

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 embedded clustering (DEC) reveals what types of dominant signals are de-
8 tected, not just when or where they are detected.
- 9 • DEC can be adapted to various kinds of data sets, enabling rapid exploration of
10 “big data” in seismology.
- 11 • Paired with environmental data, deep embedded clustering could provide insights
12 into the causes of seismicity.

Corresponding author: W. F. Jenkins II, wjenkins@ucsd.edu

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 embedded clustering (DEC) was employed to efficiently investigate these signals. DEC automatically groups these 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 DEC algorithm uses spectrograms as input and encodes their salient features into a 9-feature representation. Encoding is performed with an autoencoder, a type of deep neural network that is trained iteratively and seeks to reconstruct the input spectrograms from the encoded representation. 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

Using a machine learning technique called deep embedded clustering (DEC), we demonstrate the ability to automatically identify different types of impulsive seismic signals. The DEC algorithm encodes spectrograms into simplified representations and separates the representations into distinct clusters of signal types. The DEC technique was applied to seismic data recorded on the Ross Ice Shelf, Antarctica from 2014 to 2017. In addition to knowing when and where signals are detected, DEC enables users to determine the signal characteristics. Paired with environmental data, DEC 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 et al., 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 et al., 2019; Johnson et al., 2019). In this study we present an extension of *clustering* (Mousavi et al., 2016; Snover et al., 2020), 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

76 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 an
95 unsupervised clustering methodology 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). Hi-

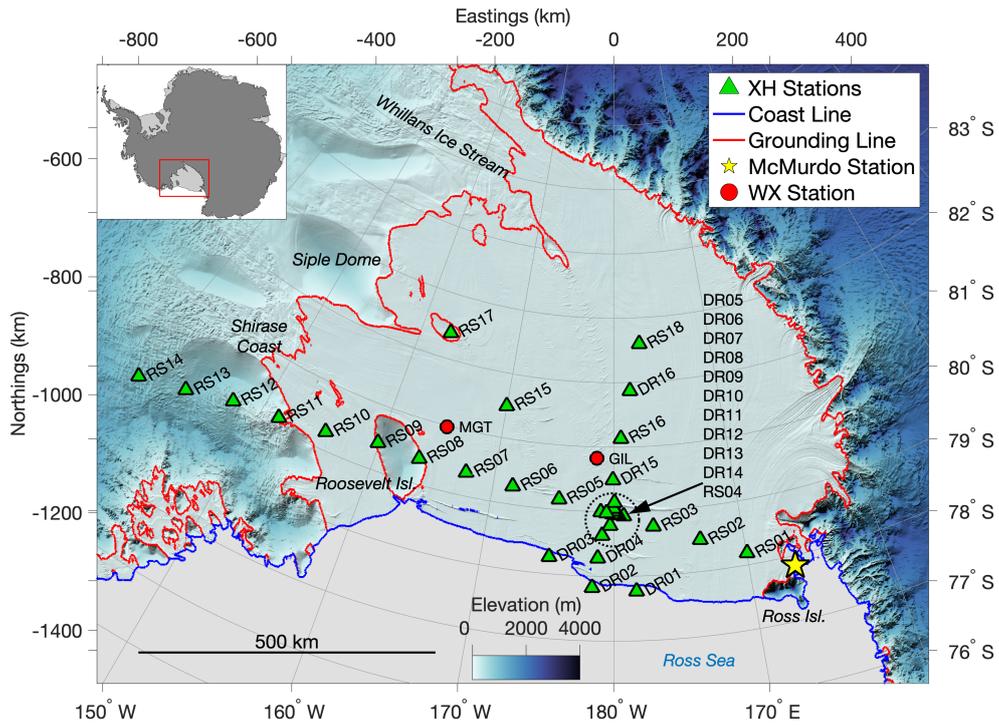


Figure 1. The passive broadband seismic array deployed from November 2014 to January 2017 consisted of 34 seismic stations and was deployed as part of the Ross Ice Shelf Dynamic Response to Wave-Induced Vibrations Project (Bromirski et al., 2015). RIS surface elevation, ice and water layer thicknesses,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

106 erarchical clustering has been used by Mousavi et al. (2016) to automatically discrim-
 107 inate between shallow and deep earthquakes, and by Trugman and Shearer (2017) to more
 108 precisely localize earthquakes. Graphical clustering has been used to localize sources in
 109 a dense seismic array by Riahi and Gerstoft (2017), and by Telesca and Chelidze (2018)
 110 to cluster seismic events in time. Distance-based clustering, like the popular k -means al-
 111 gorithm, (MacQueen, 1967; Hartigan & Wong, 1979) has been used by Chamarczuk et
 112 al. (2020) to cluster seismicity based on features extracted from seismic data. Perol et
 113 al. (2018) used k -means to define probabilistic earthquake locations as part of their con-
 114 volutional neural network (CNN) detection and localization technique. A novel approach
 115 was presented by Seydoux et al. (2020), who detect and cluster seismic signals and back-
 116 ground noise with the use of a deep scattering neural network and a Gaussian mixture
 117 model.

118 Not all clustering methods involve machine learning. Template matching, in which
 119 a matched filter is constructed from a template waveform, is used to scan through con-
 120 tinuous recordings to locate similar signals (Gibbons & Ringdal, 2006; Beaucé et al., 2018;
 121 Chamberlain et al., 2018). Yoon et al. (2015) and Bergen and Beroza (2018) present com-
 122 putationally efficient techniques in which seismic data are encoded using locally sensi-
 123 tive hashing, allowing similar signals to be identified within a hash table. Hotovec-Ellis
 124 and Jeffries (2016) developed an approach that uses correlation-based similarity search
 125 to automatically detect and cluster repeating volcanic seismicity in continuous data. Cole
 126 (2020) adopted the method of Hotovec-Ellis and Jeffries (2016) to cluster RIS array data
 127 at stations RS09, RS10, and RS11 in order to characterize tidal forcing of seismicity at
 128 these stations.

129 **2.2 Dimensionality**

130 Data are considered high-dimensional when many features are required to repre-
 131 sent or describe the data. Seismic data represented as time series, spectrograms, scalo-
 132 grams, or energy envelopes can contain thousands of features (e.g., discrete samples in
 133 a time series, or bins in a spectrogram). Clustering performed directly on such input data
 134 (Aggarwal & Reddy, 2014) is vulnerable to the “curse of dimensionality” (Bellman, 1961),
 135 i.e., as the dimensionality of the input data increases, the number of data points required
 136 to maintain sufficient sampling density increases exponentially. A further consideration
 137 is that clustering error metrics can give less meaningful results as dimensionality increases.

138 As high-dimensional data are difficult to cluster (Aggarwal et al., 2001; Steinbach
139 et al., 2004), dimensionality reduction remains a major focus of development (Yang et
140 al., 2017). To address the issue of dimensionality, it is desirable to transform the input
141 data to a lower-dimensional representation described by fewer, more salient features. Reddy
142 et al. (2012) use principal component analysis to compress seismic data to maximize fea-
143 ture variance, while Bianco and Gerstoft (2018) and Bianco et al. (2019) reduce dimen-
144 sionality using adaptive dictionary learning for high-resolution seismic tomography.

145 The approach to reducing dimensionality in this study employs an autoencoder,
146 a model whose output aims to reproduce its input via a series of non-linear transforma-
147 tions employing a deep neural network (DNN) (Hinton, 2006; Murphy, 2012; Yang et al.,
148 2017). These non-linear transformations provide greater capacity in dimension reduc-
149 tion, and can better model data with low-dimensional representations than, for exam-
150 ple, principal component analysis (PCA) (Goodfellow et al., 2016). The autoencoder first
151 encodes input data such as an image—in our case, a spectrogram—into a latent feature
152 vector. Next, the autoencoder decodes the latent features and reconstructs the original
153 image. Since the autoencoder provides a non-linear transformation of the data, it must
154 be trained using gradient descent. In this iterative training, the error between the in-
155 put and output is minimized. In doing so, the salient features of the data are learned
156 by the network weights. With the dimensionality of the input data reduced in the la-
157 tent feature space, clustering algorithms can be applied to the data’s latent feature space.

158 **2.3 Clustering in Reduced Dimensions**

159 A method that has shown improvement over traditional clustering techniques was
160 developed by Xie et al. (2016), whose *deep embedded clustering* (DEC) uses the latent
161 feature space as input to an adaptive clustering algorithm. DEC consists of two processes:
162 first, an autoencoder is trained to represent the data’s salient features; then the encod-
163 ing layers and clustering layer are jointly optimized. Yang et al. (2017) extend the ap-
164 proach in DEC by jointly optimizing the clustering step with training the entire autoen-
165 coder, not just the encoder layers.

166 Additional variations of DEC have been proposed: Xie et al. (2016) used a stacked
167 de-noising autoencoder (Vincent et al., 2010), and (Min et al., 2018) employed autoen-
168 coders composed of CNN layers and other architectures. More recently, Chazan et al.

169 (2019) developed a novel approach in which joint clustering is performed with a mixture
170 of autoencoders, each representing a cluster. Mousavi et al. (2019) used DEC to predict
171 whether seismic detections were local or teleseismic, and Snover et al. (2020) demonstrated
172 DEC’s ability to cluster anthropogenically generated seismic noise.

173 In this study, we implement DEC on RIS seismic data collected from December 2014
174 to November 2016, identifying several different classes of signals. Additionally, we demon-
175 strate the utility of DEC as an exploratory tool for large, real-world seismic data sets
176 by associating the clustering results with observed environmental factors.

177 **3 Ross Ice Shelf Seismic Array and Data**

178 Each station in the RIS seismic array consisted of 3-component Nanometrics Tril-
179 lium 120 PHQ seismometers emplaced 1 m below the surface of the ice, powered by so-
180 lar panels during the austral summers, and lithium-ion batteries during the austral win-
181 ters. Two subarrays comprised the array. The larger subarray consisted of 18 stations
182 spaced approximately 80 km apart (prefix RS), primarily oriented parallel to the RIS
183 front. The RS stations sampled short-period orthogonal components of ground veloc-
184 ity at a sampling rate of 100 Hz, except for two stations that sampled at 200 Hz. The
185 smaller subarray consisted of 16 stations (prefix DR) arranged approximately orthog-
186 onal to the icefront along the international date line, sampling ground velocity with a
187 sampling rate of 200 Hz. For this study, we were primarily interested in the detection
188 and classification of icequakes and local/regional earthquakes, using only vertical com-
189 ponent observations. Representative types of signals detected are shown in Figure 2.

190 Seismic data from each station were processed in 24-hour segments. Data were lin-
191 early de-trended, tapered with a Hann window, and decimated to a sampling rate of 50 Hz.
192 Instrument responses for all stations were removed, giving acceleration in m/s^2 . A band-
193 pass filter with cutoff frequencies at 3 and 20 Hz was applied to remove long-period sig-
194 nals originating from tides, tsunamis, infragravity waves, ocean swell, and teleseisms. An
195 event detection algorithm, the Z-detector (Swindell & Snell, 1977; Withers et al., 1998),
196 was used to detect impulsive signals, particularly icequakes and local earthquakes, with
197 a sliding window of 3 s. The detector was applied to data from each station between 2
198 December 2014 and 20 November 2016 for a total of 719 days of array data, yielding 427,798
199 detections.

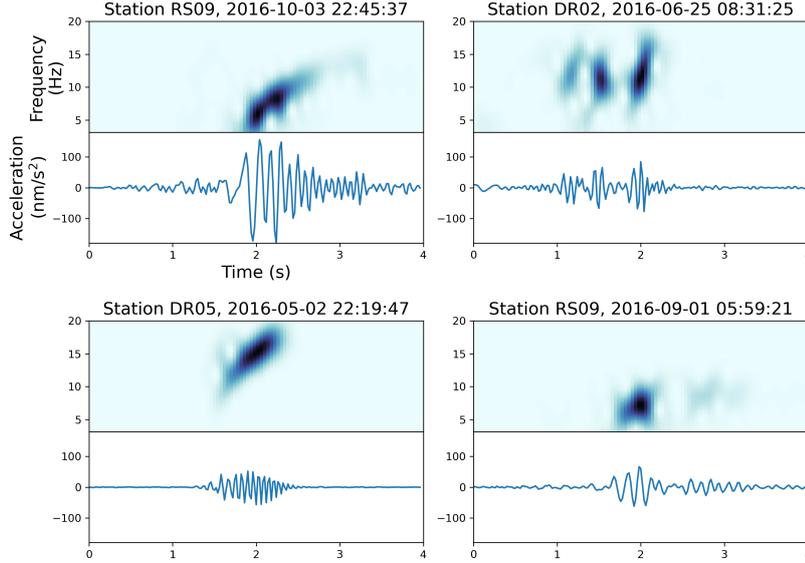


Figure 2. Seismic signals detected on the Ross Ice Shelf were diverse with variation in time, space, and source mechanism. Shown are examples of acceleration response seismograms and their respective normalized spectrograms spanning the 3-20 Hz band that were typical for the data set. The normalized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

200 Upon detection, a 4 s trace centered on the spectral peak of each triggered event
 201 was saved for processing. For each seismic trace saved, a spectrogram was computed us-
 202 ing the short-time Fourier transform with a 0.4 s Kaiser window, NFFT=256, and 90%
 203 overlap. Spectrograms contained one channel of amplitude information, 87 frequency bins,
 204 and 100 time bins for a total of 8,700 features per spectrogram. Finally, for each spec-
 205 trogram, the spectral levels were centered about the mean spectral value and normal-
 206 ized to the interval $[-1, 1]$.

