagated through the DEC model parameters,

398 which include the autoencoder as well as the cluster centroids. Second, to account for
 399 the cluster centroids changing as training progresses, the distributions q_{nk} and p_{nk} are
 400 updated at intervals. The update interval is a hyperparameter that must be tuned. Through
 401 hyperparameter tuning, an update interval of 10 per training epoch was found to be op-
 402 timal for clustering performance, minimizing DEC loss, and training within a reasonable
 403 time frame. Training is stopped after the number of samples changing assignments af-
 404 ter every update interval reaches less than 0.4% of the total number of training samples.
 405 The same mini-batch size and initial learning rate are used to train both the autoencoder
 406 and DEC model (Table 2). Figure 4b shows how losses decrease over time and the per-
 407 cent change in label assignments for every mini-batch training iteration. Though the over-
 408 all trends in the loss curves show exponential decay, periodic spikes occur at every up-
 409 date interval, when q_{nk} and p_{nk} are recalculated, and are visible since the losses are recorded
 410 after every mini-batch rather than every epoch.

411 **4.3 Selecting Optimal Number of Clusters**

412 Determining the optimal number of clusters, K , is a major challenge in unsuper-
 413 vised machine learning. In this study we treat K as a hyperparameter, iterating the deep
 414 clustering workflow over a range of values for K and evaluating the results to choose the
 415 best value. Results are evaluated both quantitatively and qualitatively. Quantitative eval-
 416 uation is performed for each class by examining cumulative distribution functions and
 417 probability density functions as functions of distance to each class centroid, $d_{n,k}$ (equa-
 418 tion (3)). Additionally, traditional statistical methods for choosing the optimal number
 419 of clusters, such as the gap statistic (Tibshirani et al., 2001) and silhouette score (Rousseeuw,
 420 1987), are consulted. The qualitative approach is to visually inspect the similarity of the
 421 latent feature vectors \mathbf{z}_n to their respective class centroids $\boldsymbol{\mu}_k$, and to see if the spec-
 422 trograms and seismograms assigned to each class likewise exhibit similarity. In general,
 423 the formation of two or more similar classes may indicate that too many classes were ini-
 424 tialized, and the data in those classes can be grouped into a single class in post-processing.
 425 Too much variance among the spectrograms within a class may indicate the need for one
 426 or more additional classes. We found that $K = 8$ was the optimal number of classes for
 427 the RIS data set.

5 Results

The following analysis of GMM and DEC performance focuses on how the clustering algorithms affect the latent space Z and whether the methods yield meaningful results in the data space X . Since the samples in the data set are unlabeled and there is no “ground truth” against which to compare results, measurements of intra-class similarity among spectrograms and latent feature vectors are examined. We conclude that neither GMM nor DEC provides a clear advantage in clustering performance. Accordingly, we recommend implementation of GMM for deep clustering of RIS seismic data. The statistical and mathematical underpinnings of GMM are well understood, and the complexity of implementation and interpretation of DEC is difficult to justify in the absence of compelling performance improvement. Furthermore, in practice GMM clustering on a graphics processing unit takes approximately one minute to cluster the entire data set, whereas one DEC hyperparameter tuning run can take several hours.

In the analyses that follow, results are presented for the entire data set of 531,407 spectrograms, including the training and validation data subsets. We mitigate the risk of the DNN in the DEC model overfitting on the training data (Murphy, 2012, p. 23) by using less than 10% of the data set for training and validation, and by drawing training samples randomly without replacement to achieve a training subset representative of the entire data set.

5.1 Clustering Performance

Deep clustering performance is qualitatively checked by comparing centroids to their respective assigned latent data samples. Results for GMM are shown in Figure 6. Each class k is represented by the columns in Figure 6, with each centroid $\boldsymbol{\mu}_k$ and its reconstruction $g_\theta(\boldsymbol{\mu}_k)$ plotted along the top row. Although the centroid is not a member of the data set, because the centroid represents the salient features of its class, its reconstruction is expected to resemble the spectrograms \boldsymbol{x}_n assigned to its class. Subsequent rows show the latent feature vectors \boldsymbol{z}_n , spectrograms \boldsymbol{x}_n , and associated seismograms of the data samples assigned to the respective classes. To inspect whether intra-class similarity holds with increasing distance from the centroid, samples \boldsymbol{z}_n and \boldsymbol{x}_n are shown for $n = \{1, 1000, 5000, 10000, 15000, 20000, 25000\}$. Near the centroid, latent feature vectors \boldsymbol{z}_n generally exhibit similar values to their class centroid $\boldsymbol{\mu}_k$, indicating that GMM

459 has successfully grouped similar latent data samples into the class, and that the centroid
 460 is representative of the data in its class. The spectrograms in each class are likewise sim-
 461 ilar to each other and to the centroid reconstruction $g_{\theta}(\boldsymbol{\mu}_k)$, confirming that the latent
 462 features embedded in the centroids are representative of the spectrograms in the class.
 463 Finally, the similarity in the latent space and time-frequency domain extends to the time
 464 domain, where seismograms in each class are similar to one another. As distance increases
 465 (i.e., with increasing n), cases of dissimilarity begin to arise as samples overlap with ad-
 466 jacent clusters.

467 In addition to checking the efficacy of the clustering, visual examination of the re-
 468 sults in Figure 6 gives indication of whether or not an appropriate number of clusters
 469 was chosen. For example, classes 4 and 8 exhibit similar characteristics in time and fre-
 470 quency, distinct from each other primarily in peak amplitude characteristics. If such dis-
 471 tinctions are not useful or if similarities are redundant, classes can be combined in post-
 472 processing. If too few clusters are selected, classes may contain widely differing signals,
 473 indicating the need to increase the number of clusters.

474 Clustering with DEC involves two steps: first, the GMM clustering algorithm ini-
 475 tializes the centroids, but the latent data are left unmodified. Second, during DEC, cen-
 476 troids are further refined while the latent data are moved much closer to their respec-
 477 tive centroids, with some data reassigned to different classes altogether. To determine
 478 to what extent this occurs, t-SNE is used to visualize the 9-dimensional latent space in
 479 two dimensions (van der Maaten & Hinton, 2008). t-SNE can illuminate possible clus-
 480 ters within data in an unsupervised manner by displaying data in geometrically sepa-
 481 rated clusters. In Figure 7a, t-SNE results of the latent feature space clustered with GMM
 482 show that the data are largely contiguous with few exceptions. Applying the labels as-
 483 signed by GMM clustering to the data points shows that, while there is some geomet-
 484 ric separation between the clusters, the embedding is characterized by overlapping and
 485 dispersed class members, indicating poor separation in the latent space. Contrast this
 486 with Figure 7b, in which t-SNE results at the conclusion of DEC show both geometric
 487 separation as well as nearly homogeneous class assignments.

488 While t-SNE offers an intuitively visual way to look for clusters in data, results are
 489 sometimes difficult to interpret and are impossible to reproduce exactly due to the in-
 490 herent randomness of the algorithm. Running t-SNE iteratively and with the same ran-

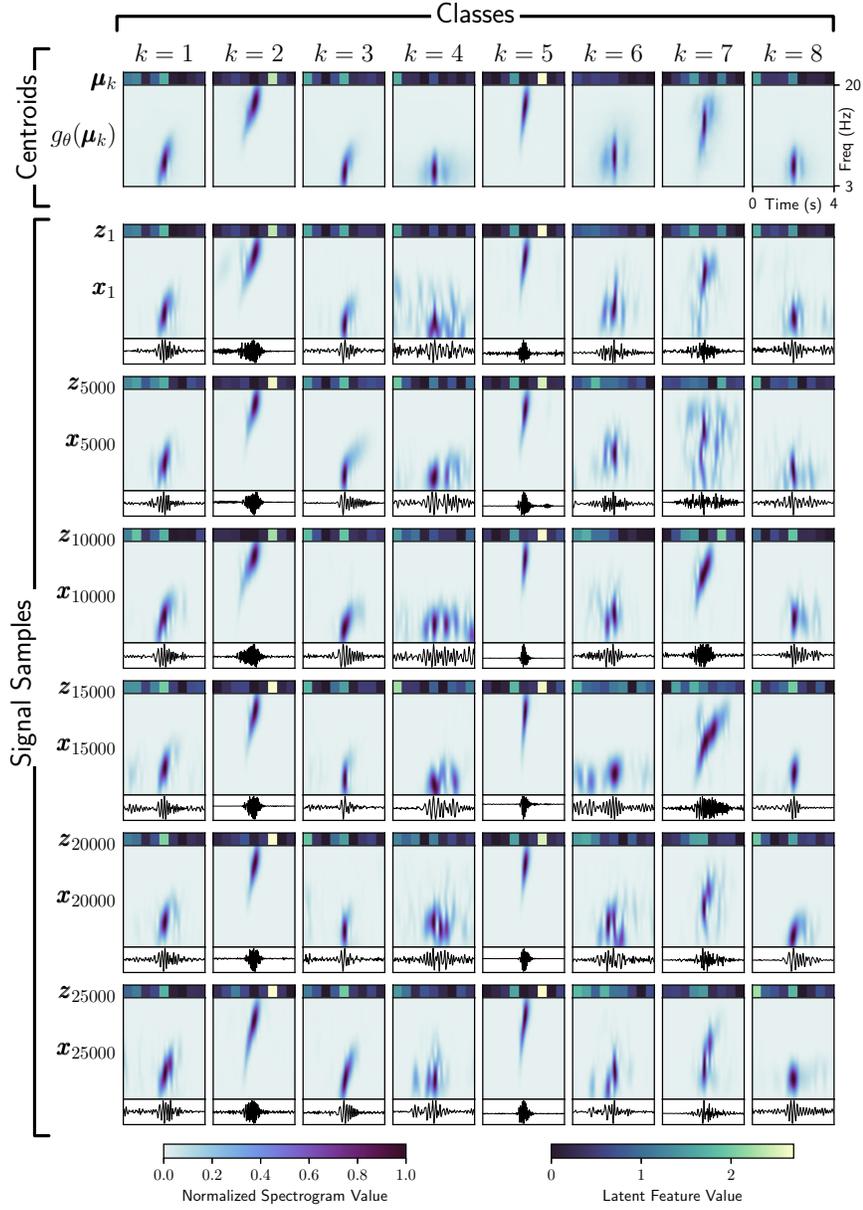


Figure 6. Gaussian mixture model (GMM) clustering results are shown, with samples z_n and x_n the n^{th} closest to their respective centroids. Within a given class k , the cluster centroids μ_k are similar to the latent feature vectors z_n , whose nine elements are shown above each spectrogram. Though the centroids are not members of the data set, their reconstructions $g_\theta(\mu_k)$ exhibit similar characteristics to the spectrograms x_n assigned to each class. Seismograms plotted below each spectrogram also exhibit similarity within each class. With increasing distance from the centroid (i.e., as n increases), dissimilarity and potential cases of mis-assignment are visible in latent feature vectors, spectrograms, and seismograms, e.g for $k = 7$, $n = 15000$.

