

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 identified classes of seismic signals with similar spec-
8 tral and temporal features.
- 9 • Deep embedded clustering can be adapted to various kinds of data sets, enabling
10 rapid exploration of “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

We demonstrate the ability of a machine learning technique called deep embedded clustering (DEC) 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 & Gerstoft, 2018; 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. A related approach based on sparse modeling, called dictionary learning, has been applied to regularizing seismic inverse problems (Bianco & Gerstoft, 2018; Bianco et al., 2019).

77 To test the applicability of clustering groups of similar signals for monitoring ice
78 shelves, we focus specifically on the Ross Ice Shelf (RIS), Antarctica, where a 34-station
79 passive seismic array was deployed from November 2014 to January 2017 to observe the
80 response of the RIS to ocean gravity wave impacts and investigate the structural dynam-
81 ics of the ice shelf (Bromirski et al., 2015). The array, shown in Figure 1, continuously
82 recorded long- and short-period seismic signals that exhibited seasonal and spatial vari-
83 ations related to the shelf’s coupling to the ocean, atmosphere, and crust (Baker et al.,
84 2019). Signals and ambient noise of interest on the RIS include tidally-driven stick-slip
85 seismicity at Whillans Ice Stream (Bindschadler, King, et al., 2003; Bindschadler, Vorn-
86 berger, et al., 2003; D. A. Wiens et al., 2008); basal micro-earthquakes and tremor (Barcheck
87 et al., 2018); tidally and thermally driven rift fractures (Olinger et al., 2019); diurnal seis-
88 micity associated with subsurface melting (MacAyeal et al., 2019); wind-generated res-
89 onance in the ice (Chaput et al., 2018); flexural and plate waves generated by ocean swell,
90 infragravity waves, and tsunamis (Bromirski & Stephen, 2012; Bromirski et al., 2017; Chen
91 et al., 2018); regional and teleseismic earthquakes (Baker et al., 2020); and icequakes gen-
92 erated by ocean gravity waves (Chen et al., 2019). Ambient seismic noise, which can be
93 used to estimate the RIS structure (Diez et al., 2016), also contains spectra from ocean
94 gravity waves, whose dispersion can be used to identify their source distance and origin
95 (Bromirski et al., 2015; Hell et al., 2019).

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

102 **2 Background**

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

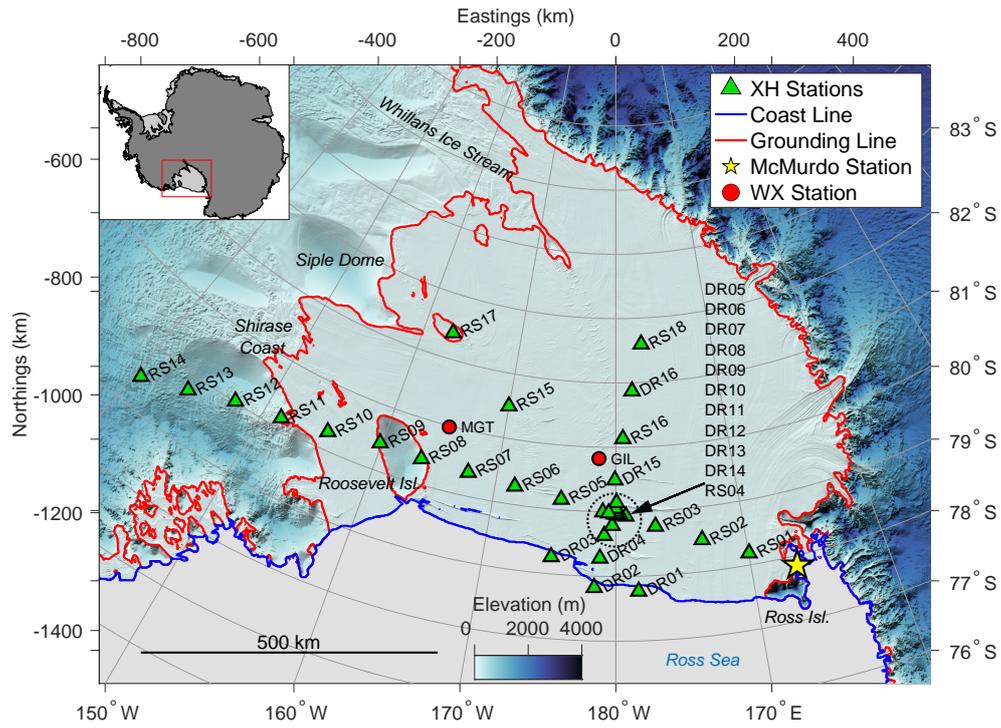


Figure 1. The passive broadband seismic array deployed from November 2014 to January 2017 consisted of 34 seismic stations and was deployed as part of the Ross Ice Shelf Dynamic Response to Wave-Induced Vibrations Project (Bromirski et al., 2015). RIS surface elevation, ice and water layer thicknesses,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