207 4 Deep Embedded Clustering Model

208 The objective of the DEC model, shown in Figure 3, is to encode the input data—
 209 in this case, spectrograms of seismic signals—into a layer containing latent (lower-dimensional)
 210 features, called the *embedded* layer, to which a clustering algorithm is applied. The out-
 211 come of the clustering performance is then used to refine both the autoencoder model
 212 and the clustering layer in an effort to obtain more accurate latent space embeddings while
 213 improving clustering performance. In the implementation that follows, the 8,700 features

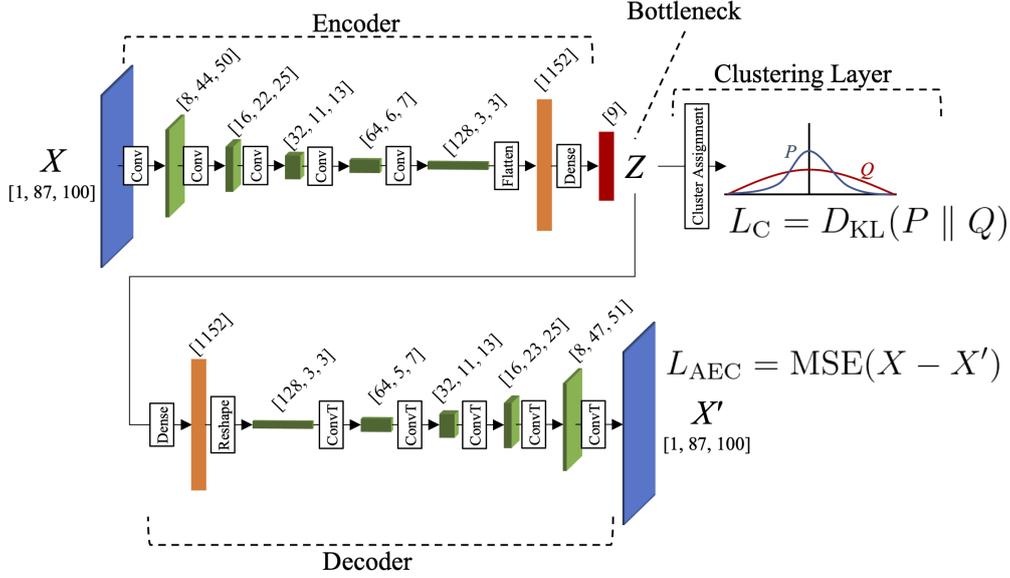


Figure 3. The deep embedded clustering model uses a convolutional autoencoder that encodes the data space X into the latent feature space Z , and a decoder that recovers the original input X from Z . The mean squared error (MSE) between the input X and the reconstruction X' is used as the autoencoder loss function. The latent feature space Z lies at the bottleneck between the encoder and decoder, providing the input to the clustering layer, which separately outputs a loss function. The two loss functions are combined and used to train the parameters that map $X \rightarrow Z \rightarrow X'$.

214 of an unencoded spectrogram are reduced to a latent feature space of just 9 embedded
 215 features with the use of a convolutional autoencoder, a type of DNN composed of con-
 216 volutional layers and their transposes.

217 4.1 Dimensionality Reduction with a Convolutional Autoencoder

218 Autoencoders provide a useful means of data approximation using a lower-dimensional
 219 representation via a sequence of non-linear transformations. The autoencoder model con-
 220 sists of three components: an *encoder*, a *bottleneck*, and a *decoder* (Murphy, 2012). First,
 221 the encoder maps input data from a data space X into a latent feature space Z , which
 222 is contained within the bottleneck of the model. Next, the decoder attempts to recon-
 223 struct X from Z . This process is performed iteratively with the objective of minimiz-
 224 ing the error between X and the decoder output, X' . In minimizing the error, the au-

225 toencoder learns the salient features of X and accurately encodes them in Z , thus re-
 226 ducing the dimensionality of the clustering task.

227 Consider a data set of spectrograms $\mathcal{D} = \{\mathbf{x}_i \in X^M\}_{i=1}^N$, where \mathbf{x}_i is a vector
 228 representation of the i^{th} spectrogram in a data set containing N spectrograms, and the
 229 number of features in \mathbf{x}_i , M , is the spectrogram size (the product of the number of fre-
 230 quency bins and time bins). In the encoder stage, the mapping of X to Z is described
 231 by $f_\theta : X \rightarrow Z$, where θ are parameters that are learned through iterative model train-
 232 ing. The decoder stage is a mirror operation of the encoder and seeks to map the latent
 233 feature space Z to the reconstruction X' by $g_\theta : Z \rightarrow X'$. The overall mapping of the
 234 autoencoder can be described as $F_\theta : X \rightarrow Z \rightarrow X'$, where $F_\theta = g_\theta \circ f_\theta$. Input spec-
 235 trograms \mathbf{x}_i map to their corresponding latent feature vectors by $\mathbf{z}_i = f_\theta(\mathbf{x}_i) \in Z^D$,
 236 where D is the number of embedded features, and to their reconstructions by $\mathbf{x}'_i = F_\theta(\mathbf{x}_i) \in$
 237 X' .

238 As the autoencoder is composed of convolutional layers and their transposes, F_θ
 239 is a nonlinear mapping that must be appropriately parameterized. This is accomplished
 240 by iteratively learning the parameters θ in order to minimize the error between the in-
 241 put and reconstructed data. The mean squared error (MSE) between an input spectro-
 242 gram with M features and its reconstruction, defined as

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

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

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

244 Performing this calculation over the entire data set at once is computationally expen-
 245 sive, memory intensive, and can lead to poor convergence. Instead, the loss is calculated
 246 in mini-batch subsets of the data space. For each mini-batch loss, the gradient of the loss
 247 function with respect to the weights comprising θ is computed, and the weights updated.
 248 When all mini-batches have been processed in this way, the next training epoch begins
 249 and the process is repeated. This mini-batch procedure, called stochastic gradient de-
 250 scent (Goodfellow et al., 2016), effectively adds noise to the gradient from sampling mini-
 251 batches, which ensures that the model does not converge on local minima. In parallel
 252 with autoencoder training, a subset of the data is used to validate the model’s perfor-
 253 mance after each training epoch. Training is performed until a specified maximum num-

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]	-	-	[1,152]	0
Encoded	Fully Connected	[1,152]	-	ReLU	[9]	10,377
FC	Fully Connected	[9]	-	ReLU	[1,152]	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

254 ber of epochs is reached, or if no improvement in the validation data’s MSE is observed
255 after ten epochs.

256 The design choice of autoencoder architecture can be informed by *a priori* knowl-
257 edge of a data set and its features, as well as practical considerations such as computa-
258 tional resources available. Our DNN architecture, detailed in Table 1, was designed to
259 be computationally efficient, simple to construct, and robust enough to learn salient fea-
260 tures from a noisy seismic data set. The model’s encoder is composed of five, two-dimensional
261 convolutional layers with rectified linear unit (ReLU) activation functions after each layer.
262 The kernel shape of each convolution is [3, 3], with a stride of two pixels. The number
263 of filters doubles at each layer. The output of the final convolution is flattened and lin-
264 early transformed with a dense layer into the encoded layer, which contains the embed-
265 ded latent feature space. For decoding, the inverse of the encoder operations are performed

Table 2. *Deep Embedded Clustering Model Parameters with Selected Hyper-parameters*

N	N_{train}	N_{val}	Learning Rate	Batch Size	Clusters	λ
427,798	100,000	25,000	0.001	1024	8	0.05

266 through linear transformation, reshaping, and transpose convolutions. In total, θ con-
 267 tains 218,250 trainable parameters under this architecture.

268 During training, the trainable parameters are optimized using the Adaptive Mo-
 269 ment Estimation (Adam) algorithm (Kingma & Ba, 2017), which propagates the loss af-
 270 ter each batch iteration backwards through the model. In this configuration, there are
 271 two principal hyper-parameters to address. First is the learning rate, which controls the
 272 step size used to step down the gradient of the loss. A learning rate which is too small
 273 will cause the model to converge extremely slowly or settle into a local minimum, while
 274 one that is too large may converge slowly or not at all. The second hyper-parameter is
 275 the batch size, which sets the number of spectrograms to be passed through the model
 276 at one time. The optimal configuration is found through hyper-parameter tuning, a pro-
 277 cess in which the autoencoder is trained with the various hyper-parameter permutations,
 278 with the hyper-parameters that yield the best results kept.

279 Autoencoder training was performed using spectrograms randomly selected with-
 280 out replacement. Of the selected spectrograms, 80% were used for training and 20% for
 281 validation. The number of spectrograms used and the optimal hyper-parameters are listed
 282 in Table 2. As seen in Figure 4a, training and validation losses fall off exponentially with
 283 each training epoch. To prevent the autoencoder from over-fitting, training is stopped
 284 early if the validation loss has not decreased in 10 epochs. The effectiveness of the au-
 285 toencoder’s ability to reconstruct the input spectrogram is illustrated in Figure 5. Be-
 286 cause the convolutional layers and their transposes consist of down-sampling and up-sampling
 287 operations, some loss of resolution in time and frequency is expected. Nevertheless, the
 288 structure of the spectrogram is largely preserved, with the salient information of the in-
 289 put encoded into the latent feature space.

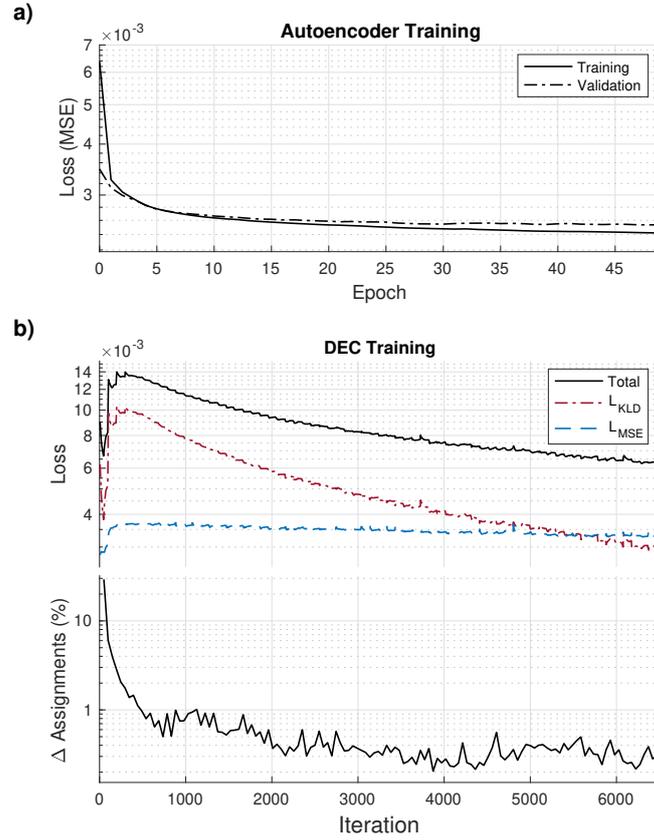


Figure 4. (a) Training and validation losses during autoencoder training. To avoid over-fitting the model, training was stopped at 49 epochs when the validation error began to increase. (b) In the upper plot, loss curves are shown for deep embedded clustering (DEC). In the lower plot, the percentage of samples which undergo class reassignment at each update interval is shown; training is stopped once the percent change is less than 0.2%

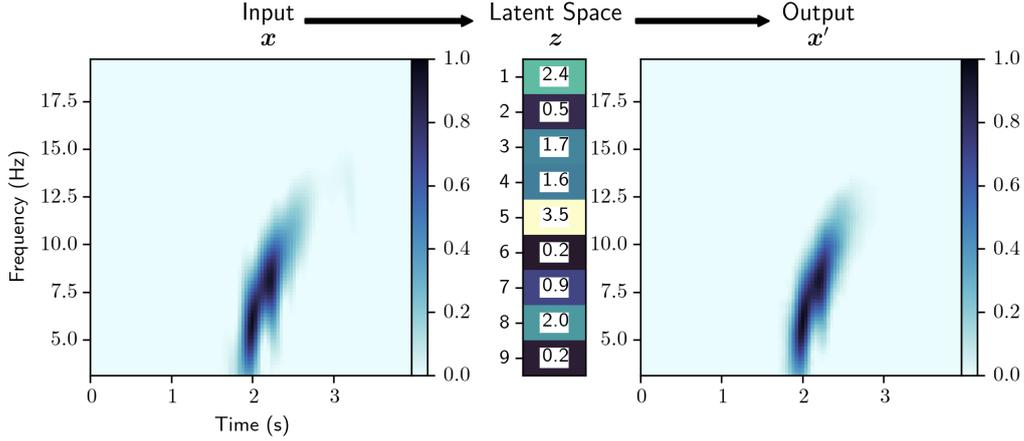


Figure 5. A trained autoencoder takes an input spectrogram \mathbf{x} , encodes it into an embedded, 9-dimensional latent feature vector \mathbf{z} , then reconstructs the input as \mathbf{x}' . The autoencoder preserves features correlated within a given cluster and discards the remaining signal. This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

290

4.2 Clustering Layer

291

292

293

294

295

In the DEC framework, clustering is performed in the latent feature space, Z , with the goal of finding 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 $\{\boldsymbol{\mu}_j \in Z\}_{j=1}^K$, i.e., centroids around which other similar signals may be found.