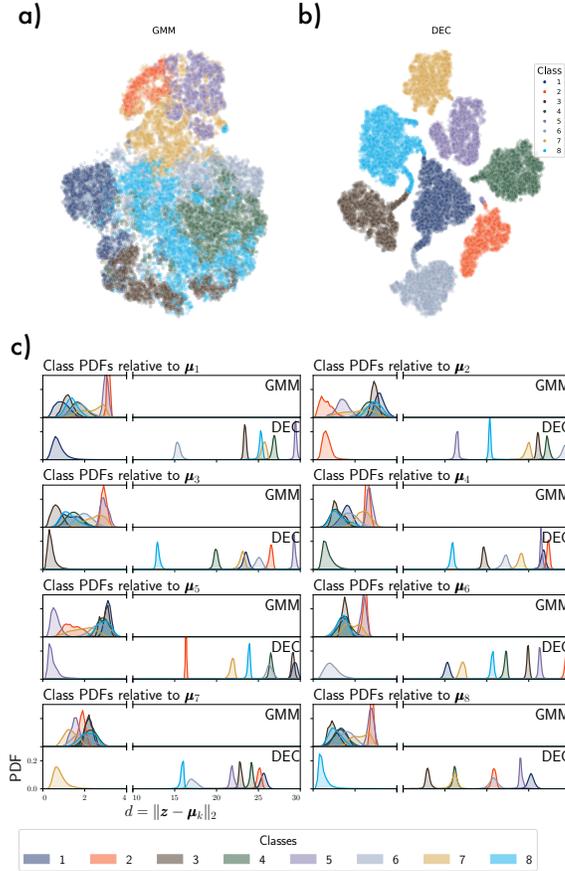


Figure 7. (a) Visualization of the 9-dimensional latent data space is shown in two dimensions using the t-distributed stochastic neighbor embedding (t-SNE) plot for Gaussian mixture model (GMM) clustering. GMM exhibits limited separation within the data and overlapping classes. (b) t-SNE plot for deep embedded clustering (DEC), whose clusters are well separated and contain nearly homogeneous class members. (c) The effects of DEC in the latent feature space are evident for each class probability density function (PDF) with respect to the distance from the centroids. In addition to moving the assigned class members closer to the centroid, DEC increases the distance between the other class centroids and PDFs.

491 dom seed can mitigate these limitations, but examination of the effects of deep cluster-
 492 ing on the densities of the clusters provides a more concrete visualization. Of interest
 493 to the ability for the clustering algorithms to identify clusters is the distance of each cluster
 494 to the others. In Figure 7c, the probability density functions (PDF) of all clusters
 495 are shown as functions of distance to each centroid. Before DEC, though GMM cluster-
 496 ing usually results in the PDF of each class being closest to its centroid, there is signif-
 497 icant overlap with other clusters, and the clusters themselves are not particularly dense.
 498 With DEC, the PDF of each class is closer to its centroid, denser, and farther removed
 499 from the other clusters. Thus, DEC effectively separates each cluster from the others,
 500 allowing for better distinction between clusters in the latent space.

501 The effects of DEC become readily apparent when the latent feature vectors are
 502 stacked and sorted according to their distance from each centroid, as shown in Figure 8.
 503 By sorting the latent space by sample index n such that $d_{n+1,k} > d_{n,k}$, cluster sepa-
 504 ration can be visualized directly in the latent space. Before DEC, centroids are initial-
 505 ized with the GMM clustering algorithm without modification to the latent data. Clos-
 506 est to each class centroid, the latent feature vectors are similar in appearance to the cen-
 507 troid, but transition continuously to different patterns as the sorted index n increases.
 508 The contrast with the latent feature space after DEC is stark: because DEC moves la-
 509 tent data assigned to a particular class closer to the centroid, the effect is that the la-
 510 tent feature vectors take on similar values, and therefore appearance, to the centroid.
 511 The result is that the latent space appears more sharply segmented after DEC, with the
 512 samples closest to the centroid of nearly uniform appearance to the centroid itself. For
 513 reference, the relative location of the other class centroids are marked with white ver-
 514 tical lines. With GMM, the latent feature vectors belonging to the other classes are not
 515 readily apparent, whereas after DEC, most of the other centroid locations are associated
 516 with their distinctive latent feature vectors.

517 While DEC effectively transforms the latent feature space Z by moving latent fea-
 518 ture vectors closer to their centroids, less clear is whether this transformation causes a
 519 corresponding improvement in clustering quality in the data space X . To evaluate intra-
 520 class similarity among spectrograms, four pairwise metrics are used to compare the cluster-
 521 ing assignments obtained from GMM and DEC.

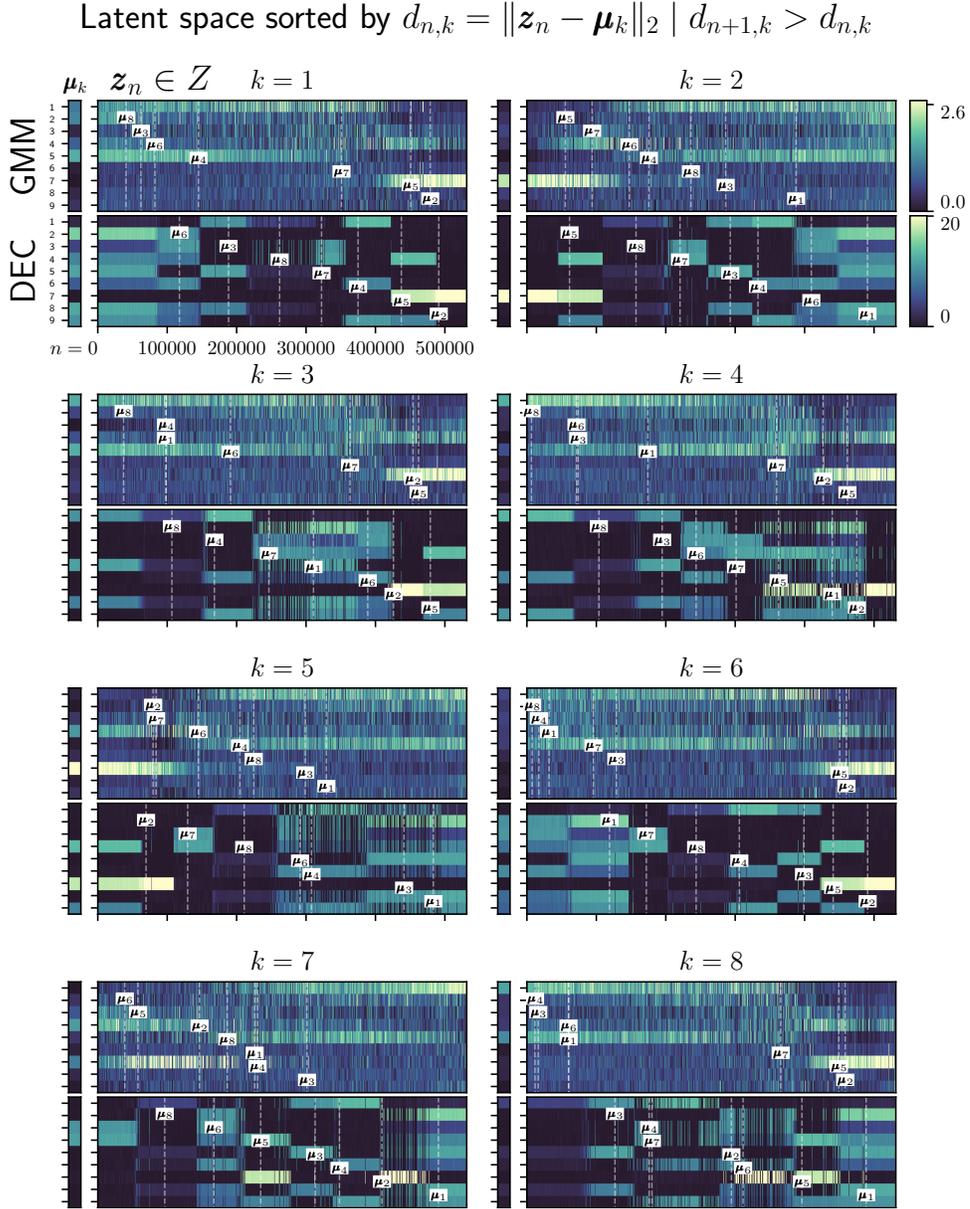


Figure 8. For each class k , latent data samples \mathbf{z}_n are shown stacked according to their distance $\|\mathbf{z}_n - \boldsymbol{\mu}_k\|$ from the centroid $\boldsymbol{\mu}_k$ (shown to the left). Distance of the other cluster centroids relative to the selected class k are indicated with vertical dotted lines. Deep embedded clustering (DEC) brings assigned data \mathbf{z}_n closer to the class centroid, resulting in homogeneity among the latent feature vectors assigned to that class.