106 **2.1 Clustering**

107 There are numerous methods to cluster data, (Aggarwal & Reddy, 2014), many of
108 which have been adapted for use in seismology and geophysics (Kong et al., 2019). Hi-
109 erarchical clustering has been used by Mousavi et al. (2016) to automatically discrim-
110 inate between shallow and deep earthquakes, and by Trugman and Shearer (2017) to more
111 precisely localize earthquakes. Graphical clustering has been used to localize sources in
112 a dense seismic array by Riahi and Gerstoft (2017), and by Telesca and Chelidze (2018)
113 to cluster seismic events in time. Distance-based clustering, like the popular k -means al-
114 gorithm, (MacQueen, 1967; Hartigan & Wong, 1979) has been used by Chamarczuk et
115 al. (2020) to cluster seismicity based on features extracted from seismic data. Perol et
116 al. (2018) used k -means to define probabilistic earthquake locations as part of their con-
117 volutional neural network (CNN) detection and localization technique. A novel approach
118 was presented by Seydoux et al. (2020), who detect and cluster seismic signals and back-
119 ground noise with the use of a deep scattering neural network and a Gaussian mixture
120 model.

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

132 **2.2 Dimensionality**

133 Data are considered high-dimensional when many features are required to repre-
134 sent or describe the data. Seismic data represented as time series, spectrograms, scalo-
135 grams, or energy envelopes can contain thousands of features (e.g., discrete samples in
136 a time series, or bins in a spectrogram). Clustering performed directly on such input data

(Aggarwal & Reddy, 2014) is vulnerable to the “curse of dimensionality” (Bellman, 1961), i.e., as the dimensionality of the input data increases, the number of data points required to maintain sufficient sampling density increases exponentially. A further consideration is that clustering error metrics can give less meaningful results as dimensionality increases.

As high-dimensional data are difficult to cluster (Aggarwal et al., 2001; Steinbach et al., 2004), dimensionality reduction remains a major focus of development (Yang et al., 2017). It is often desirable to transform the input data to a lower-dimensional representation described by fewer, more salient features. A popular approach is to use principal component analysis (PCA), which projects higher dimensional data into lower dimensional space (Goodfellow et al., 2016) and was used by Reddy et al. (2012) to compress seismic data to maximize feature variance.

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

2.3 Clustering in Reduced Dimensions

A method that has shown improvement over traditional clustering techniques was developed by Xie et al. (2016), whose *deep embedded clustering* (DEC) uses the latent feature space as input to an adaptive clustering algorithm. DEC consists of two processes: (1) An autoencoder is trained to represent the data’s salient features; and (2) the encoding layers and clustering layer are jointly optimized. Yang et al. (2017) extend the

167 approach in DEC by jointly optimizing the clustering step with training the entire au-
168 toencoder, not just the encoder layers.

169 Additional variations of DEC have been proposed: Xie et al. (2016) used a stacked
170 de-noising autoencoder (Vincent et al., 2010), and (Min et al., 2018) employed autoen-
171 coders composed of CNN layers and other architectures. More recently, Chazan et al.
172 (2019) developed a novel approach in which joint clustering is performed with a mixture
173 of autoencoders, each representing a cluster. Mousavi et al. (2019) used DEC to predict
174 whether seismic detections were local or teleseismic, and Snover et al. (2020) demonstrated
175 DEC’s ability to cluster anthropogenically generated seismic noise.

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

180 **3 Ross Ice Shelf (RIS) Seismic Array and Data**

181 Each station in the RIS seismic array consisted of 3-component Nanometrics Tril-
182 lium 120 PHQ seismometers emplaced 1 m below the surface of the ice, powered by so-
183 lar panels during the austral summers, and lithium-ion batteries during the austral win-
184 ters. Two subarrays comprised the array. The larger subarray consisted of 18 stations
185 spaced approximately 80 km apart (prefix RS), primarily oriented parallel to the RIS
186 front. The RS stations sampled short-period orthogonal components of ground veloc-
187 ity at a sampling rate of 100 Hz, except for two stations that sampled at 200 Hz. The
188 smaller subarray consisted of 16 stations (prefix DR) arranged approximately orthog-
189 onal to the icefront along the international date line, sampling ground velocity with a
190 sampling rate of 200 Hz. For this study, we were primarily interested in the detection
191 and classification of icequakes and local/regional earthquakes, using only vertical com-
192 ponent observations with frequencies of interest occurring between 3 and 20 Hz. Rep-
193 resentative types of signals detected are shown in Figure 2.