The Euclidean distance between a centroid and a latent feature vector is given by

$$d_{i,j} = \|\mathbf{z}_i - \boldsymbol{\mu}_j\|_2. \quad (3)$$

296

297

The distance metric $d_{i,j}$ is a measure of the similarity between features indexed by i and j , based on their Euclidean distance in the latent space.

298

299

300

301

Centroids are initialized with the k -means clustering algorithm (Hartigan & Wong, 1979), which randomly seeds the centroids and then iteratively updates their positions by minimizing the inertia, defined as the sum of the Euclidean distance between the centroids and the data points within each cluster:

$$\sum_{i=1}^n \min_{\boldsymbol{\mu}_j \in k} (d_{i,j}). \quad (4)$$

302

303

304

Because the k -means solution is guaranteed to converge on a local minimum, the algorithm is sensitive to the random seeding in its initialization. Two steps are taken to mitigate this. First, to improve the random seeding procedure and to speed up convergence,

305 an improved initialization scheme, k -means++ (Arthur & Vassilvitskii, 2007), is used
 306 instead of purely random seeding. Second, the k -means algorithm is repeated 100 times,
 307 and the solution that yields the minimum inertia is used to initialize the centroids.

308 With the centroids initialized by k -means, DEC seeks to further improve cluster-
 309 ing by using the difference between the embedded spectrograms and the cluster centroids
 310 as a loss function that provides feedback to the model parameters. Because the input
 311 data is unlabeled, a self-supervised method is required. We implement the method de-
 312 veloped by Xie et al. (2016), who, drawing from the t-Distributed Stochastic Neighbor
 313 Embedding (t-SNE) algorithm (van der Maaten & Hinton, 2008), propose measuring the
 314 difference between a t-Student’s distribution kernel of the embedded spectrograms and
 315 an auxiliary target distribution. A simplified Student’s t-distribution is used to measure
 316 the similarity between an embedded spectrogram, \mathbf{z}_i , and the cluster centroids $\boldsymbol{\mu}_j$:

$$q_{ij} = \frac{(1 + \|\mathbf{z}_i - \boldsymbol{\mu}_j\|^2)^{-1}}{\sum_j (1 + \|\mathbf{z}_i - \boldsymbol{\mu}_j\|^2)^{-1}}. \quad (5)$$

317 Equation (5) results in a set of soft class assignments, i.e., the probability that embed-
 318 ded spectrogram i will be assigned to class j , which are used to compute the auxiliary
 319 target distribution, p . The form of p is designed to improve clustering performance, em-
 320 phasize embeddings with high-confidence assignments, and normalize each cluster cen-
 321 troid’s contribution to the loss function so that large clusters minimally distort Z (Xie
 322 et al., 2016):

$$p_{ij} = \frac{q_{ij}^2 / \sum_i q_{ij}}{\sum_j (q_{ij}^2 / \sum_i q_{ij})}. \quad (6)$$

323 The dissimilarity between the distributions given by equations (5) and (6) is measured
 324 using the Kullback-Leibler divergence (Kullback & Leibler, 1951). From the divergence
 325 the clustering layer’s loss function is obtained:

$$L_C = D_{\text{KL}}(P \parallel Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}}. \quad (7)$$

326 In DEC, the clustering layer is attached to the trained autoencoder’s bottleneck,
 327 as shown in Figure 3. During training of the DEC model, the loss functions from equa-
 328 tions (1) and (7) are combined into a total loss function,

$$L = L_{\text{AEC}} + \lambda L_C, \quad (8)$$

329 where λ is a hyper-parameter that balances the contributions of the two losses, since they
 330 are of differing magnitudes. λ must be tuned: if it is too large, the clustering loss will

331 cause model instability and lead to distortion of the latent space, in which case the la-
 332 tent space will no longer represent the salient features of the data. If λ is too small, the
 333 effect on clustering performance will be minimal. We found that $\lambda = 0.05$ yielded op-
 334 timal performance for model training and clustering.

335 Two constituent processes occur simultaneously during DEC model training. First,
 336 the full loss from equation (8) is backpropagated through the DEC model parameters,
 337 which include the autoencoder as well as the cluster centroids. Second, to account for
 338 the cluster centroids changing as training progresses, the distributions q_{ij} and p_{ij} are up-
 339 dated at intervals. The update interval is a hyper-parameter that must be tuned. Through
 340 hyper-parameter tuning, an update interval of twice per training epoch was found to be
 341 optimal for clustering performance, minimizing DEC loss, and training within a reason-
 342 able time frame. Training is stopped after the number of samples changing assignments
 343 after every update interval reaches less than 0.2% of the total number of samples.

344 4.3 Selecting Optimal Number of Clusters

345 Determining the optimal number of clusters, K , is a major challenge in unsuper-
 346 vised machine learning. Although there are statistical methods available for choosing the
 347 optimal number of clusters (Rousseeuw, 1987; Tibshirani et al., 2001), in this study we
 348 treat K as a hyper-parameter, iterating the DEC workflow over a range of values for K
 349 and evaluating the results to choose the best value. Results are evaluated both quan-
 350 titatively and qualitatively. Quantitative evaluation is performed for each class by ex-
 351 amining cumulative distribution functions and probability density functions as functions
 352 of distance to each class centroid, $d_{i,j}$ (equation (3)). The qualitative approach is to vi-
 353 sually inspect the similarity of the latent feature vectors \mathbf{z}_i to their respective class cen-
 354 troids $\boldsymbol{\mu}_j$, and to see if the spectrograms and seismograms assigned to each class like-
 355 wise exhibit similarity. In general, the formation of two or more similar classes may in-
 356 dicate that too many classes were initialized, and the data can be grouped into a sin-
 357 gular class in post-processing. Too much variance among the spectrograms within a class
 358 may indicate the need for one or more additional classes. We found that $K = 8$ was
 359 the optimal number of classes for the full RIS data set.

5 Results

5.1 Clustering Performance

Clustering with DEC results in two distinct phases: first, the k -means algorithm sets the initial centroids, but the latent data are left unmodified. Second, during DEC, centroids are further refined while the latent data are moved much closer to their respective centroids, with some data reassigned to different classes altogether.

The performance of DEC is qualitatively checked by comparing centroids to their respective assigned latent data samples. Results for the training data set are shown in Figure 6. Each class j is represented by the columns in Figure 6, with each centroid $\boldsymbol{\mu}_j$ and its reconstruction $g_\theta(\boldsymbol{\mu}_j)$ 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}_i assigned to its class. Subsequent rows show the latent feature vectors \boldsymbol{z}_i , spectrograms \boldsymbol{x}_i , and associated seismograms of the data samples assigned to the respective classes.

For each class, latent feature vectors \boldsymbol{z}_i exhibit similar values to the class centroid $\boldsymbol{\mu}_j$, indicating that DEC has successfully grouped similar latent data samples into a class, and that the centroid is representative of the data in its class. The spectrograms in each class are likewise similar to each other and to the centroid reconstruction $g_\theta(\boldsymbol{\mu}_j)$, confirming that the latent features embedded in the centroids are representative of the spectrograms in the class. Finally, the similarity in the latent space and the time-frequency domain extends to the time domain, where seismograms in each class are similar to one another.

In addition to checking the efficacy of the clustering, visual examination of the results in Figure 6 permits an indication of whether or not an appropriate number of clusters was chosen. For example, classes 4 and 8 exhibit similar characteristics in time and frequency, distinct from each other primarily in peak frequency. If such distinctions are not useful or if similarities are redundant, classes can be combined in post-processing. If too few clusters are selected, classes may contain widely differing signals, indicating the need to increase the number of clusters.

To determine to what extent DEC further improves clustering over k -means clustering, t-SNE is used to visualize the 9-dimensional latent space in two dimensions (van