522 The first metric used is the silhouette coefficient, which uses the mean intra-cluster
 523 and nearest-cluster distances to express whether a sample belongs in its assigned cluster
 524 or if it is more similar to another cluster (Rousseeuw, 1987). The silhouette coefficient
 525 exists on the interval $[-1, 1]$, with positive values indicating a sample has likely been
 526 correctly assigned, values near 0 indicating overlapping clusters, and negative values indicating
 527 a sample may have been placed in the wrong cluster. Coefficients are calculated
 528 for every sample, and the silhouette score is defined as the mean of all the coefficients.
 529 A summary of class and total silhouette scores is given in Table 3. In Figure 9, silhouette
 530 analyses are shown stacked by cluster assignment for the latent feature data in Z
 531 for GMM (Figure 9a) and DEC (Figure 9b), and for the spectrograms in the data space
 532 X for GMM (Figure 9c) and DEC (Figure 9d). In Figure 9a, classes 1-3 and 5 are distinctly
 533 clustered, classes 4, 6, and 7 are likely in a region of overlap, and class 8 is not
 534 well clustered; the silhouette score for this data is 0.08. In contrast, every class in Figure
 535 9b is well clustered with a silhouette score of 0.90, results which are consistent with
 536 those presented in Figures 7 and 8. To determine whether these analyses correspond to
 537 meaningful results in the data space, we examine the correlation between the silhouette
 538 analyses of the latent space Z and data space X . The silhouette analysis for GMM in
 539 the data space is shown in Figure 9c with a silhouette score of 0.05. These results are
 540 consistent with the GMM latent space results in Figure 9a and indicate a proper mapping
 541 from the data space into the latent space with the autoencoder. The silhouette analysis
 542 for DEC in the data space is shown in Figure 9d with a silhouette score of 0.13, which
 543 is inconsistent with its corresponding latent space analysis in Figure 9b. Comparison between
 544 Figures 9c and Figures 9d might lead us to conclude that DEC provides superior
 545 clustering performance, and this may be true. However, the inconsistency observed for
 546 DEC between the latent space and the data space require that additional metrics be examined.
 547

548 For the remaining metrics, spectrograms \mathbf{x}_n are vectorized and divided by their vector
 549 norm, resulting in unit vectors projected onto an n -sphere. The second metric is obtained
 550 by taking the inner product between two such unit vectors, which provides a measure of the
 551 angle between them and thus a proxy for similarity. The third metric is MSE, but to mitigate
 552 its tendency to exaggerate the effects of outliers by squaring the error, the mean absolute
 553 error (MAE) is used as a fourth metric. For each of these metrics, an intra-class mean
 554 vector is calculated against which all other vectors in the class are mea-

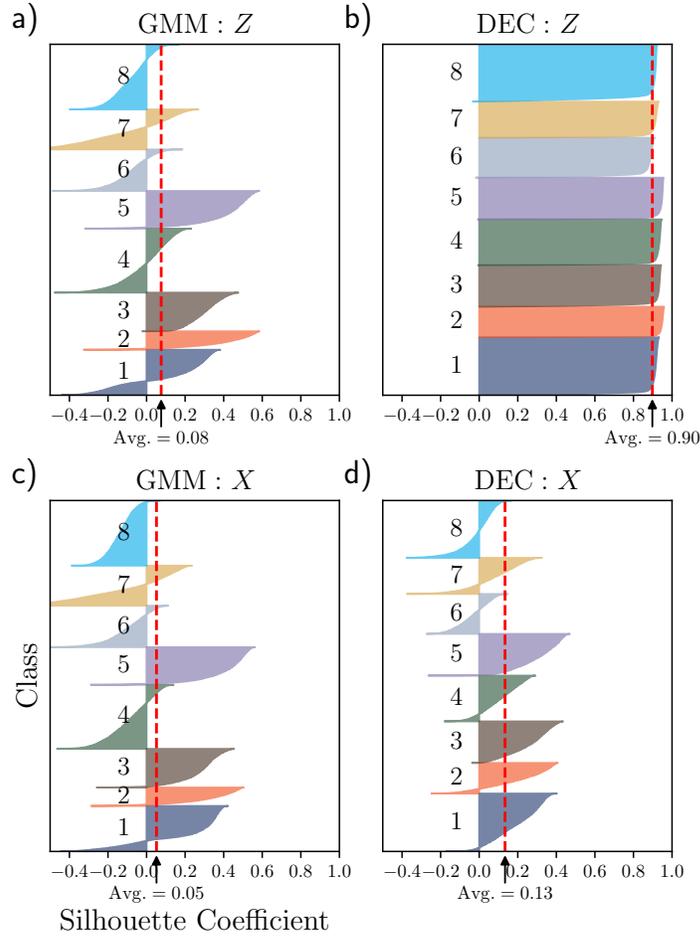


Figure 9. Silhouette analyses for (a,c) Gaussian mixture model (GMM) clustering and (b,d) deep embedded clustering (DEC) for the (a, b) latent feature space Z and (c,d) data space X .

Table 3. Comparison of Clustering Metrics for Gaussian Mixture Model (GMM) Clustering and Deep Embedded Clustering (DEC)

Class	N	Data Space			Latent Space	
		Mean Inner Product	Mean MSE ($\times 10^{-5}$)	Mean MAE ($\times 10^{-3}$)	Silhouette Score	Silhouette Score
1	66817 / 85789	0.82 / 0.80	0.26 / 0.28	0.20 / 0.23	0.19 / 0.20	0.11 / 0.89
2	27568 / 45607	0.88 / 0.81	0.44 / 0.55	0.36 / 0.44	0.31 / 0.20	0.39 / 0.93
3	59131 / 63725	0.86 / 0.87	0.64 / 0.74	0.53 / 0.61	0.27 / 0.26	0.30 / 0.90
4	95323 / 68521	0.61 / 0.73	1.21 / 1.13	0.90 / 0.90	-0.08 / 0.11	0.00 / 0.92
5	57318 / 64235	0.91 / 0.85	1.33 / 1.35	1.01 / 1.05	0.41 / 0.30	0.41 / 0.93
6	63326 / 59925	0.49 / 0.64	2.06 / 1.87	1.48 / 1.43	-0.10 / -0.03	-0.08 / 0.85
7	61430 / 55699	0.48 / 0.57	2.81 / 2.49	1.88 / 1.82	-0.09 / -0.08	-0.08 / 0.89
8	98494 / 87906	0.67 / 0.76	3.29 / 2.84	2.16 / 2.10	-0.14 / -0.01	-0.08 / 0.87
Overall Mean:		0.71* / 0.75*	1.50* / 1.41*	1.06* / 1.07*	0.05 / 0.13	0.08 / 0.90

All table values read as GMM / DEC. *Weighted mean.

555 sured. The class and total mean values for each metric for GMM and DEC are given in
556 Table 3, with better scores in bold. While Figures 7, 8, and 9b,d may lead us to favor
557 DEC performance, the data space metrics in Table 3 offer a more nuanced understand-
558 ing. On average, DEC slightly outperforms GMM in the mean inner product, MSE, and
559 silhouette score. Importantly, however, the inconsistencies among the metrics within each
560 class preclude a definitive decision regarding which clustering method is better. Of par-
561 ticular concern is the disparity in latent space and data space results for DEC. The la-
562 tent space transformation in DEC is substantial and does lead to sharp, distinct clus-
563 ters in the latent space. However, it appears these results do not map into the data space
564 so readily. We assess that this disparity arises when the DEC model is training: as the
565 model parameters are updated, the latent space is continually manipulated to conform
566 to the class centroids, effectively distorting the latent space. Even through hyperparam-
567 eter tuning, we were unable to obtain results that provided a compelling reason to jus-
568 tify the complexity of DEC, especially within the context of initial data exploration, in

569 which GMM is more efficient. Consequently, results shown in the subsequent sections
 570 are from the GMM deep clustering workflow.

571 **5.2 Deep Clustering Methodology Considerations**

572 One of the key strengths of the deep clustering implementation in this study is the
 573 employment of an autoencoder to reduce the dimensionality of the input data to obtain
 574 more effective clustering performance. By reducing the dimensionality of the data space,
 575 the complexity of the clustering problem is similarly decreased and the distance metrics
 576 gain relevance. The ability of the autoencoder to quickly learn the salient features of the
 577 data and embed them into the latent space makes the technique adaptable to new data
 578 sets. While the autoencoder design choice for this study was sufficiently robust, autoen-
 579 coder design presents opportunities for further experimentation and improvement. De-
 580 sign variables that could be altered in the DNN architecture include the number and types
 581 of layers, dimensions of the latent feature space, activation function types, incorporation
 582 of max-pooling and drop-out layers, and filter size, depth, and stride.

583 The selection of an appropriate algorithm for the clustering layer largely depends
 584 on the type and properties of the data set. Though in this study we use GMM and DEC,
 585 as described in Section 2, there are numerous clustering algorithms of which some may
 586 be applicable to a deep clustering workflow. Regardless of the choice of clustering algo-
 587 rithm, careful consideration must be given towards understanding whether clustering in
 588 the latent space maps to meaningful results in the data space.