194 Seismic data from each station were processed in 24-hour segments as follows: 1) Data
195 were linearly de-trended and tapered with a Hann window. 2) Instrument responses for
196 all stations were removed, giving acceleration in m/s^2 . 3) Since the bandwidth of inter-
197 est was from 3 to 20 Hz, data were decimated to 50 Hz, using low-pass filtering followed-

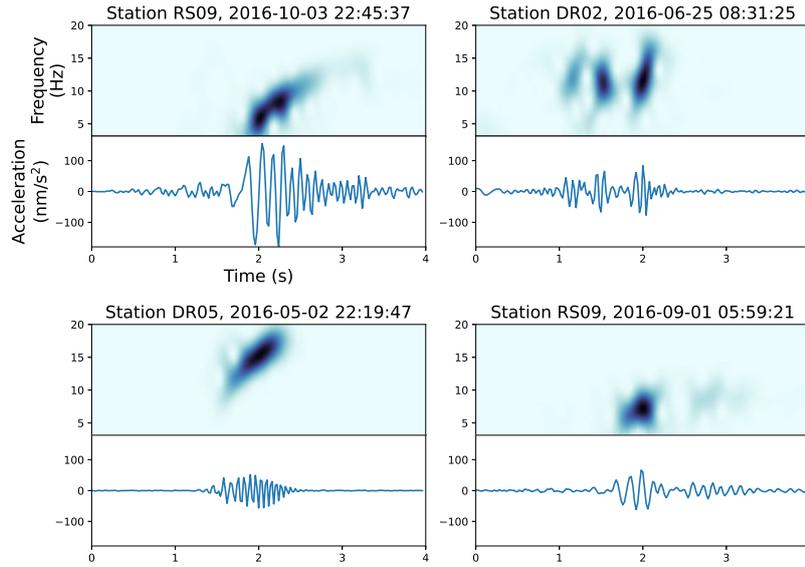


Figure 2. Seismic signals detected on the Ross Ice Shelf were diverse with variation in time, space, and source mechanism. Shown are examples of acceleration response seismograms and their respective normalized spectrograms spanning the 3-20 Hz band that were typical for the data set. The normalized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

198 by downsampling. 4) A band-pass filter with cutoff frequencies at 3 and 20 Hz was ap-
 199 plied to remove long-period signals originating from tides, tsunamis, infragravity waves,
 200 ocean swell, and teleseisms. 5) An event detection algorithm, the Z-detector (Swindell
 201 & Snell, 1977; Withers et al., 1998), was used to detect impulsive signals, particularly
 202 icequakes and local earthquakes, with a sliding window of 3 s. The detector was applied
 203 to data from each station between 2 December 2014 and 20 November 2016 for a total
 204 of 719 days of array data, yielding 427,798 detections.

205 Upon detection, a 4 s trace centered on the spectral peak of each triggered event
 206 was saved for processing. For each seismic trace saved, a spectrogram was computed us-
 207 ing the short-time Fourier transform with a 0.4 s Kaiser window, NFFT=256, and 90%
 208 overlap. Spectrograms (samples) contained one channel of amplitude information, 87 fre-
 209 quency bins, and 100 time bins for a total of 8,700 features per spectrogram. Finally, sam-
 210 plewise normalization was performed by subtracting the mean spectral level from each
 211 spectrogram, subsequently normalizing to the interval [0, 1].