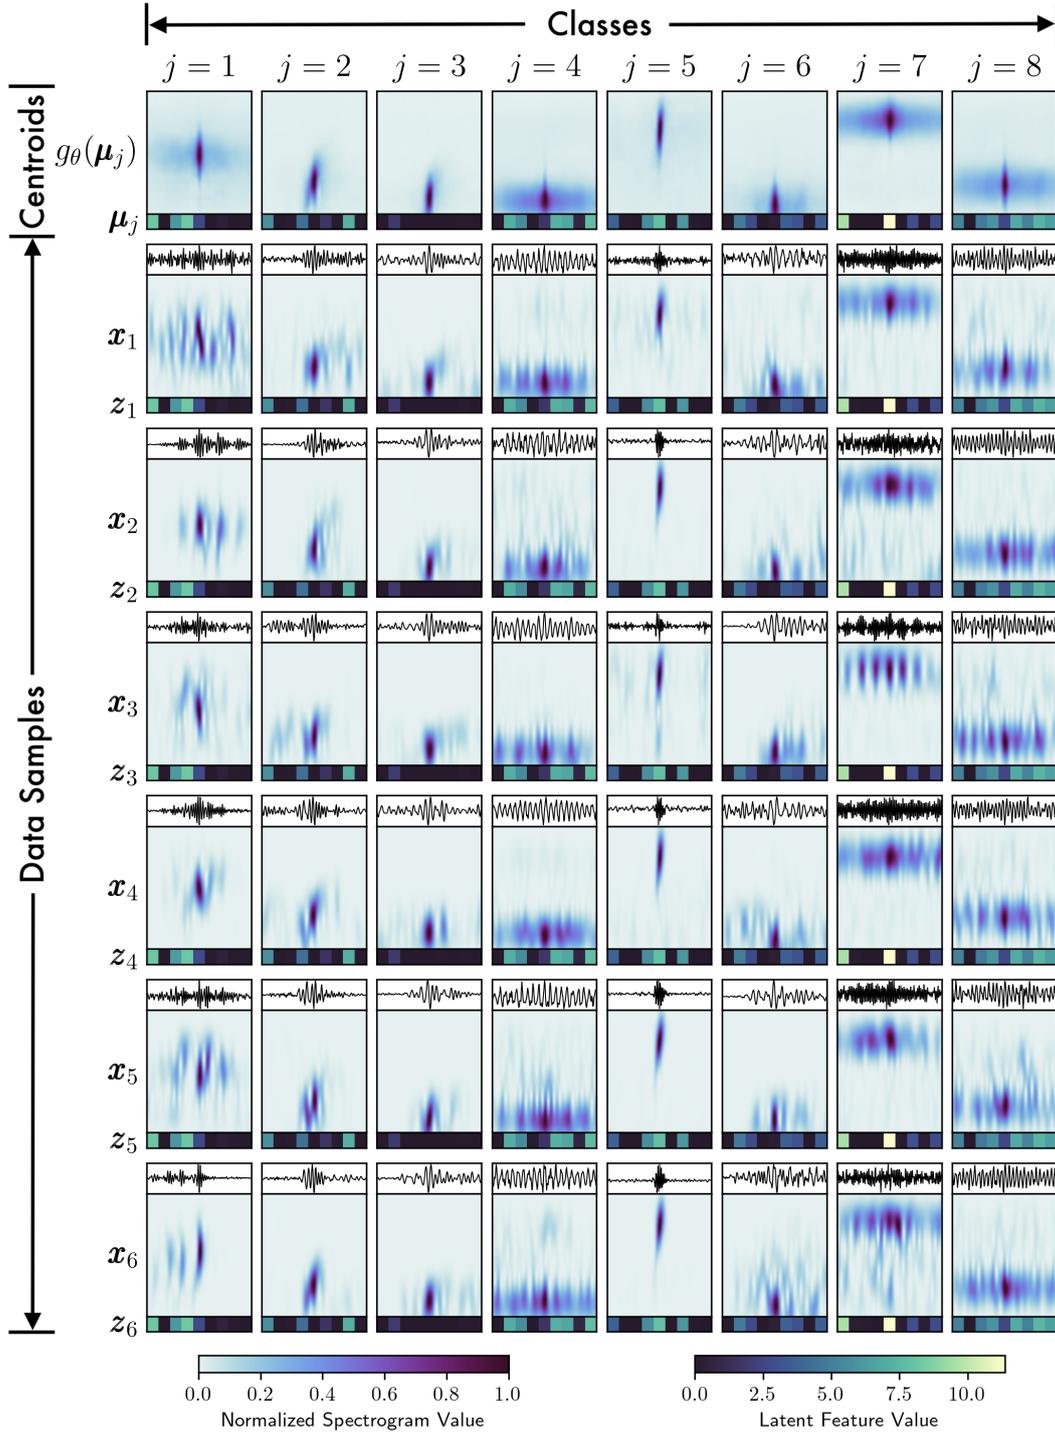


Figure 6. Within a given class j , the cluster centroids μ_j are similar to the latent feature data z_i . Though the centroids are not members of the data set, their reconstructions $g_{\theta}(\mu_j)$ exhibit similar characteristics to the spectrograms x_i assigned to each class. Seismograms plotted above each spectrogram also exhibit similarity within each class.

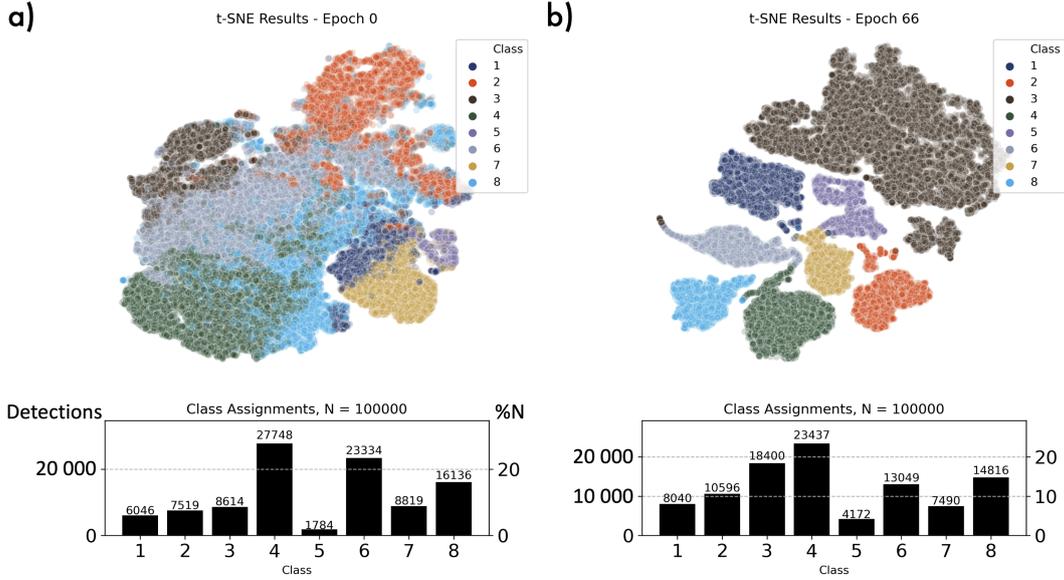


Figure 7. Visualization of the 9-dimensional latent data space is shown in two dimensions using the t-Student Stochastic Neighbor Embedding (t-SNE) plot. (a) Clustering and label assignment is performed with k -means clustering, with relatively poor separation within the data and overlapping classes. (b) After DEC, clusters are well separated and contain nearly homogeneous class members. Class histograms (a) before and (b) after DEC reveal the extent to which DEC reassigns latent data samples to different classes.

391 der Maaten & Hinton, 2008). t-SNE can illuminate possible clusters within data in
 392 an unsupervised manner by displaying data in geometrically separated clusters. In Figure 7,
 393 t-SNE results of the latent feature space clustered with k -means show that the data are
 394 largely contiguous with few exceptions. Applying the labels assigned by k -means to the
 395 data points shows that, while there is some geometric separation between the clusters,
 396 the embedding is characterized by overlapping and dispersed class members, indicating
 397 poor separation in the latent space and potentially incorrect assignment of samples to
 398 classes. Contrast this with Figure 7, in which t-SNE results at the conclusion of DEC
 399 show both geometric separation as well as homogeneous class assignments.

400 While t-SNE offers an intuitively visual way to look for clusters in data, results are
 401 sometimes difficult to interpret and are impossible to reproduce exactly due to the in-
 402herent randomness of the algorithm. Running t-SNE iteratively and with the same ran-
 403-dom seed can mitigate these limitations, but examination of the effects of DEC on the
 404-densities of the clusters provides a more concrete visualization. In Figure 8a, the cumu-

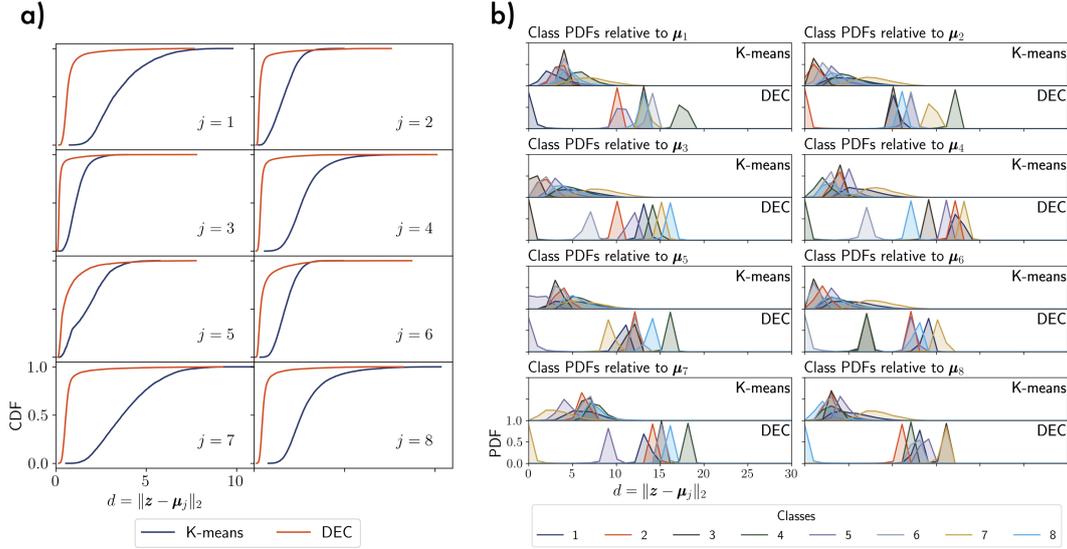


Figure 8. (a) Within each class, deep embedded clustering (DEC) reduces the mean distance of the assigned latent data to the centroid. The variance of the distance also decreases. As a result, the cumulative distribution functions shift to the left and have a steeper slope. (b) The effects of DEC are also 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 also increases the distance to the other class centroids and PDFs. The total effect is to separate the latent data samples of one class from the other classes, resulting in better clustering performance than k -means clustering.

405 relative distribution functions (CDF) for each class are shown as functions of distance to
 406 the centroid (equation (3)). For each class, the latent data move substantially closer to
 407 their assigned centroid, as evidenced by the decreased mean and variance of the CDF.
 408 Of interest to the ability for DEC to distinguish between clusters is the relation of each
 409 cluster to the others. In Figure 8b, the probability density functions (PDF) of all clusters
 410 are shown as functions of distance to each centroid. Before DEC, though k -means
 411 clustering results in the PDF of each class being closest to its centroid, there is significant
 412 overlap with other clusters, and the clusters themselves are not particularly dense.
 413 After DEC, the PDF of each class is closer to its centroid, denser, and farther removed
 414 from the other clusters. Thus, DEC effectively separates each cluster from the others,
 415 allowing for better distinction between clusters in the latent space.

416 The effects of DEC become readily apparent when the latent feature vectors are
 417 stacked and sorted according to their distance from each centroid, as shown in Figure 9.
 418 By sorting the latent space by sample index i such that $d_{i+1,j} > d_{i,j}$, cluster separa-
 419 tion can be visualized directly in the latent space. Before DEC, centroids are initialized
 420 with the k -means algorithm without modification to the latent data. Closest to each class
 421 centroid, the latent feature vectors are similar in appearance to the centroid, but tran-
 422 sition continuously to different patterns as the sorted index i increases. The contrast with
 423 the latent space after DEC is stark: because DEC moves latent data assigned to a par-
 424 ticular class closer to the centroid, the effect is that the latent feature vectors take on
 425 similar values, and therefore appearance, to the centroid. The result is that the latent
 426 space appears more sharply segmented after DEC, with the samples closest to the cen-
 427 troid of nearly uniform appearance to the centroid itself. For reference, the relative lo-
 428 cation of the other class centroids are marked with white vertical lines. Before DEC, the
 429 latent feature vectors belonging to the other classes are not readily apparent, whereas
 430 after DEC, most of the other centroid locations are associated with their distinctive la-
 431 tent feature vectors.

432 5.2 DEC Methodology Considerations

433 One of the key strengths of DEC is its employment of an autoencoder to reduce
 434 the dimensionality of the input data to obtain more effective clustering performance. By
 435 reducing the dimensionality of the data space, the complexity of the clustering problem
 436 is similarly decreased and the distance metrics gain relevance. The ability of the autoen-
 437 coder to quickly learn the salient features of the data and embed them into the latent
 438 space makes the technique adaptable to new data sets. While the autoencoder design
 439 choice for this study was sufficiently robust, autoencoder design presents opportunities
 440 for further experimentation and improvement. Design variables that could be altered in
 441 the DNN architecture include the number of layers, dimensions of the latent feature space,
 442 activation function types, incorporation of max-pooling and drop-out layers, and filter
 443 size, depth, and stride.

444 A second key strength of DEC is that clustering improvement and model optimiza-
 445 tion occur simultaneously. The outcome is denser, more separated clusters. This is a de-
 446 sirable effect in distance-based clustering, but it introduces a vulnerability: the success
 447 of the results may depend on the quality of the initial centroids. The k -means algorithm