589 The flexibility afforded by deep clustering extends not only to model design, but
 590 also to data pre- and post-processing. Whereas model design is largely concerned with
 591 *how* the salient features are learned, data pre-processing is concerned with *what* is sup-
 592 plied to the model. This information is dependent on the choice of signal processing pa-
 593 rameters, particularly signal duration, filter cutoff frequencies, and seismic event detec-
 594 tion algorithm. Additionally, various data transforms commonly used to characterize seis-
 595 mic waveforms can be used as input to deep clustering workflows (Mousavi et al., 2016).
 596 In our case, we used spectrograms, but other transforms, such as continuous wavelet trans-
 597 form scalograms, could just as easily be used as inputs. In post-processing, redundant
 598 or similar results can be combined.

6 Discussion: Glaciological Implications

The spatial and temporal distribution of signals from the eight classes identified gives information on the response of the RIS to various climatological forcings, including from oceanographic and atmospheric variability. Importantly, two years of continuous seismic monitoring allows identification of seasonal and interannual patterns of variability, particularly allowing examination of the effects of the strong 2016 El Niño on RIS seismicity by comparisons with 2015 levels.

The two-year RIS array data set contains 531,407 seismic detections. A summary of the data set statistics and class characteristics (Table 4) shows the total number of detections for each class, as well as the percentage of detections occurring in the austral summers (January, February, and March) versus the austral winters (June, July and August). Classes 2, 4, 5, 6, and 8 have pronounced differences (more than 10%) between the number of detections occurring in the summers versus the winters, while differences for classes 1, 3, and 7 are less pronounced (between 5% and 10%). Interannual comparisons for each season show that classes 5, 6, and 7 experienced an increase in activity in the 2016 austral summer over the 2015 austral summer, with classes 5 and 7 exhibiting the largest changes.

The seasonal changes are investigated in more detail in Figure 10a, where detection occurrences shown as a function of station and month exhibit spatiotemporal patterns that reveal associations between environmental forcing and seismicity. Clustering enables these patterns to be further explored by class and month (Figure 10b), and by class and station (Figure 10c).

From Figure 10a, certain patterns are readily apparent, such as increased seismic detections during the austral summer months at stations DR01, DR02, and DR03. These three stations were located approximately 2 km from the ice front (Figure 1) and detected seismicity associated with ocean gravity waves impacting the shelf front that cause fracturing (icequakes) and calving (Chen et al., 2019). Seismicity at these stations during the 2016 austral summer was higher than the same period in 2015, and across the array, a substantial increase in seismicity was observed in the months immediately following the 2016 austral summer, indicative of the impact of El Niño on Antarctic ice shelf fronts (Nicolas et al., 2017).

Table 4. Austral Summer (January-February-March) and Winter (June-July-August) Detection Statistics, Average Peak Frequencies, and Amplitude Characteristics for Each Signal Class over the Entire Seismic Array

Class	Detections						Freq (Hz)		Amplitude (accel., nm/s ²)			
	N	%N Summer (JFM)			%N Winter (JJA)			Mean peak	Mean	Median	Std. dev.	Max.
		Total	2015	2016	Total	2015	2016					
1	66,817	27	13	13	22	11	11	7.3	46	37	45	3,242
2	27,568	1	0	1	27	0	27	16.7	60	27	95	2,222
3	59,131	30	16	14	21	11	10	5.9	61	37	130	12,825
4	95,323	37	17	20	23	10	13	5.4	112	32	488	41,924
5	57,318	13	0	12	29	1	28	16.6	124	42	368	33,623
6	63,326	39	16	23	19	8	11	8.1	155	34	6,533	1,632,100
7	61,430	24	6	18	19	3	16	13.7	169	30	3,277	461,205
8	98,494	46	22	24	16	7	9	6.3	210	46	1,388	268,633

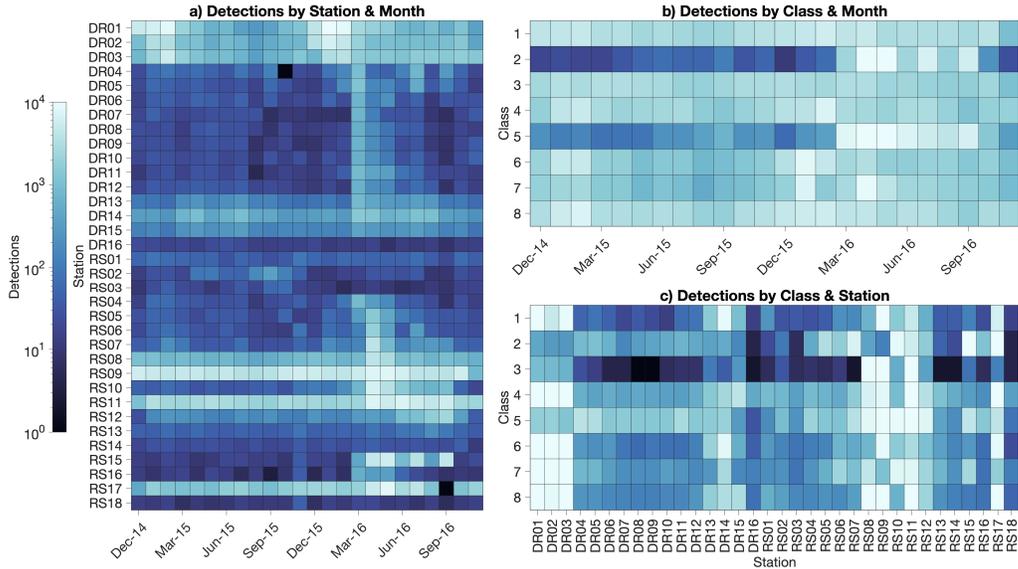


Figure 10. (a) The frequency of detections comprising the Ross Ice Shelf data set is shown by station and month. Clustering provides a further breakdown by (b) class and month for all stations, and (c) class and station.

630 Some of the most seismically active stations were located near grounding zones: sta-
 631 tion RS09 (118,105 detections) on the eastern flank of Roosevelt Island; station RS11
 632 (81,138 detections) on the Shirase Coast; station RS17 (50,385 detections) on Steershead
 633 Ice Rise; and station RS08 (25,500 detections) on the western flank of Roosevelt Island.
 634 These stations were on either fully or partially grounded ice, suggesting that the seis-
 635 micity results from interactions of basal ice with the solid earth. Increases in seismic-
 636 ity during the 2016 winter at floating stations RS10 (between Roosevelt Island and the
 637 Shirase Coast) and RS15 (over a bathymetric high) may result from El Niño related changes
 638 in water layer thickness that affect flexural gravity wave amplitudes (Bromirski et al.,
 639 2017). The RIS front stations DR01 (64,311 detections), DR02 (39,822 detections), and
 640 DR03 (39,176 detections) were also active. All of these active stations exhibited persis-
 641 tent seismicity throughout the two deployment years, with the exception of station RS17,
 642 which was offline for several weeks from August to September 2016.

643 Some classes of signal detections exhibit temporal patterns that are visible in Fig-
 644 ure 10b. Classes 2 and 5 have increased detection frequencies in the austral winter of 2015
 645 when local storms are more intense, suggesting meteorological forcing. The remaining
 646 classes have increased detections in the austral summers. The clustering results reveal
 647 that the large increase in seismicity in classes 2 and 5 occurs following the 2016 austral
 648 summer. A further dimension to the analysis is shown in Figure 10c, which shows the
 649 distribution of classes by station. Classes 1, 4, 5, and 8 are prominent signal types at sta-
 650 tions near grounding zones (RS08, RS09, RS11, and RS17), and classes 1, 4, 6, 7, and
 651 8 are prominent at the RIS front (DR01, DR02, DR03).

652 An important caveat for the detection statistics shown in Table 4 and Figure 10
 653 arises from the physics governing seismic propagation. For a given amplitude, low fre-
 654 quency seismic energy propagates farther than high frequency seismic energy. We thus
 655 expect the seismometers in the RIS array to detect low-frequency signals originating far-
 656 ther away than high-frequency signals. For example, from Figure 6, class 1 is similar to
 657 class 3, with the notable difference in that class 1 contains more energy at frequencies
 658 slightly higher than class 3 and has lower amplitude. Thus, class 3 may be generated by
 659 a similar source mechanism as class 1 but have a longer propagation path.

660 Factoring in signal amplitude also affects the range at which seismic energy is de-
 661 tected. From Table 4, class 2 has an average spectral peak at 16.7 Hz, the highest of the

662 classes, with a total of 27,568 detections, the lowest of the classes. Similarly, class 5 has
663 the second-highest average spectral peak at 16.6 Hz, with the second lowest amount of
664 detections among the classes. These two classes are nevertheless distinct from each other
665 in amplitude and waveform type: from Table 4, class 2 has a mean amplitude of 46 nm/s²,
666 while class 5 has a mean amplitude of 124 nm/s². From Figure 6, class 2 consists of high
667 frequency signals experiencing dispersion, while class 5 signals are more impulsive; both
668 likely result from fracturing.

669 Detection statistics are affected by signal-to-noise ratios at the seismometers and
670 by limitations of the automated seismic event detector, such as the inability to separate
671 signals from different classes that are received nearly simultaneously. Consideration should
672 also be given to determining if classes are duplicates of the same seismic source mech-
673 anism. Seismic surface waves in the ice undergo dispersion as they propagate, which DEC
674 may interpret as separate signal classes. This may be the case with classes 2 and 5.
675 The longer wave train for class 5 signals is consistent with Rayleigh wave propagation
676 of class 2 signals. Propagation modeling can be used to calculate expected dispersion re-
677 lations to confirm if this is the case. Such distinctions could be useful in identifying com-
678 mon propagation paths or providing source range discrimination.

679 Though the sources of uncertainty in the detection statistics are nontrivial, with
680 a proper understanding of these limitations and when paired with environmental data,
681 the clustering results can nevertheless be used to analyze the association of potential seis-
682 mic source mechanisms that may be related to ice shelf dynamics. In the following sec-
683 tions, we provide vignettes using stations DR02 and RS09 to demonstrate the utility of
684 deep clustering in exploring data and identifying potential causes of seismicity when ex-
685 amined in conjunction with environmental data.

686 **6.1 Seasonal seismicity at the RIS front**

687 Approximately 2 km from the RIS front on Nascent Iceberg, station DR02 exhibits
688 a seasonal pattern of seismicity associated with changes in air temperature and sea ice
689 concentration in the Ross Sea. During the austral winter, sea ice coverage (Figure 11a)
690 reaches nearly 100%, damping ocean swell. During the austral summer, sea ice concen-
691 tration decreases to approximately 25%, permitting ocean gravity waves to directly im-
692 pact the ice shelf front and cause iceberg calving. Additionally, warmer air temperatures

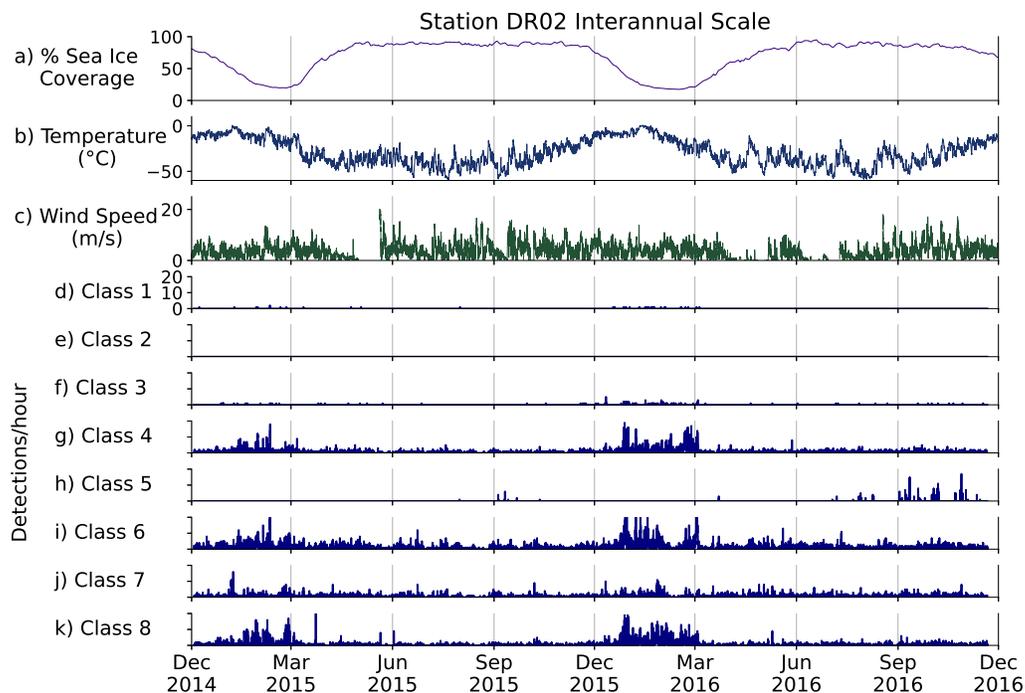


Figure 11. Two years of (a) sea ice coverage on the Ross Sea, (b) temperature and (c) wind speed observations at Gill automated weather station (approximately 223 km south of DR02, Figure 1), and (d-k) icequake detection statistics for each signal class. Classes 4, 6, 7, and 8 exhibit increased seismicity during the austral summers. Sea ice concentration data were obtained from NSIDC (Cavalieri et al., 1996, updated yearly); weather station data from AMRC, SSEC, UW–Madison.

693 (Figure 11b) may promote calving with associated increased icequake activity (Chen et
 694 al., 2019).

695 Increased levels of seismicity at DR02 are observed for all classes except 2 and 5
 696 at DR02 (Figure 11d,f,g,i-k) during the austral summers. Classes 4, 6, and 8 are espe-
 697 cially active during the 2016 austral summer, when strong El Niño conditions led to anoma-
 698 lously persistent high temperatures across West Antarctica (Nicolas et al., 2017) and ocean-
 699 ice shelf interactions were enhanced. Patterns similar to the seismicity at DR02 were ob-
 700 served at stations DR01 and DR03, also located near the RIS front, and can be seen in
 701 the total detections by station and month in Figure 10a. Widespread surface melt on
 702 the RIS was observed between 10-21 January 2016 (Nicolas et al., 2017; Chaput et al.,

2018), which affects firn layer properties and seismicity through freeze/thaw cycles (MacAyeal et al., 2019).

Although class 6 has elevated activity during the summers, it maintains activity throughout the winter months, suggesting that gravity wave activity is not the dominant forcing. The persistence of class 1 signals, which often consist of impulse trains, suggests they may be caused by icequakes resulting from the motion of the ice shelf itself (Klein et al., 2020), as the ice flow velocity in the vicinity of station DR02 is among the highest observed on the RIS. Class 5 (Figure 11h) is more active during the coldest periods of the year (April-September), suggesting that these signals may be associated with extremely cold temperatures or strong wind events. Cold-weather enhanced seismicity occurs at a rift approximately 140 km south of the ice front (Olinger et al., 2019). Alternatively, from Table 4, these classes are lower amplitude than those most active during the austral summer, which suggests that these detections may be masked by higher amplitude signals associated with the other classes. Across all classes, discrete instances of high seismicity occur that do not correspond to environmental forcing. Such instances may indicate the occurrence of fracturing ice (icequakes) or events associated with crevasse expansion.

6.2 Diurnal seismicity on Roosevelt Island

Station RS09 on the eastern flank of Roosevelt Island experienced the most detections across the array, comprising 22% of detections in the full data set. In Figure 12, potential environmental sources of seismicity are compared to the seismicity of each class. Temperature and wind speed (Figure 12a,b) were recorded at a nearby automated weather station, Margaret, 122 km southwest of RS09. Tides (Figure 12c) were realized from the CATS2008 model (Padman et al., 2002) at station RS10, which is on floating ice and approximates the tidal signal in the basin between Roosevelt Island and the Shirase Coast. Seismicity for class 1 (Figure 12d) dominates the detections at RS09 and is active throughout the year, comprising 52.8% of the detections. Classes 3, 4, 6, and 8 (Figure 12f,g,i,k) are also active throughout the year. Classes 5 and 7 (Figure 12h,j) are comparatively sparse, with seismicity limited to what appear to be discrete signals that could be associated with large fracture or crevasse events. No class 2 (Figure 12e) signals were recorded at RS09, even though elevated class 5 seismicity occurred during the 2016 winter.

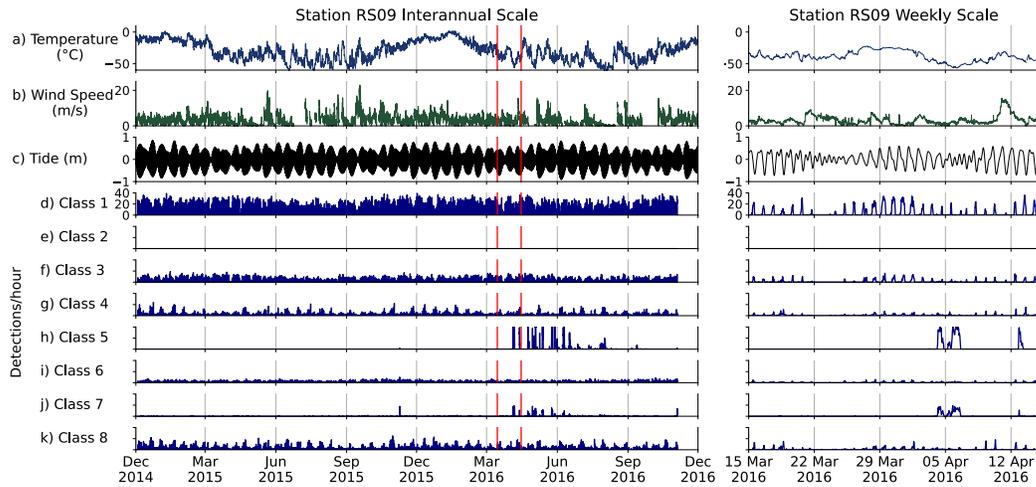


Figure 12. Two years of (a) temperature and (b) wind speed observations at Margaret automated weather station (MGT, approximately 122 km southwest of RS09, Figure 1), (c) model-derived tides calculated at station RS10, and (d-k) icequake detection statistics for each signal class. Interannual timescale is shown at left with vertical red lines indicating the subset weekly time-scale at right. The diurnal tidal signal correlates with seismicity for classes 2, 3, and 6. Tidal model from (Padman et al., 2002); weather station data from AMRC, SSEC, UW–Madison.

734 Of particular interest at station RS09 is evidence of seismicity associated with the
735 diurnal tide (Figure 12). On an interannual timescale, classes 4 and 8 exhibit a periodic
736 modulation of seismicity which tends to correlate with spring tides. Variability over fort-
737 night tidal cycles is shown between 15 March 2016 and 15 April 2016. This weekly timescale
738 shows that classes 1 and 3 correlate with diurnal tides. Even some relatively non-active
739 classes (4, 6, and 8) show signs of diurnal seismicity. These results are consistent with
740 a previous study that found more than 95% of detections at RS09 were from tidally in-
741 duced swarms of icequakes that occur throughout the year (Cole, 2020). The weekly timescale
742 also reveals the sudden onset and termination of winter seismicity in classes 5 and 7, sug-
743 gesting association with discrete ice shelf events such as crevasse expansion or major ice
744 fracture. This onset is consistent with the substantial increase in seismicity detected across
745 the RIS array visible in Figure 10 beginning in March 2016.