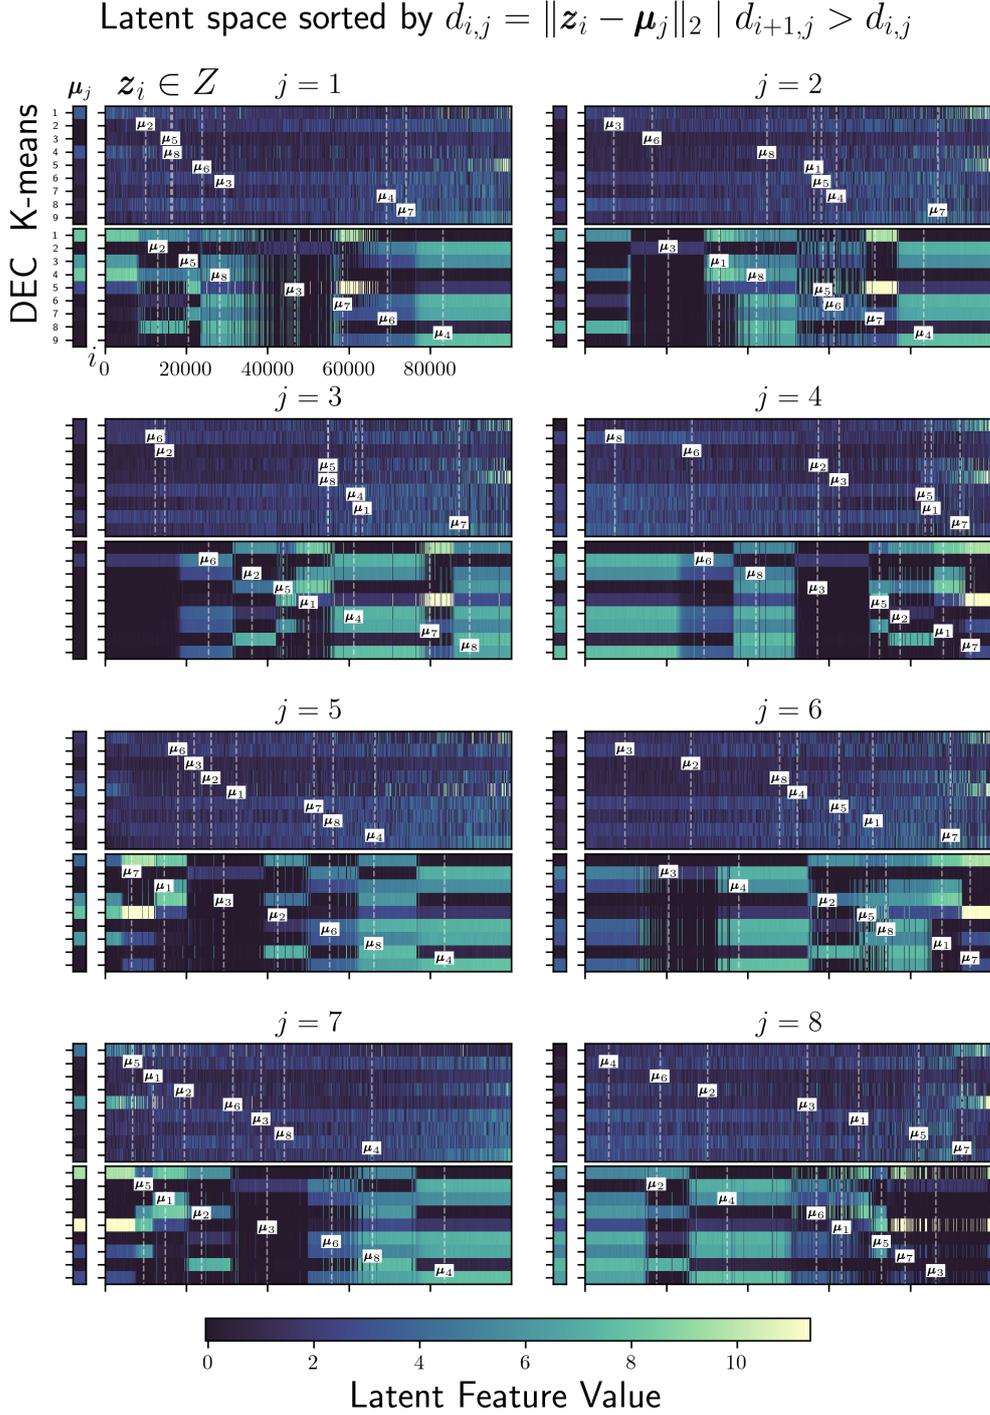


Figure 9. For each class j , latent data samples \mathbf{z}_i are shown stacked according to their distance $\|\mathbf{z}_i - \boldsymbol{\mu}_j\|$ from the centroid $\boldsymbol{\mu}_j$ (shown to the left). Distance of the other cluster centroids relative to the selected class j are indicated with vertical dotted lines. Deep embedded clustering (DEC) brings assigned data \mathbf{z}_i closer to the class centroid, resulting in homogeneity among the latent feature vectors assigned to that class.

448 assumes that all classes identified have approximately equal sample size and variance,
 449 yet we cannot expect naturally occurring seismic sources to exhibit such statistics, par-
 450 ticularly given the highly diverse nature of the data in time and space. However, Xie et
 451 al. (2016) demonstrated that DEC is robust against imbalances in class sample sizes, in-
 452 cluding size disparities up to a factor of 10. Further refinement of the centroid initial-
 453 ization step could be accomplished with a clustering algorithm that takes into account
 454 differing sample sizes and variances; for example, Gaussian mixture models are capable
 455 of characterizing clusters that are elongated and of varying size. Alternatively, an algo-
 456 rithm like k -medians, which can use the “Manhattan” distance $d_{i,j} = \|z_i - \mu_j\|_1$ to
 457 optimize centroids, could be used to overcome the inherent challenges of clustering in
 458 high-dimensional space. However, we assess that further refinement of centroid initial-
 459 ization is not necessary to achieve a workflow capable of performing satisfactory data
 460 exploration.

461 The flexibility afforded by DEC extends not only to model design, but also to data
 462 pre- and post-processing. Whereas model design is largely concerned with *how* the salient
 463 features are learned, data pre-processing is concerned with *what* is supplied to the model.
 464 This information is dependent on the choice of signal processing parameters, particularly
 465 signal duration, filter cutoff frequencies, and seismic event detection algorithm. Addi-
 466 tionally, various data transforms commonly used to characterize seismic waveforms can
 467 be used as input to DEC (Mousavi et al., 2016). In our case, we used spectrograms, but
 468 other transforms, such as continuous wavelet transform scalograms, could just as eas-
 469 ily be used as input to the DEC model. In post-processing, redundant or similar results
 470 can be combined.

471 **6 Discussion: Glaciological Implications**

472 The full RIS array data set contains 427,798 seismic detections. A summary of the
 473 data set statistics and class characteristics (Table 3) shows the total number of detec-
 474 tions for each class, as well as the percentage of detections occurring in the austral sum-
 475 mers (January, February, and March) versus the austral winters (June, July and August).
 476 Classes 1, 4, and 7 have pronounced differences (more than 10%) between the number
 477 of detections occurring in the summers versus the winters, while differences for classes
 478 2, 3, 5, and 8 are less pronounced (between 5% and 10%). Class 6 appears to have lit-
 479 tle difference (less than 5%) between austral summers and winters. Inter-annual com-

Table 3. *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				Amplitude (accel., nm/s ²)							
	N	%N Summer (JFM)		%N Winter (JJA)		Mean Peak Freq (Hz)	Mean	Median	Standard Deviation	Maximum		
		Total	2015	2016	Total						2015	2016
1	34,919	19	9	10	31	10	21	10.5	53	5	400	26,780
2	45,079	31	15	16	23	11	12	7.2	150	46	992	69,410
3	78,861	32	17	15	24	14	11	5.3	187	38	6,706	1,632,000
4	100,009	19	10	9	31	17	13	4.8	13	6	129	23,650
5	18,268	18	7	11	25	6	18	14.8	409	15	5,684	461,200
6	55,633	30	17	13	27	16	11	4.4	54	8	315	25,166
7	32,276	11	7	4	43	27	16	16.1	7	4	22	2,709
8	62,753	23	10	13	29	10	19	6.9	18	4	334	41,923

480 parisons for each season show that classes 1, 2, 5, and 8 experienced an increase in ac-
481 tivity in the 2016 austral summer over the 2015 austral summer, with classes 5 and 8
482 exhibiting the largest changes. Classes 5 and 8 also increase by factors of three and two,
483 respectively, in the 2016 austral winter over the 2015 austral winter. These trends can
484 be investigated in more detail from Figure 10a, where the frequencies of detections shown
485 as a function of station and month exhibit spatiotemporal patterns that may reveal as-
486 sociations between environmental forcing and seismicity. Clustering enables these pat-
487 terns to be further explored by class and month (Figure 10b), and by class and station
488 (Figure 10c).

489 From Figure 10a, certain patterns are readily apparent, such as increased seismic
490 detections during the austral summer months at stations DR01, DR02, and DR03. These
491 three stations were located approximately 2 km from the ice front and detected seismic-
492 ity associated with ocean gravity waves impacting at the shelf front that cause fractur-
493 ing (icequakes) and calving (Chen et al., 2019). Furthermore, seismicity at these stations
494 during the 2016 austral summer is higher than the same period in 2015. The remain-
495 ing DR stations and stations RS01 through RS07 exhibit the opposite pattern: austral

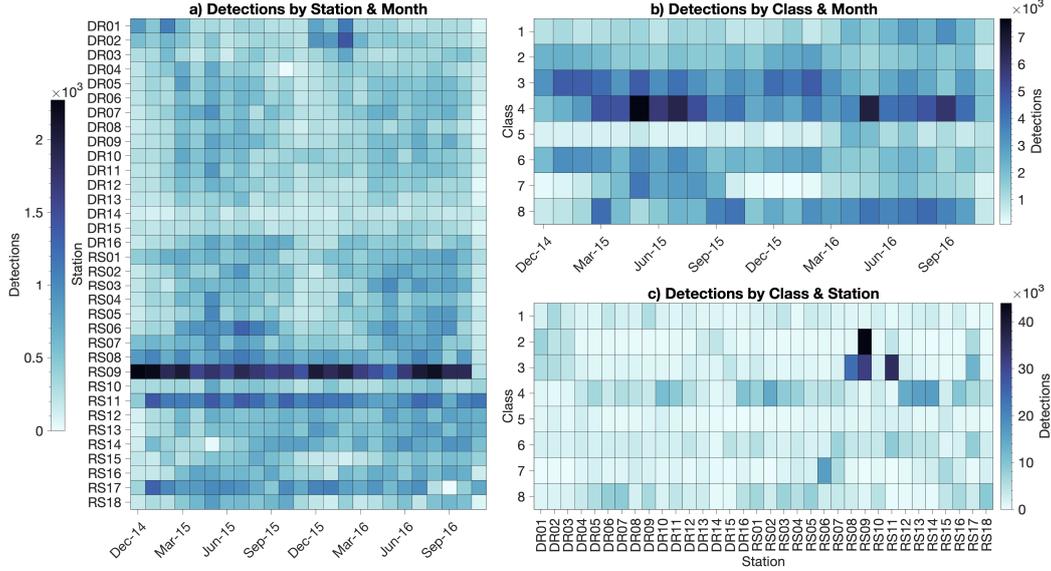


Figure 10. (a) The frequency of detections comprising the Ross Ice Shelf data set is shown by station and month. Deep embedded clustering (DEC) provides a further breakdown by (b) class and month for all stations, and (c) class and station.

496 summers are relatively quiet, with increased detection frequencies in the austral winters.
 497 The four most active stations were located near grounding zones: station RS09 (41,615
 498 detections) on the eastern flank of Roosevelt Island; station RS11 (25,884 detections) on
 499 the Shirase Coast; station RS08 (18,655 detections) on the western flank of Roosevelt
 500 Island; and station RS17 (18,653 detections) on Steershead Ice Rise. All of these stations
 501 exhibited persistent seismicity throughout the two deployment years, with the exception
 502 of station RS17, which was offline for several weeks from August to September 2016.

503 Some classes of signal detections exhibit temporal patterns that are visible in Fig-
 504 ure 10b. Classes 1 and 5 have increased detection frequencies in the austral winter of 2016,
 505 while classes 2 and 3 have increased detections in the austral summers. Class 6 also ex-
 506 hibits increased detections in the austral summers, though detections in the austral win-
 507 ter of 2015 are also high. Meanwhile, classes 4, 7, and 8 exhibit low seismicity during the
 508 austral summers. A further dimension to the analysis is shown in Figure 10c, which shows
 509 the distribution of classes by station. The most frequently occurring class in the data
 510 set is class 4, which from Figure 6 resembles seismic tremor, and occurs across the ar-

511 ray with peak activity in the austral winters. Class 3 is the prominent signal type at sta-
 512 tions near grounding zones (RS08, RS09, RS11, and RS17), as is class 2 at RS09.

513 An important caveat for the detection statistics shown in Table 3 and Figure 10
 514 arises from the physics governing seismic propagation. For a given amplitude, low fre-
 515 quency seismic energy propagates farther than high frequency seismic energy. We thus
 516 expect the seismometers in the RIS array to detect low-frequency signals from farther
 517 away than high-frequency signals, skewing the total detections in favor of the lower-frequency
 518 detections. Factoring in signal amplitude also affects the range at which seismic energy
 519 is detected. From Table 3, class 7 has an average spectral peak at 16.1 Hz, the highest
 520 of the classes, with a total of 32,276 detections, the second lowest of the classes. Sim-
 521 ilarly, class 5 has the second-highest average spectral peak at 14.8 Hz, with the least amount
 522 of detections among the classes. These two classes are nevertheless distinct from each
 523 other in amplitude and waveform type: from Table 3, class 7 has a mean amplitude of
 524 7.0 nm/s^2 , while class 5 has a mean amplitude of 408.5 nm/s^2 ; and from Figure 6, class
 525 7 consists of continuous signals while class 5 consists of high-frequency impulses likely
 526 resulting from fracturing. Detection statistics are further affected by signal-to-noise ra-
 527 tios at the seismometers and by limitations of the automated seismic event detector, such
 528 as the inability to separate signals from different classes that are received simultaneously.

529 Though the sources of uncertainty in the detection statistics are nontrivial, with
 530 a proper understanding of these limitations and when paired with environmental data,
 531 the clustering results can nevertheless be used to analyze the association of potential seis-
 532 mic source mechanisms that may be related to ice shelf dynamics. In the following sec-
 533 tions, we provide vignettes using stations DR02 and RS09 to demonstrate the utility of
 534 DEC in exploring data and identifying potential causes of seismicity when examined in
 535 conjunction with environmental data.

536 **6.1 Seasonal seismicity at the RIS front**

537 Approximately 2 km from the RIS front on Nascent Iceberg, station DR02 exhibits
 538 a seasonal pattern of seismicity associated with changes in air temperature and sea ice
 539 concentration in the Ross Sea. During the austral summer, sea ice melts and the con-
 540 centration (Figure 11a) decreases from nearly 100% to approximately 25%, permitting
 541 ocean gravity waves to impinge directly on the ice shelf front. Additionally, warmer air

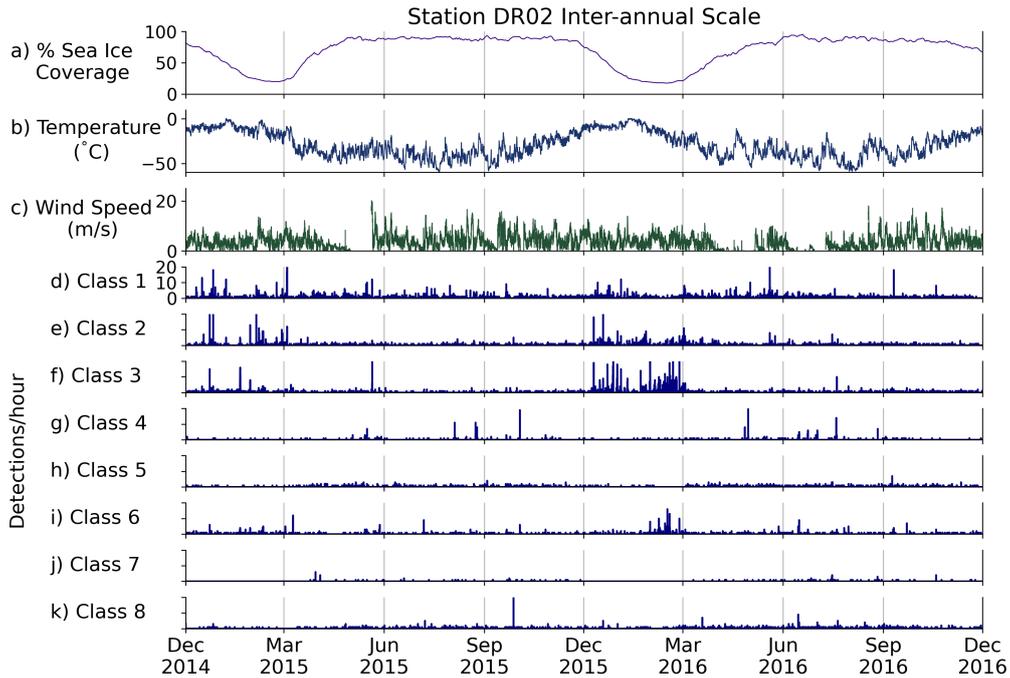


Figure 11. Two years of (a) sea ice coverage on the Ross Sea, (b) temperature and (c) wind speed at Gill automated weather station (approximately 223 km south of DR02), and (d-k) icequake detection statistics for each signal class. Classes 2, 3, and 6 exhibit increased seismicity during the austral summers. Sea ice concentration from NSIDC (Cavalieri et al., 1996, updated yearly); weather station data from AMRC, SSEC, UW–Madison.

542 temperatures (Figure 11b) accelerate calving and ice shelf structural failure at the front,
 543 processes which generate seismic activity (Chen et al., 2019). During the austral win-
 544 ter, sea ice coverage on the Ross Sea reaches nearly 100%, damping higher frequency ocean
 545 gravity waves such as swell.

546 Increased levels of seismicity are observed for classes 2, 3, and 6 at DR02 (Figure 11e,f,i)
 547 during the austral summers. Classes 3 and 6 are especially active during the 2016 aus-
 548 tral summer, when strong El Niño conditions led to anomalously persistent high tem-
 549 peratures across West Antarctica (Nicolas et al., 2017) and ocean-ice shelf interactions
 550 were likely enhanced. Patterns similar to the seismicity at DR02 were observed at sta-
 551 tions DR01 and DR03, also located near the RIS front, and can be seen in the total de-
 552 tections by station and month in Figure 10a. Widespread surface melt on the RIS was

553 observed between 10-21 January 2016 (Nicolas et al., 2017; Chaput et al., 2018), which
554 can lead to hydrofracture and contribute to ice shelf disintegration (Hubbard et al., 2016;
555 Alley et al., 2018).

556 Although class 1 has elevated activity during the summers, it maintains activity
557 throughout the winter months, suggesting that gravity wave activity is not the dominant
558 forcing. The persistence of class 1 signals, which often consist of impulse trains, suggests
559 they may be caused by icequakes resulting from the motion of the ice shelf itself, as the
560 ice flow velocity in the vicinity of station DR02 is among the highest observed on the
561 RIS (Klein et al., 2020). Classes 4, 5, 7, and 8 (Figure 11g,h,j,k) are more active dur-
562 ing the coldest periods of the year, suggesting that these signals may be associated with
563 extremely cold temperatures or strong wind events. Cold-weather enhanced seismicity
564 has been documented at a rift approximately 140 km south of the ice front Olinger et
565 al. (2019). Across all classes, discrete instances of high seismicity occur that do not cor-
566 respond to environmental forcing. Such instances may indicate the occurrence of frac-
567 turing ice (icequakes) or events associated with crevasse expansion.

568 **6.2 Diurnal seismicity on Roosevelt Island**

569 Station RS09 on the eastern flank of Roosevelt Island experienced the most detec-
570 tions on the array, comprising 9.7% of detections in the full data set. In Figure 12, po-
571 tential environmental sources of seismicity are compared to the seismicity of each class.
572 Temperature and wind speed (Figure 12a,b) were recorded at a nearby automated weather
573 station, Margaret, 122 km southwest of RS09. Tides (Figure 12c) were realized from the
574 CATS2008 model (Padman et al., 2002) and calculated at station RS10, which is on float-
575 ing ice and approximates the tidal signal in the basin between Roosevelt Island and the
576 Shirase Coast. Seismicity for classes 2, 3, and 6 (Figure 12d,e,i) dominate the detection
577 results and are active throughout the year, with classes 2 and 3 comprising 52.8% and
578 38.0% of the detections, respectively. Classes 1, 4, 5, 7, and 8 (Figure 12d,g,h,j,k) are
579 comparatively sparse, with seismicity limited to what appear to be discrete events that
580 could be associated with large fracture or crevasse events.

581 Of particular interest at station RS09 is evidence of seismicity corresponding with
582 the diurnal tide. On the inter-annual timescale, class 6 exhibits a periodic modulation
583 of seismicity which tends to correspond with spring tides. To examine the data on a finer

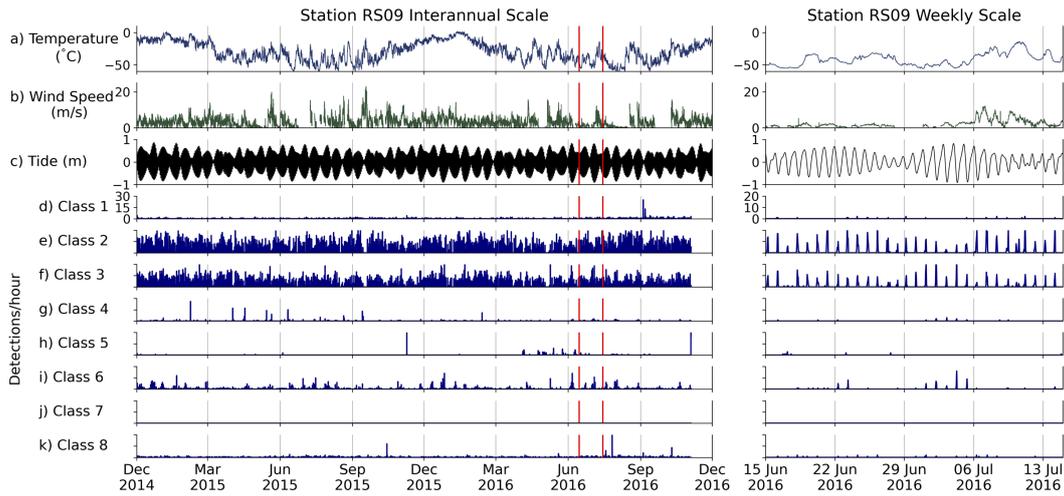


Figure 12. Two years of (a) temperature and (b) wind speed at Margaret automated weather station (MGT, approximately 122 km southwest of RS09, Figure 1), (c) model-derived tides calculated at station DR10, and (d-k) icequake detection statistics for each signal class. Inter-annual timescale is shown at left with vertical red lines indicating the subset weekly time-scale at right. The diurnal tidal signal corresponds to seismicity for classes 2, 3, and 6. Tidal model from (Padman et al., 2002); weather station data from AMRC, SSEC, UW–Madison.

584 timescale, data are shown between 15 June 2016 and 15 July 2016. This weekly timescale
 585 shows that classes 2 and 3, the dominant signal classes, correspond to diurnal tides. Even
 586 some relatively non-active classes (1, 4, and 8) show signs of diurnal seismicity. These
 587 results are consistent with a previous study that found more than 95% of detections at
 588 RS09 were from tidally induced swarms of icequakes that occur throughout the year (Cole,
 589 2020).

590 Other stations located in grounding zones exhibit similar patterns of seismicity, though
 591 to a lesser extent than RS09. Station RS11, located east of RS09 across an expanse of
 592 floating ice, exhibits similar patterns of seismicity to RS09. This suggests that ice shelf
 593 seismicity at Roosevelt Island and the Shirase Coast grounding zones are associated with
 594 similar ice shelf processes. RS08, on the western flank of the RIS, and RS17, at Steer-
 595 ahead Ice Rise, also exhibit diurnal seismicity, suggesting a dynamic diurnal process com-
 596 mon to the grounding zones and corresponding to the tidal cycle. Class 3 signals are the
 597 most common among the four stations in the grounding zone. With a mean peak fre-

598 quency of 5.3 Hz and a mean amplitude of 187 nm/s², these signals are among the strongest
599 detected on the array.

600 **7 Conclusions**

601 Deep embedded clustering (DEC) is an effective way to explore large seismic data
602 sets, particularly in its ability to identify dominant types of seismicity. The results pro-
603 vided by DEC, when contextualized with non-seismic environmental data, can assist in
604 the identification or correlation of seismic source mechanisms, as demonstrated with the
605 RIS environmental data. Additionally, DEC can be readily tailored to investigate dif-
606 ferent aspects of the same or new data sets. Combined with its effectiveness at cluster-
607 ing seismic detections, this flexibility suggests that DEC can be incorporated into ex-
608 isting seismic workflows in order to speed up exploratory data analysis.

609 Application of DEC to the Ross Ice Shelf array data set identified eight classes of
610 impulsive signals, with linkage of three of the classes to tidal variability near grounding
611 zones. Additionally, stations near the front showed increased icequake activity during
612 the 2016 El Niño austral summer. The highest seismicity was observed at grounding zones,
613 particularly along the eastern flank of Roosevelt Island.

614 As seismic data sets grow ever larger, novel machine learning techniques will be nec-
615 essary to enable researchers to fully utilize this data. DEC has the potential to become
616 an important tool for exploring these large data sets, and to complement other machine
617 learning-based tools as well as conventional signal processing approaches. The incorpo-
618 ration of such tools will enable more thorough and timely geophysical data analysis, thus
619 improving the response of geophysical research to the needs of society in a rapidly chang-
620 ing earth.

621 **Acknowledgments**

622 This work was funded by the Office of Naval Research through the National Defense Sci-
623 ence and Engineering Graduate Fellowship Program, and by National Science Founda-
624 tion grant PLR 1246151. Seismic data from network XH (D. Wiens & Bromirski, 2014)
625 were downloaded through IRIS Web Services (<https://service.iris.edu/irisws/>).
626 Seismic data were processed using Obspy software (Beyreuther et al., 2010). Figures were
627 generated in MATLAB (<https://www.mathworks.com>) and with Matplotlib (<https://>

628 matplotlib.org). The DEC model was produced using PyTorch (<https://pytorch.org>).
 629 Antarctica elevation data, grounding line, and coast line were obtained from Bedmachine
 630 (Morlighem et al., 2017) and plotted using Antarctic Mapping Tools for MATLAB (Greene
 631 et al., 2017). Surface temperatures were obtained from AMRC, SSEC, University of Wisconsin–
 632 Madison (<https://amrc.ssec.wisc.edu>). Tide data were generated by the CATS2008
 633 model (Padman et al., 2002). Ross Sea ice coverage was obtained from NASA NSIDC
 634 (Cavalieri et al., 1996, updated yearly). Code for this workflow is available at [https://](https://github.com/NeptuneProjects/RISClusterPT)
 635 github.com/NeptuneProjects/RISClusterPT.