746 Other stations located at grounding zones exhibit similar patterns of seismicity, though
747 to a lesser extent than RS09. Station RS11, located east of RS09 on the Shirase Coast,
748 exhibits patterns of seismicity similar to RS09. These similarities indicate that ice shelf
749 seismicity at grounding zones is associated with similar ice shelf processes. RS08, on the
750 western flank of Roosevelt Island, and RS17, at Steershead Ice Rise, also exhibit diur-
751 nal seismicity, suggesting a dynamic diurnal process common to the grounding zones.
752 These patterns of seismicity indicate that the interaction of the ice shelf with the solid
753 earth at grounding zones is modulated by tides. Among the four stations at grounding
754 zones, classes 1, 4, and 8 are the most common signals, with class 8 signals occurring most
755 frequently at these stations. With a mean peak frequency of 6.3 Hz and a mean ampli-
756 tude of 210 nm/s², class 8 signals are among the strongest detected across the array.

757 **7 Conclusions**

758 Deep clustering of the Ross Ice Shelf (RIS) seismic array data set using a Gaus-
759 sian mixture model identified eight classes of impulsive signals, with linkage of at least
760 two of the classes to tidal variability near grounding zones. Additionally, compared to
761 2015, stations near the RIS front showed increased icequake activity during the 2016 El
762 Niño austral summer. A sudden increase in seismicity was also observed across the ar-
763 ray during the transition to the 2016 austral winter. The highest seismicity was observed
764 at grounding zones, particularly along the eastern flank of Roosevelt Island.

765 Deep clustering is an effective way to explore large seismic data sets, particularly
766 in its ability to identify dominant types of seismicity. The results provided by deep clus-
767 tering, when contextualized with non-seismic environmental data, can assist in the iden-
768 tification or correlation of seismic source mechanisms, as demonstrated with the RIS en-
769 vironmental data. Additionally, deep clustering can be readily tailored to investigate dif-
770 ferent aspects of the same or new data sets. Combined with its effectiveness at cluster-
771 ing seismic detections, this flexibility suggests that deep clustering can be incorporated
772 into existing seismic workflows to speed up exploratory data analysis.

773 As seismic data sets grow ever larger, novel machine learning techniques will be nec-
774 essary to enable researchers to fully utilize this data. Deep clustering has the potential
775 to become an important tool for exploring these large data sets, and to complement other
776 machine learning-based tools as well as conventional signal processing approaches. The
777 incorporation of such tools will enable more thorough and timely geophysical data anal-
778 ysis, thus improving the response of geophysical research to the needs of society in a rapidly
779 changing earth.

780 **Acknowledgments**

781 This work was funded by the Office of Naval Research through the National Defense Sci-
782 ence and Engineering Graduate Fellowship Program, and by National Science Founda-
783 tion (NSF) grant PLR 1246151, with support for Bromirski by NSF 1744856. Seismic
784 data from network XH (D. Wiens & Bromirski, 2014) were downloaded through IRIS
785 Web Services (<https://service.iris.edu/irisws/>). Seismic data were processed us-
786 ing Obspy software (Beyreuther et al., 2010). Figures were generated in MATLAB ([https://](https://www.mathworks.com)
787 www.mathworks.com) and with Matplotlib (<https://matplotlib.org>). The DEC model
788 was produced using PyTorch (<https://pytorch.org>). Antarctica elevation data, ground-
789 ing line, and coast line were obtained from Bedmachine (Morlighem et al., 2017) and plot-
790 ted using Antarctic Mapping Tools for MATLAB (Greene et al., 2017). Surface temper-
791 atures were obtained from AMRC, SSEC, University of Wisconsin–Madison ([https://](https://amrc.ssec.wisc.edu)
792 amrc.ssec.wisc.edu). Tide data were generated by the CATS2008 model (Padman et
793 al., 2002). Ross Sea ice coverage was obtained from NASA NSIDC (Cavalieri et al., 1996,
794 updated yearly). Code for this workflow is available at [https://github.com/NeptuneProjects/](https://github.com/NeptuneProjects/RISClusterPT)
795 [RISClusterPT](https://github.com/NeptuneProjects/RISClusterPT).

796 **References**

- 797 Aggarwal, C. C., Hinneburg, A., & Keim, D. A. (2001). On the Surprising Behavior
 798 of Distance Metrics in High Dimensional Space. In G. Goos, J. Hartmanis,
 799 J. van Leeuwen, J. Van den Bussche, & V. Vianu (Eds.), *Database Theory*
 800 — *ICDT 2001* (Vol. 1973, pp. 420–434). Berlin, Heidelberg: Springer Berlin
 801 Heidelberg. doi: 10.1007/3-540-44503-X_27
- 802 Aggarwal, C. C., & Reddy, C. K. (Eds.). (2014). *Data clustering: Algorithms and*
 803 *applications*. Boca Raton: Chapman and Hall/CRC.
- 804 Allen, R. (1982, December). Automatic phase pickers: Their present use and fu-
 805 ture prospects. *Bulletin of the Seismological Society of America*, 72(6B), S225-
 806 S242. doi: 10.1785/BSSA07206B0225
- 807 Aster, R. C., & Winberry, J. P. (2017, December). Glacial seismology. *Reports on*
 808 *Progress in Physics*, 80(12), 126801. doi: 10.1088/1361-6633/aa8473
- 809 Baker, M. G., Aster, R. C., Anthony, R. E., Chaput, J., Wiens, D. A., Nyblade, A.,
 810 ... Stephen, R. A. (2019, December). Seasonal and spatial variations in the
 811 ocean-coupled ambient wavefield of the Ross Ice Shelf. *Journal of Glaciology*,
 812 65(254), 912–925. doi: 10.1017/jog.2019.64
- 813 Baker, M. G., Aster, R. C., Wiens, D. A., Nyblade, A., Bromirski, P. D., Ger-
 814 stoft, P., & Stephen, R. A. (2020, October). Teleseismic earthquake wave-
 815 fields observed on the Ross Ice Shelf. *Journal of Glaciology*, 1–17. doi:
 816 10.1017/jog.2020.83
- 817 Barcheck, C. G., Tulaczyk, S., Schwartz, S. Y., Walter, J. I., & Winberry, J. P.
 818 (2018, March). Implications of basal micro-earthquakes and tremor for ice
 819 stream mechanics: Stick-slip basal sliding and till erosion. *Earth and Planetary*
 820 *Science Letters*, 486, 54–60. doi: 10.1016/j.epsl.2017.12.046
- 821 Beaucé, E., Frank, W. B., & Romanenko, A. (2018, January). Fast Matched Fil-
 822 ter (FMF): An Efficient Seismic Matched-Filter Search for Both CPU and
 823 GPU Architectures. *Seismological Research Letters*, 89(1), 165–172. doi:
 824 10.1785/0220170181
- 825 Bellman, R. E. (1961). *Adaptive Control Processes: A Guided Tour*. Rand Corpora-
 826 tion.
- 827 Bergen, K. J., & Beroza, G. C. (2018, June). Detecting earthquakes over a seismic
 828 network using single-station similarity measures. *Geophysical Journal Interna-*

- 829 *tional*, 213(3), 1984–1998. doi: 10.1093/gji/ggy100
- 830 Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann,
831 J. (2010, May). ObsPy: A Python Toolbox for Seismology. *Seismological*
832 *Research Letters*, 81(3), 530–533. doi: 10.1785/gssrl.81.3.530
- 833 Bianco, M. J., & Gerstoft, P. (2018, December). Travel Time Tomography With
834 Adaptive Dictionaries. *IEEE Transactions on Computational Imaging*, 4(4),
835 499–511. doi: 10.1109/TCI.2018.2862644
- 836 Bianco, M. J., Gerstoft, P., Olsen, K. B., & Lin, F.-C. (2019, December). High-
837 resolution seismic tomography of Long Beach, CA using machine learning. *Sci-*
838 *entific Reports*, 9(1), 14987. doi: 10.1038/s41598-019-50381-z
- 839 Bindschadler, R. A., King, M. A., Alley, R. B., Anandakrishnan, S., & Padman, L.
840 (2003, August). Tidally Controlled Stick-Slip Discharge of a West Antarctic
841 Ice Stream. *Science*, 301(5636), 1087–1089. doi: 10.1126/science.1087231
- 842 Bindschadler, R. A., Vornberger, P. L., King, M. A., & Padman, L. (2003). Tidally
843 driven stick–slip motion in the mouth of Whillans Ice Stream, Antarctica. *An-*
844 *nals of Glaciology*, 36, 263–272. doi: 10.3189/172756403781816284
- 845 Bishop, C. (2006). *Pattern Recognition and Machine Learning* (First ed.). Springer-
846 Verlag New York.
- 847 Boubekki, A., Kampffmeyer, M., Brefeld, U., & Jenssen, R. (2021, July). Joint op-
848 timization of an autoencoder for clustering and embedding. *Machine Learning*,
849 110(7), 1901–1937. doi: 10.1007/s10994-021-06015-5
- 850 Bromirski, P. D., Chen, Z., Stephen, R. A., Gerstoft, P., Arcas, D., Diez, A., . . . Ny-
851 blade, A. (2017, July). Tsunami and infragravity waves impacting A ntarctic
852 ice shelves. *Journal of Geophysical Research: Oceans*, 122(7), 5786–5801. doi:
853 10.1002/2017JC012913
- 854 Bromirski, P. D., Diez, A., Gerstoft, P., Stephen, R. A., Bolmer, T., Wiens, D. A.,
855 . . . Nyblade, A. (2015, September). Ross ice shelf vibrations. *Geophysical*
856 *Research Letters*, 42(18), 7589–7597. doi: 10.1002/2015GL065284
- 857 Bromirski, P. D., & Stephen, R. A. (2012). Response of the Ross Ice Shelf, Antarc-
858 tica, to ocean gravity-wave forcing. *Annals of Glaciology*, 53(60), 163–172. doi:
859 10.3189/2012AoG60A058
- 860 Cavalieri, D. J., Parkinson, C. L., Gloersen, P., & Zwally, H. J. (1996, updated
861 yearly). *Sea Ice Concentrations from Nimbus-7 SMMR and DMSP SSM/I-*