636 References

- 637 Aggarwal, C. C., Hinneburg, A., & Keim, D. A. (2001). On the Surprising Behavior
 638 of Distance Metrics in High Dimensional Space. In G. Goos, J. Hartmanis,
 639 J. van Leeuwen, J. Van den Bussche, & V. Vianu (Eds.), *Database Theory*
 640 — *ICDT 2001* (Vol. 1973, pp. 420–434). Berlin, Heidelberg: Springer Berlin
 641 Heidelberg. doi: 10.1007/3-540-44503-X_27
- 642 Aggarwal, C. C., & Reddy, C. K. (Eds.). (2014). *Data clustering: Algorithms and*
 643 *applications*. Boca Raton: Chapman and Hall/CRC.
- 644 Alley, K., Scambos, T., Miller, J., Long, D., & MacFerrin, M. (2018, June). Quan-
 645 tifying vulnerability of Antarctic ice shelves to hydrofracture using microwave
 646 scattering properties. *Remote Sensing of Environment*, *210*, 297–306. doi:
 647 10.1016/j.rse.2018.03.025
- 648 Arthur, D., & Vassilvitskii, S. (2007). K-Means++: The advantages of careful
 649 seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on*
 650 *discrete algorithms* (pp. 1027–1035). USA: Society for Industrial and Applied
 651 Mathematics.
- 652 Aster, R. C., & Winberry, J. P. (2017, December). Glacial seismology. *Reports on*
 653 *Progress in Physics*, *80*(12), 126801. doi: 10.1088/1361-6633/aa8473
- 654 Baker, M. G., Aster, R. C., Anthony, R. E., Chaput, J., Wiens, D. A., Nyblade, A.,
 655 ... Stephen, R. A. (2019, December). Seasonal and spatial variations in the
 656 ocean-coupled ambient wavefield of the Ross Ice Shelf. *Journal of Glaciology*,
 657 *65*(254), 912–925. doi: 10.1017/jog.2019.64
- 658 Baker, M. G., Aster, R. C., Wiens, D. A., Nyblade, A., Bromirski, P. D., Ger-
 659 stoft, P., & Stephen, R. A. (2020, October). Teleseismic earthquake wave-

- 660 fields observed on the Ross Ice Shelf. *Journal of Glaciology*, 1–17. doi:
661 10.1017/jog.2020.83
- 662 Barcheck, C. G., Tulaczyk, S., Schwartz, S. Y., Walter, J. I., & Winberry, J. P.
663 (2018, March). Implications of basal micro-earthquakes and tremor for ice
664 stream mechanics: Stick-slip basal sliding and till erosion. *Earth and Planetary*
665 *Science Letters*, 486, 54–60. doi: 10.1016/j.epsl.2017.12.046
- 666 Beaucé, E., Frank, W. B., & Romanenko, A. (2018, January). Fast Matched Fil-
667 ter (FMF): An Efficient Seismic Matched-Filter Search for Both CPU and
668 GPU Architectures. *Seismological Research Letters*, 89(1), 165–172. doi:
669 10.1785/0220170181
- 670 Bellman, R. E. (1961). *Adaptive Control Processes: A Guided Tour*. Rand Corpora-
671 tion.
- 672 Bergen, K. J., & Beroza, G. C. (2018, June). Detecting earthquakes over a seismic
673 network using single-station similarity measures. *Geophysical Journal Interna-*
674 *tional*, 213(3), 1984–1998. doi: 10.1093/gji/ggy100
- 675 Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann,
676 J. (2010, May). ObsPy: A Python Toolbox for Seismology. *Seismological*
677 *Research Letters*, 81(3), 530–533. doi: 10.1785/gssrl.81.3.530
- 678 Bianco, M. J., & Gerstoft, P. (2018, December). Travel Time Tomography With
679 Adaptive Dictionaries. *IEEE Transactions on Computational Imaging*, 4(4),
680 499–511. doi: 10.1109/TCI.2018.2862644
- 681 Bianco, M. J., Gerstoft, P., Olsen, K. B., & Lin, F.-C. (2019, December). High-
682 resolution seismic tomography of Long Beach, CA using machine learning. *Sci-*
683 *entific Reports*, 9(1), 14987. doi: 10.1038/s41598-019-50381-z
- 684 Bindschadler, R. A., King, M. A., Alley, R. B., Anandakrishnan, S., & Padman, L.
685 (2003, August). Tidally Controlled Stick-Slip Discharge of a West Antarctic
686 Ice Stream. *Science*, 301(5636), 1087–1089. doi: 10.1126/science.1087231
- 687 Bindschadler, R. A., Vornberger, P. L., King, M. A., & Padman, L. (2003). Tidally
688 driven stick-slip motion in the mouth of Whillans Ice Stream, Antarctica. *An-*
689 *nals of Glaciology*, 36, 263–272. doi: 10.3189/172756403781816284
- 690 Bishop, C. (2006). *Pattern Recognition and Machine Learning* (First ed.). Springer-
691 Verlag New York.
- 692 Bromirski, P. D., Chen, Z., Stephen, R. A., Gerstoft, P., Arcas, D., Diez, A., . . . Ny-

- 693 blade, A. (2017, July). Tsunami and infragravity waves impacting Antarctic
694 ice shelves. *Journal of Geophysical Research: Oceans*, *122*(7), 5786–5801. doi:
695 10.1002/2017JC012913
- 696 Bromirski, P. D., Diez, A., Gerstoft, P., Stephen, R. A., Bolmer, T., Wiens, D. A.,
697 ... Nyblade, A. (2015, September). Ross ice shelf vibrations. *Geophysical*
698 *Research Letters*, *42*(18), 7589–7597. doi: 10.1002/2015GL065284
- 699 Bromirski, P. D., & Stephen, R. A. (2012). Response of the Ross Ice Shelf, Antarc-
700 tica, to ocean gravity-wave forcing. *Annals of Glaciology*, *53*(60), 163–172. doi:
701 10.3189/2012AoG60A058
- 702 Cavalieri, D. J., Parkinson, C. L., Gloersen, P., & Zwally, H. J. (1996, updated
703 yearly). *Sea Ice Concentrations from Nimbus-7 SMMR and DMSP SSM/I-*
704 *SSMIS Passive Microwave Data, Version 1*. Boulder, Colorado USA. NASA
705 National Snow and Ice Data Center Distributed Active Archive Center.
- 706 Chamarczuk, M., Nishitsuji, Y., Malinowski, M., & Draganov, D. (2020, January).
707 Unsupervised Learning Used in Automatic Detection and Classification of
708 Ambient-Noise Recordings from a Large-N Array. *Seismological Research*
709 *Letters*, *91*(1), 370–389. doi: 10.1785/0220190063
- 710 Chamberlain, C. J., Hopp, C. J., Boese, C. M., Warren-Smith, E., Chambers, D.,
711 Chu, S. X., ... Townend, J. (2018, January). EQcorrscan: Repeating and
712 Near-Repeating Earthquake Detection and Analysis in Python. *Seismological*
713 *Research Letters*, *89*(1), 173–181. doi: 10.1785/0220170151
- 714 Chaput, J., Aster, R. C., McGrath, D., Baker, M., Anthony, R. E., Gerstoft, P., ...
715 Stevens, L. A. (2018, October). Near-Surface Environmentally Forced Changes
716 in the Ross Ice Shelf Observed With Ambient Seismic Noise. *Geophysical*
717 *Research Letters*, *45*(20). doi: 10.1029/2018GL079665
- 718 Chazan, S. E., Gannot, S., & Goldberger, J. (2019, March). Deep Clustering Based
719 on a Mixture of Autoencoders. *arXiv:1812.06535 [cs, stat]*.
- 720 Chen, Z., Bromirski, P. D., Gerstoft, P., Stephen, R. A., Lee, W. S., Yun, S., ...
721 Nyblade, A. A. (2019, August). Ross Ice Shelf Icequakes Associated With
722 Ocean Gravity Wave Activity. *Geophysical Research Letters*, *46*(15), 8893–
723 8902. doi: 10.1029/2019GL084123
- 724 Chen, Z., Bromirski, P. D., Gerstoft, P., Stephen, R. A., Wiens, D. A., Aster, R. C.,
725 & Nyblade, A. A. (2018, October). Ocean-excited plate waves in the Ross and

- 726 Pine Island Glacier ice shelves. *Journal of Glaciology*, *64*(247), 730–744. doi:
727 10.1017/jog.2018.66
- 728 Cole, H. M. (2020). *Tidally Induced Seismicity at the Grounded Margins of the*
729 *Ross Ice Shelf, Antarctica* (Master’s Thesis). Colorado State University, Fort
730 Collins, Colorado.
- 731 De Angelis, H., & Skvarca, P. (2003, March). Glacier Surge After Ice Shelf Collapse.
732 *Science*, *299*(5612), 1560–1562. doi: 10.1126/science.1077987
- 733 Diez, A., Bromirski, P., Gerstoft, P., Stephen, R., Anthony, R., Aster, R., . . . Wiens,
734 D. (2016, May). Ice shelf structure derived from dispersion curve analysis
735 of ambient seismic noise, Ross Ice Shelf, Antarctica. *Geophysical Journal*
736 *International*, *205*(2), 785–795. doi: 10.1093/gji/ggw036
- 737 Dupont, T. K., & Alley, R. B. (2005). Assessment of the importance of ice-shelf but-
738 tressing to ice-sheet flow. *Geophysical Research Letters*, *32*(4). doi: 10.1029/
739 2004GL022024
- 740 Fürst, J. J., Durand, G., Gillet-Chaulet, F., Tavard, L., Rankl, M., Braun, M., &
741 Gagliardini, O. (2016, May). The safety band of Antarctic ice shelves. *Nature*
742 *Climate Change*, *6*(5), 479–482. doi: 10.1038/nclimate2912
- 743 Gibbons, S. J., & Ringdal, F. (2006, April). The detection of low magnitude seis-
744 mic events using array-based waveform correlation. *Geophysical Journal Inter-*
745 *national*, *165*(1), 149–166. doi: 10.1111/j.1365-246X.2006.02865.x
- 746 Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- 747 Greene, C. A., Gwyther, D. E., & Blankenship, D. D. (2017, July). Antarctic Map-
748 ping Tools for Matlab. *Computers & Geosciences*, *104*, 151–157. doi: 10.1016/
749 j.cageo.2016.08.003
- 750 Hartigan, J. A., & Wong, M. A. (1979). Algorithm AS 136: A K-Means Clustering
751 Algorithm. *Applied Statistics*, *28*(1), 100. doi: 10.2307/2346830
- 752 Hell, M. C., Cornelle, B. D., Gille, S. T., Miller, A. J., & Bromirski, P. D. (2019,
753 November). Identifying Ocean Swell Generation Events from Ross Ice Shelf
754 Seismic Data. *Journal of Atmospheric and Oceanic Technology*, *36*(11), 2171–
755 2189. doi: 10.1175/JTECH-D-19-0093.1
- 756 Hinton, G. E. (2006, July). Reducing the Dimensionality of Data with Neural Net-
757 works. *Science*, *313*(5786), 504–507. doi: 10.1126/science.1127647
- 758 Holtzman, B. K., Paté, A., Paisley, J., Waldhauser, F., & Repetto, D. (2018,

- 759 May). Machine learning reveals cyclic changes in seismic source spec-
 760 tra in Geysers geothermal field. *Science Advances*, 4(5), eaao2929. doi:
 761 10.1126/sciadv.aao2929
- 762 Hotovec-Ellis, A. J., & Jeffries, C. (2016, April). *Near Real-time Detection, Cluster-*
 763 *ing, and Analysis of Repeating Earthquakes: Application to Mount St. Helens*
 764 *and Redoubt Volcanoes* [Invited]. Reno, NV, USA.
- 765 Hubbard, B., Luckman, A., Ashmore, D. W., Bevan, S., Kulesa, B.,
 766 Kuipers Munneke, P., . . . Rutt, I. (2016, September). Massive subsurface
 767 ice formed by refreezing of ice-shelf melt ponds. *Nature Communications*, 7(1),
 768 11897. doi: 10.1038/ncomms11897
- 769 Johnson, C. W., Ben-Zion, Y., Meng, H., & Vernon, F. (2020, August). Identifying
 770 Different Classes of Seismic Noise Signals Using Unsupervised Learning. *Geo-*
 771 *physical Research Letters*, 47(15). doi: 10.1029/2020GL088353
- 772 Johnson, C. W., Meng, H., Vernon, F., & Ben-Zion, Y. (2019, August). Characteris-
 773 tics of Ground Motion Generated by Wind Interaction With Trees, Structures,
 774 and Other Surface Obstacles. *Journal of Geophysical Research: Solid Earth*,
 775 124(8), 8519–8539. doi: 10.1029/2018JB017151
- 776 Kingma, D. P., & Ba, J. (2017, January). Adam: A Method for Stochastic Opti-
 777 mization. *arXiv:1412.6980 [cs]*.
- 778 Klein, E., Mosbeux, C., Bromirski, P. D., Padman, L., Bock, Y., Springer, S. R., &
 779 Fricker, H. A. (2020, October). Annual cycle in flow of Ross Ice Shelf, Antarc-
 780 tica: Contribution of variable basal melting. *Journal of Glaciology*, 66(259),
 781 861–875. doi: 10.1017/jog.2020.61
- 782 Kong, Q., Trugman, D. T., Ross, Z. E., Bianco, M. J., Meade, B. J., & Gerstoft, P.
 783 (2019, January). Machine Learning in Seismology: Turning Data into Insights.
 784 *Seismological Research Letters*, 90(1), 3–14. doi: 10.1785/0220180259
- 785 Kullback, S., & Leibler, R. A. (1951, March). On Information and Sufficiency.
 786 *The Annals of Mathematical Statistics*, 22(1), 79–86. doi: 10.1214/aoms/
 787 1177729694
- 788 MacAyeal, D. R., Banwell, A. F., Okal, E. A., Lin, J., Willis, I. C., Goodsell, B.,
 789 & MacDonald, G. J. (2019, September). Diurnal seismicity cycle linked to
 790 subsurface melting on an ice shelf. *Annals of Glaciology*, 60(79), 137–157. doi:
 791 10.1017/aog.2018.29

- 792 MacQueen, J. (1967). Some methods for classification and analysis of multivariate
 793 observations. In *Proceedings of the fifth berkeley symposium on mathematical*
 794 *statistics and probability, volume 1: Statistics* (pp. 281–297). Berkeley, Calif.:
 795 University of California Press.
- 796 Min, E., Guo, X., Liu, Q., Zhang, G., Cui, J., & Long, J. (2018). A Survey of Clus-
 797 tering With Deep Learning: From the Perspective of Network Architecture.
 798 *IEEE Access*, *6*, 39501–39514. doi: 10.1109/ACCESS.2018.2855437
- 799 Morlighem, M., Williams, C. N., Rignot, E., An, L., Arndt, J. E., Bamber, J. L., ...
 800 Zinglensen, K. B. (2017, November). BedMachine v3: Complete Bed Topog-
 801 raphy and Ocean Bathymetry Mapping of Greenland From Multibeam Echo
 802 Sounding Combined With Mass Conservation. *Geophysical Research Letters*,
 803 *44*(21). doi: 10.1002/2017GL074954
- 804 Mousavi, S. M., Horton, S. P., Langston, C. A., & Samei, B. (2016, October).
 805 Seismic features and automatic discrimination of deep and shallow induced-
 806 microearthquakes using neural network and logistic regression. *Geophysical*
 807 *Journal International*, *207*(1), 29–46. doi: 10.1093/gji/ggw258
- 808 Mousavi, S. M., Zhu, W., Ellsworth, W., & Beroza, G. (2019, November). Un-
 809 supervised Clustering of Seismic Signals Using Deep Convolutional Autoen-
 810 coders. *IEEE Geoscience and Remote Sensing Letters*, *16*(11), 1693–1697. doi:
 811 10.1109/LGRS.2019.2909218
- 812 Murphy, K. P. (2012). *Machine learning: A probabilistic perspective*. Cambridge,
 813 MA: MIT Press.
- 814 Nicolas, J. P., Vogelmann, A. M., Scott, R. C., Wilson, A. B., Cadeddu, M. P.,
 815 Bromwich, D. H., ... Wille, J. D. (2017, August). January 2016 extensive
 816 summer melt in West Antarctica favoured by strong El Niño. *Nature Commu-*
 817 *nications*, *8*(1), 15799. doi: 10.1038/ncomms15799
- 818 Olinger, S. D., Lipovsky, B. P., Wiens, D. A., Aster, R. C., Bromirski, P. D., Chen,
 819 Z., ... Stephen, R. A. (2019, June). Tidal and Thermal Stresses Drive Seismic-
 820 ity Along a Major Ross Ice Shelf Rift. *Geophysical Research Letters*, *46*(12),
 821 6644–6652. doi: 10.1029/2019GL082842
- 822 Padman, L., Fricker, H. A., Coleman, R., Howard, S., & Erofeeva, L. (2002). A new
 823 tide model for the Antarctic ice shelves and seas. *Annals of Glaciology*, *34*,
 824 247–254. doi: 10.3189/172756402781817752

- 825 Paolo, F. S., Fricker, H. A., & Padman, L. (2015). Volume loss from Antarctic ice
826 shelves is accelerating. *Science*, *348*(6232), 327–331. doi: 10.1126/science
827 .aaa0940
- 828 Perol, T., Gharbi, M., & Denolle, M. (2018, February). Convolutional neural net-
829 work for earthquake detection and location. *Science Advances*, *4*(2), e1700578.
830 doi: 10.1126/sciadv.1700578
- 831 Pritchard, H. D., Ligtenberg, S. R. M., Fricker, H. A., Vaughan, D. G., van den
832 Broeke, M. R., & Padman, L. (2012, April). Antarctic ice-sheet loss
833 driven by basal melting of ice shelves. *Nature*, *484*(7395), 502–505. doi:
834 10.1038/nature10968
- 835 Reddy, T. A., Devi, K. R., & Gangashetty, S. V. (2012, March). Nonlinear principal
836 component analysis for seismic data compression. In *2012 1st International
837 Conference on Recent Advances in Information Technology (RAIT)* (pp. 927–
838 932). Dhanbad, India: IEEE. doi: 10.1109/RAIT.2012.6194558
- 839 Riahi, N., & Gerstoft, P. (2017, March). Using graph clustering to locate sources
840 within a dense sensor array. *Signal Processing*, *132*, 110–120. doi: 10.1016/j
841 .sigpro.2016.10.001
- 842 Rignot, E., Mouginot, J., Morlighem, M., Seroussi, H., & Scheuchl, B. (2014).
843 Widespread, rapid grounding line retreat of Pine Island, Thwaites, Smith, and
844 Kohler glaciers, West Antarctica, from 1992 to 2011. *Geophysical Research
845 Letters*, *41*(10), 3502–3509. doi: 10.1002/2014GL060140
- 846 Rousseeuw, P. J. (1987, November). Silhouettes: A graphical aid to the interpre-
847 tation and validation of cluster analysis. *Journal of Computational and Applied
848 Mathematics*, *20*, 53–65. doi: 10.1016/0377-0427(87)90125-7
- 849 Scambos, T. A., Bohlander, J. A., Shuman, C. A., & Skvarca, P. (2004). Glacier
850 acceleration and thinning after ice shelf collapse in the Larsen B embay-
851 ment, Antarctica. *Geophysical Research Letters*, *31*(18). doi: 10.1029/
852 2004GL020670
- 853 Seydoux, L., Balestrieri, R., Poli, P., de Hoop, M., Campillo, M., & Baraniuk, R.
854 (2020, December). Clustering earthquake signals and background noises in
855 continuous seismic data with unsupervised deep learning. *Nature Communica-
856 tions*, *11*(1), 3972. doi: 10.1038/s41467-020-17841-x
- 857 Smith, B., Fricker, H. A., Gardner, A. S., Medley, B., Nilsson, J., Paolo, F. S., . . .

- 858 Zwally, H. J. (2020, June). Pervasive ice sheet mass loss reflects compet-
 859 ing ocean and atmosphere processes. *Science*, *368*(6496), 1239–1242. doi:
 860 10.1126/science.aaz5845
- 861 Snover, D., Johnson, C. W., Bianco, M. J., & Gerstoft, P. (2020). Deep clustering
 862 to identify sources of urban seismic noise in Long Beach, CA. *Geophysical Re-*
 863 *search Letters*, *33*.
- 864 Steinbach, M., Ertöz, L., & Kumar, V. (2004). The Challenges of Clustering High
 865 Dimensional Data. In L. T. Wille (Ed.), *New Directions in Statistical Physics:*
 866 *Econophysics, Bioinformatics, and Pattern Recognition* (pp. 273–309). Berlin,
 867 Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/978-3-662-08968-2_16
- 868 Swindell, H., & Snell, N. (1977). *Station processor automatic signal detection sys-*
 869 *tem, phase I: Final Report, Station Processor Software Development, Texas*
 870 *Instruments Report No. ALEX (01)-FR-77-01*. (Tech. Rep.). Texas Instru-
 871 ments.
- 872 Telesca, L., & Chelidze, T. (2018, November). Visibility Graph Analysis of Seismic-
 873 ity around Enguri High Arch Dam, Caucasus. *Bulletin of the Seismological So-*
 874 *ciety of America*, *108*(5B), 3141–3147. doi: 10.1785/0120170370
- 875 Thoma, M., Jenkins, A., Holland, D., & Jacobs, S. (2008). Modelling circumpolar
 876 deep water intrusions on the amundsen sea continental shelf, antarctica. *Geo-*
 877 *physical Research Letters*, *35*(18). doi: 10.1029/2008GL034939
- 878 Tibshirani, R., Walther, G., & Hastie, T. (2001, May). Estimating the number
 879 of clusters in a data set via the gap statistic. *Journal of the Royal Statistical*
 880 *Society: Series B (Statistical Methodology)*, *63*(2), 411–423. doi: 10.1111/1467
 881 -9868.00293
- 882 Trugman, D. T., & Shearer, P. M. (2017, March). GrowClust: A Hierarchical Clus-
 883 tering Algorithm for Relative Earthquake Relocation, with Application to the
 884 Spanish Springs and Sheldon, Nevada, Earthquake Sequences. *Seismological*
 885 *Research Letters*, *88*(2A), 379–391. doi: 10.1785/0220160188
- 886 van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of*
 887 *Machine Learning Research*.
- 888 Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., & Manzagol, P.-A. (2010, Decem-
 889 ber). Stacked Denoising Autoencoders: Learning Useful Representations in a
 890 Deep Network with a Local Denoising Criterion. *Journal of Machine Learning*

- 891 *Research*, 38.
- 892 Wiens, D., & Bromirski, P. (2014). *Collaborative Research: Dynamic Response of the*
893 *Ross Ice Shelf to Wave-Induced Vibrations, and Collaborative Research: Mantle*
894 *Structure and Dynamics of the Ross Sea from a Passive Seismic Deployment*
895 *on the Ross Ice Shelf.* International Federation of Digital Seismograph Net-
896 works.
- 897 Wiens, D. A., Anandakrishnan, S., Winberry, J. P., & King, M. A. (2008, June). Si-
898 multaneous teleseismic and geodetic observations of the stick–slip motion of an
899 Antarctic ice stream. *Nature*, 453(7196), 770–774. doi: 10.1038/nature06990
- 900 Withers, M., Aster, R., Young, C., Beiriger, J., Harris, M., Moore, S., & Trujillo, J.
901 (1998, February). A Comparison of Select Trigger Algorithms for Automated
902 Global Seismic Phase and Event Detection. *Bulletin of the Seismological*
903 *Society of America*, 88(1), 12.
- 904 Xie, J., Girshick, R., & Farhadi, A. (2016). Unsupervised Deep Embedding for Clus-
905 tering Analysis. *Proceedings of the 33rd international conference on machine*
906 *learning*, 10.
- 907 Yang, B., Fu, X., Sidiropoulos, N. D., & Hong, M. (2017, June). Towards
908 K-means-friendly Spaces: Simultaneous Deep Learning and Clustering.
909 *arXiv:1610.04794 [cs]*.
- 910 Yoon, C. E., O’Reilly, O., Bergen, K. J., & Beroza, G. C. (2015, December). Earth-
911 quake detection through computationally efficient similarity search. *Science*
912 *Advances*, 1(11), e1501057. doi: 10.1126/sciadv.1501057