- 862 *SSMIS Passive Microwave Data, Version 1.* Boulder, Colorado USA. NASA
 863 National Snow and Ice Data Center Distributed Active Archive Center.
- 864 Chamarczuk, M., Nishitsuji, Y., Malinowski, M., & Draganov, D. (2020, January).
 865 Unsupervised Learning Used in Automatic Detection and Classification of
 866 Ambient-Noise Recordings from a Large-N Array. *Seismological Research*
 867 *Letters*, *91*(1), 370–389. doi: 10.1785/0220190063
- 868 Chamberlain, C. J., Hopp, C. J., Boese, C. M., Warren-Smith, E., Chambers, D.,
 869 Chu, S. X., . . . Townend, J. (2018, January). EQcorrscan: Repeating and
 870 Near-Repeating Earthquake Detection and Analysis in Python. *Seismological*
 871 *Research Letters*, *89*(1), 173–181. doi: 10.1785/0220170151
- 872 Chaput, J., Aster, R. C., McGrath, D., Baker, M., Anthony, R. E., Gerstoft, P., . . .
 873 Stevens, L. A. (2018, October). Near-Surface Environmentally Forced Changes
 874 in the Ross Ice Shelf Observed With Ambient Seismic Noise. *Geophysical*
 875 *Research Letters*, *45*(20). doi: 10.1029/2018GL079665
- 876 Chazan, S. E., Gannot, S., & Goldberger, J. (2019, March). Deep Clustering Based
 877 on a Mixture of Autoencoders. *arXiv:1812.06535 [cs, stat]*.
- 878 Chen, Z., Bromirski, P. D., Gerstoft, P., Stephen, R. A., Lee, W. S., Yun, S., . . .
 879 Nyblade, A. A. (2019, August). Ross Ice Shelf Icequakes Associated With
 880 Ocean Gravity Wave Activity. *Geophysical Research Letters*, *46*(15), 8893–
 881 8902. doi: 10.1029/2019GL084123
- 882 Chen, Z., Bromirski, P. D., Gerstoft, P., Stephen, R. A., Wiens, D. A., Aster, R. C.,
 883 & Nyblade, A. A. (2018, October). Ocean-excited plate waves in the Ross and
 884 Pine Island Glacier ice shelves. *Journal of Glaciology*, *64*(247), 730–744. doi:
 885 10.1017/jog.2018.66
- 886 Cole, H. M. (2020). *Tidally Induced Seismicity at the Grounded Margins of the*
 887 *Ross Ice Shelf, Antarctica* (Master’s Thesis). Colorado State University, Fort
 888 Collins, Colorado.
- 889 De Angelis, H., & Skvarca, P. (2003, March). Glacier Surge After Ice Shelf Collapse.
 890 *Science*, *299*(5612), 1560–1562. doi: 10.1126/science.1077987
- 891 Diez, A., Bromirski, P., Gerstoft, P., Stephen, R., Anthony, R., Aster, R., . . . Wiens,
 892 D. (2016, May). Ice shelf structure derived from dispersion curve analysis
 893 of ambient seismic noise, Ross Ice Shelf, Antarctica. *Geophysical Journal*
 894 *International*, *205*(2), 785–795. doi: 10.1093/gji/ggw036

- 895 Dupont, T. K., & Alley, R. B. (2005). Assessment of the importance of ice-shelf but-
 896 tressing to ice-sheet flow. *Geophysical Research Letters*, *32*(4). doi: 10.1029/
 897 2004GL022024
- 898 Fürst, J. J., Durand, G., Gillet-Chaulet, F., Tavard, L., Rankl, M., Braun, M., &
 899 Gagliardini, O. (2016, May). The safety band of Antarctic ice shelves. *Nature*
 900 *Climate Change*, *6*(5), 479–482. doi: 10.1038/nclimate2912
- 901 Gibbons, S. J., & Ringdal, F. (2006, April). The detection of low magnitude seis-
 902 mic events using array-based waveform correlation. *Geophysical Journal Inter-*
 903 *national*, *165*(1), 149–166. doi: 10.1111/j.1365-246X.2006.02865.x
- 904 Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- 905 Greene, C. A., Gwyther, D. E., & Blankenship, D. D. (2017, July). Antarctic Map-
 906 ping Tools for Matlab. *Computers & Geosciences*, *104*, 151–157. doi: 10.1016/
 907 j.cageo.2016.08.003
- 908 Hartigan, J. A., & Wong, M. A. (1979). Algorithm AS 136: A K-Means Clustering
 909 Algorithm. *Applied Statistics*, *28*(1), 100. doi: 10.2307/2346830
- 910 Hell, M. C., Cornelle, B. D., Gille, S. T., Miller, A. J., & Bromirski, P. D. (2019,
 911 November). Identifying Ocean Swell Generation Events from Ross Ice Shelf
 912 Seismic Data. *Journal of Atmospheric and Oceanic Technology*, *36*(11), 2171–
 913 2189. doi: 10.1175/JTECH-D-19-0093.1
- 914 Hinton, G. E. (2006, July). Reducing the Dimensionality of Data with Neural Net-
 915 works. *Science*, *313*(5786), 504–507. doi: 10.1126/science.1127647
- 916 Holtzman, B. K., Paté, A., Paisley, J., Waldhauser, F., & Repetto, D. (2018,
 917 May). Machine learning reveals cyclic changes in seismic source spec-
 918 tra in Geysers geothermal field. *Science Advances*, *4*(5), eaao2929. doi:
 919 10.1126/sciadv.aao2929
- 920 Hotovec-Ellis, A. J., & Jeffries, C. (2016, April). *Near Real-time Detection, Cluster-*
 921 *ing, and Analysis of Repeating Earthquakes: Application to Mount St. Helens*
 922 *and Redoubt Volcanoes* [Invited]. Reno, NV, USA.
- 923 Johnson, C. W., Ben-Zion, Y., Meng, H., & Vernon, F. (2020, August). Identifying
 924 Different Classes of Seismic Noise Signals Using Unsupervised Learning. *Geo-*
 925 *physical Research Letters*, *47*(15). doi: 10.1029/2020GL088353
- 926 Johnson, C. W., Meng, H., Vernon, F., & Ben-Zion, Y. (2019, August). Characteris-
 927 tics of Ground Motion Generated by Wind Interaction With Trees, Structures,

- 928 and Other Surface Obstacles. *Journal of Geophysical Research: Solid Earth*,
 929 *124*(8), 8519–8539. doi: 10.1029/2018JB017151
- 930 Kingma, D. P., & Ba, J. (2017, January). Adam: A Method for Stochastic Opti-
 931 mization. *arXiv:1412.6980 [cs]*.
- 932 Klein, E., Mosbeux, C., Bromirski, P. D., Padman, L., Bock, Y., Springer, S. R., &
 933 Fricker, H. A. (2020, October). Annual cycle in flow of Ross Ice Shelf, Antarc-
 934 tica: Contribution of variable basal melting. *Journal of Glaciology*, *66*(259),
 935 861–875. doi: 10.1017/jog.2020.61
- 936 Kong, Q., Trugman, D. T., Ross, Z. E., Bianco, M. J., Meade, B. J., & Gerstoft, P.
 937 (2019, January). Machine Learning in Seismology: Turning Data into Insights.
 938 *Seismological Research Letters*, *90*(1), 3–14. doi: 10.1785/0220180259
- 939 Kullback, S., & Leibler, R. A. (1951, March). On Information and Sufficiency.
 940 *The Annals of Mathematical Statistics*, *22*(1), 79–86. doi: 10.1214/aoms/
 941 1177729694
- 942 LeCun, Y. A., Bottou, L., Orr, G. B., & Müller, K.-R. (2012). Efficient BackProp.
 943 In G. Montavon, G. B. Orr, & K.-R. Müller (Eds.), *Neural Networks: Tricks*
 944 *of the Trade: Second Edition* (pp. 9–48). Berlin, Heidelberg: Springer Berlin
 945 Heidelberg. doi: 10.1007/978-3-642-35289-8_3
- 946 MacAyeal, D. R., Banwell, A. F., Okal, E. A., Lin, J., Willis, I. C., Goodsell, B.,
 947 & MacDonald, G. J. (2019, September). Diurnal seismicity cycle linked to
 948 subsurface melting on an ice shelf. *Annals of Glaciology*, *60*(79), 137–157. doi:
 949 10.1017/aog.2018.29
- 950 MacQueen, J. (1967). Some methods for classification and analysis of multivariate
 951 observations. In *Proceedings of the fifth berkeley symposium on mathematical*
 952 *statistics and probability, volume 1: Statistics* (pp. 281–297). Berkeley, Calif.:
 953 University of California Press.
- 954 Min, E., Guo, X., Liu, Q., Zhang, G., Cui, J., & Long, J. (2018). A Survey of Clus-
 955 tering With Deep Learning: From the Perspective of Network Architecture.
 956 *IEEE Access*, *6*, 39501–39514. doi: 10.1109/ACCESS.2018.2855437
- 957 Morlighem, M., Williams, C. N., Rignot, E., An, L., Arndt, J. E., Bamber, J. L., . . .
 958 Zinglensen, K. B. (2017, November). BedMachine v3: Complete Bed Topog-
 959 raphy and Ocean Bathymetry Mapping of Greenland From Multibeam Echo
 960 Sounding Combined With Mass Conservation. *Geophysical Research Letters*,

- 961 44(21). doi: 10.1002/2017GL074954
- 962 Mousavi, S. M., Horton, S. P., Langston, C. A., & Samei, B. (2016, October).
 963 Seismic features and automatic discrimination of deep and shallow induced-
 964 microearthquakes using neural network and logistic regression. *Geophysical*
 965 *Journal International*, 207(1), 29–46. doi: 10.1093/gji/ggw258
- 966 Mousavi, S. M., Zhu, W., Ellsworth, W., & Beroza, G. (2019, November). Un-
 967 supervised Clustering of Seismic Signals Using Deep Convolutional Autoen-
 968 coders. *IEEE Geoscience and Remote Sensing Letters*, 16(11), 1693–1697. doi:
 969 10.1109/LGRS.2019.2909218
- 970 Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. Cambridge,
 971 MA: MIT Press.
- 972 Nicolas, J. P., Vogelmann, A. M., Scott, R. C., Wilson, A. B., Cadeddu, M. P.,
 973 Bromwich, D. H., ... Wille, J. D. (2017, August). January 2016 extensive
 974 summer melt in West Antarctica favoured by strong El Niño. *Nature Commu-*
 975 *nications*, 8(1), 15799. doi: 10.1038/ncomms15799
- 976 Olinger, S. D., Lipovsky, B. P., Wiens, D. A., Aster, R. C., Bromirski, P. D., Chen,
 977 Z., ... Stephen, R. A. (2019, June). Tidal and Thermal Stresses Drive Seismic-
 978 ity Along a Major Ross Ice Shelf Rift. *Geophysical Research Letters*, 46(12),
 979 6644–6652. doi: 10.1029/2019GL082842
- 980 Ozanich, E., Thode, A., Gerstoft, P., Freeman, L. A., & Freeman, S. (2021). Deep
 981 embedded clustering of coral reef bioacoustics. *J. Acoust. Soc. Am.*, 16.
- 982 Padman, L., Fricker, H. A., Coleman, R., Howard, S., & Erofeeva, L. (2002). A new
 983 tide model for the Antarctic ice shelves and seas. *Annals of Glaciology*, 34,
 984 247–254. doi: 10.3189/172756402781817752
- 985 Paolo, F. S., Fricker, H. A., & Padman, L. (2015). Volume loss from Antarctic ice
 986 shelves is accelerating. *Science*, 348(6232), 327–331. doi: 10.1126/science
 987 .aaa0940
- 988 Perol, T., Gharbi, M., & Denolle, M. (2018, February). Convolutional neural net-
 989 work for earthquake detection and location. *Science Advances*, 4(2), e1700578.
 990 doi: 10.1126/sciadv.1700578
- 991 Pritchard, H. D., Ligtenberg, S. R. M., Fricker, H. A., Vaughan, D. G., van den
 992 Broeke, M. R., & Padman, L. (2012, April). Antarctic ice-sheet loss
 993 driven by basal melting of ice shelves. *Nature*, 484(7395), 502–505. doi:

- 994 10.1038/nature10968
- 995 Reddy, T. A., Devi, K. R., & Gangashetty, S. V. (2012, March). Nonlinear principal
996 component analysis for seismic data compression. In *2012 1st International
997 Conference on Recent Advances in Information Technology (RAIT)* (pp. 927–
998 932). Dhanbad, India: IEEE. doi: 10.1109/RAIT.2012.6194558
- 999 Riahi, N., & Gerstoft, P. (2017, March). Using graph clustering to locate sources
1000 within a dense sensor array. *Signal Processing, 132*, 110–120. doi: 10.1016/j
1001 .sigpro.2016.10.001
- 1002 Rignot, E., Mouginot, J., Morlighem, M., Seroussi, H., & Scheuchl, B. (2014, May).
1003 Widespread, rapid grounding line retreat of Pine Island, Thwaites, Smith, and
1004 Kohler glaciers, West Antarctica, from 1992 to 2011. *Geophysical Research
1005 Letters, 41*(10), 3502–3509. doi: 10.1002/2014GL060140
- 1006 Rousseeuw, P. J. (1987, November). Silhouettes: A graphical aid to the interpre-
1007 tation and validation of cluster analysis. *Journal of Computational and Applied
1008 Mathematics, 20*, 53–65. doi: 10.1016/0377-0427(87)90125-7
- 1009 Scambos, T. A. (2004). Glacier acceleration and thinning after ice shelf collapse
1010 in the Larsen B embayment, Antarctica. *Geophysical Research Letters, 31*(18),
1011 L18402. doi: 10.1029/2004GL020670
- 1012 Seydoux, L., Balestrieri, R., Poli, P., de Hoop, M., Campillo, M., & Baraniuk, R.
1013 (2020, December). Clustering earthquake signals and background noises in
1014 continuous seismic data with unsupervised deep learning. *Nature Communica-
1015 tions, 11*(1), 3972. doi: 10.1038/s41467-020-17841-x
- 1016 Smith, B., Fricker, H. A., Gardner, A. S., Medley, B., Nilsson, J., Paolo, F. S., ...
1017 Zwally, H. J. (2020, June). Pervasive ice sheet mass loss reflects compet-
1018 ing ocean and atmosphere processes. *Science, 368*(6496), 1239–1242. doi:
1019 10.1126/science.aaz5845
- 1020 Snover, D., Johnson, C. W., Bianco, M. J., & Gerstoft, P. (2021, March). Deep
1021 Clustering to Identify Sources of Urban Seismic Noise in Long Beach,
1022 California. *Seismological Research Letters, 92*(2A), 1011–1022. doi:
1023 10.1785/0220200164
- 1024 Steinbach, M., Ertöz, L., & Kumar, V. (2004). The Challenges of Clustering High
1025 Dimensional Data. In L. T. Wille (Ed.), *New Directions in Statistical Physics:
1026 Econophysics, Bioinformatics, and Pattern Recognition* (pp. 273–309). Berlin,

- 1027 Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/978-3-662-08968-2_16
- 1028 Telesca, L., & Chelidze, T. (2018, November). Visibility Graph Analysis of Seismic-
1029 ity around Enguri High Arch Dam, Caucasus. *Bulletin of the Seismological So-*
1030 *ciety of America*, 108(5B), 3141–3147. doi: 10.1785/0120170370
- 1031 Thoma, M., Jenkins, A., Holland, D., & Jacobs, S. (2008, September). Mod-
1032 elling Circumpolar Deep Water intrusions on the Amundsen Sea continen-
1033 tal shelf, Antarctica. *Geophysical Research Letters*, 35(18), L18602. doi:
1034 10.1029/2008GL034939
- 1035 Tibshirani, R., Walther, G., & Hastie, T. (2001, May). Estimating the number
1036 of clusters in a data set via the gap statistic. *Journal of the Royal Statistical*
1037 *Society: Series B (Statistical Methodology)*, 63(2), 411–423. doi: 10.1111/1467
1038 -9868.00293
- 1039 Trugman, D. T., & Shearer, P. M. (2017, March). GrowClust: A Hierarchical Clus-
1040 tering Algorithm for Relative Earthquake Relocation, with Application to the
1041 Spanish Springs and Sheldon, Nevada, Earthquake Sequences. *Seismological*
1042 *Research Letters*, 88(2A), 379–391. doi: 10.1785/0220160188
- 1043 van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of*
1044 *Machine Learning Research*.
- 1045 Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., & Manzagol, P.-A. (2010, Decem-
1046 ber). Stacked Denoising Autoencoders: Learning Useful Representations in a
1047 Deep Network with a Local Denoising Criterion. *Journal of Machine Learning*
1048 *Research*, 38.
- 1049 Wallet, B. C., & Hardisty, R. (2019, August). Unsupervised seismic facies using
1050 Gaussian mixture models. *Interpretation*, 7(3), 19.
- 1051 Wiens, D., & Bromirski, P. (2014). *Collaborative Research: Dynamic Response of the*
1052 *Ross Ice Shelf to Wave-Induced Vibrations, and Collaborative Research: Mantle*
1053 *Structure and Dynamics of the Ross Sea from a Passive Seismic Deployment*
1054 *on the Ross Ice Shelf*. International Federation of Digital Seismograph Net-
1055 works.
- 1056 Wiens, D. A., Anandkrishnan, S., Winberry, J. P., & King, M. A. (2008, June). Si-
1057 multaneous teleseismic and geodetic observations of the stick–slip motion of an
1058 Antarctic ice stream. *Nature*, 453(7196), 770–774. doi: 10.1038/nature06990
- 1059 Xie, J., Girshick, R., & Farhadi, A. (2016). Unsupervised Deep Embedding for Clus-

- 1060 tering Analysis. *Proceedings of the 33rd international conference on machine*
1061 *learning*, 10.
- 1062 Yang, B., Fu, X., Sidiropoulos, N. D., & Hong, M. (2017, June). Towards
1063 K-means-friendly Spaces: Simultaneous Deep Learning and Clustering.
1064 *arXiv:1610.04794 [cs]*.
- 1065 Yoon, C. E., O'Reilly, O., Bergen, K. J., & Beroza, G. C. (2015, December). Earth-
1066 quake detection through computationally efficient similarity search. *Science*
1067 *Advances*, 1(11), e1501057. doi: 10.1126/sciadv.1501057