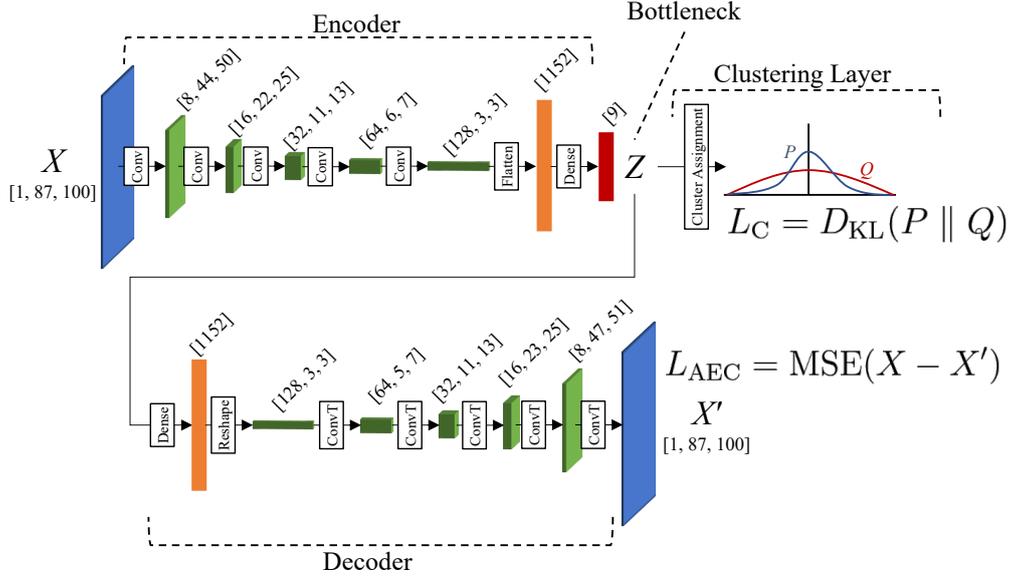


Figure 3. The deep embedded clustering model uses a convolutional autoencoder that encodes the data space X into the latent feature space Z , and a decoder that recovers the original input X from Z . The mean squared error (MSE) between the input X and the reconstruction X' is used as the autoencoder 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 $X \rightarrow Z \rightarrow X'$.

212 4 Deep Embedded Clustering Model

213 The objective of the DEC model, shown in Figure 3, is to encode the input data—
 214 in this case, spectrograms of seismic signals—into a layer containing latent (lower-dimensional)
 215 features, called the *embedded* layer, to which a clustering algorithm is applied. The out-
 216 come of the clustering performance is then used to refine both the autoencoder model
 217 and the clustering layer in an effort to obtain more accurate latent space embeddings while
 218 improving clustering performance. In the implementation that follows, the 8,700 features
 219 of an input spectrogram are reduced to a latent feature space of just 9 embedded fea-
 220 tures with the use of a convolutional autoencoder, a type of DNN composed of convo-
 221 lutional layers and their transposes.

222

4.1 Dimensionality Reduction with a Convolutional Autoencoder

223

224

225

226

227

228

229

230

231

Autoencoders provide a useful means of data approximation using a lower-dimensional representation via a sequence of non-linear transformations. The autoencoder model consists of three components: an *encoder*, a *bottleneck*, and a *decoder* (Murphy, 2012). First, the encoder maps input data from a data space X into a latent feature space Z , which is contained within the bottleneck of the model. Next, the decoder attempts to reconstruct X from Z . This process is performed iteratively with the objective of minimizing the error between X and the decoder output, X' . In minimizing the error, the autoencoder learns the salient features of X and accurately encodes them in Z , thus reducing the dimensionality of the clustering task.

232

233

234

235

236

237

238

239

240

241

242

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

243

244

245

246

247

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

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

248

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

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

249

250

Performing this calculation over the entire data set at once is computationally expensive, memory intensive, and can lead to poor convergence. Instead, the loss is calculated

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

251 in mini-batch subsets of the data space. For each mini-batch loss, stochastic gradient de-
252 scent (Goodfellow et al., 2016) is used to update the weights. When all mini-batches have
253 been processed, the next training epoch begins and the process is repeated. After each
254 epoch, a subset of the data separate from the training data is used to validate the model’s
255 performance without updating the weights, yielding a validation MSE. Training is per-
256 formed until a specified maximum number of epochs is reached, or stopped early if the
257 validation MSE fails to decrease below its minimum value after ten epochs. The early
258 stopping criterion prevents the autoencoder from overfitting the training data.

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

Table 2. *Sample Sizes and Hyperparameters used to Train the 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)
427,798	100,000	25,000	0.001	1024	8	0.05

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

265 Autoencoder training is implemented using 125,000 spectrograms randomly selected
 266 without replacement from the 427,298 detections. Of the selected spectrograms, 80% are
 267 used for training and 20% for validation. The trainable parameters are optimized using
 268 the Adaptive Moment Estimation (Adam) algorithm (Kingma & Ba, 2017). In training,
 269 there are two principal hyperparameters to address. First is the initial learning rate, which
 270 controls the initial step size used by Adam to step down the gradient of the loss. The
 271 second hyperparameter is the mini-batch size, which sets the number of spectrograms
 272 to be passed through the model at one time. The optimal configuration is found through
 273 a grid search of the hyperparameters. A summary of the optimal hyperparameters and
 274 the number of spectrograms used are listed in Table 2. As seen in Figure 4a, training
 275 and validation losses fall off exponentially with each training epoch until the early stop-
 276 ping criterion is met; in this case, after 49 epochs. The effectiveness of the autoencoder’s
 277 ability to reconstruct the input spectrogram is illustrated in Figure 5. Though some loss
 278 of resolution in time and frequency is expected due to the convolutional layers and their
 279 transposes, the structure of the spectrogram is largely preserved, with the salient infor-
 280 mation of the input encoded into the latent feature space.

281 4.2 Clustering Layer

282 In the DEC framework, clustering is performed in the latent feature space, Z , with
 283 the goal of finding K distinct classes of signals within the data. We assume that the data
 284 form clusters which are separable in Z space, and that these clusters coalesce around unique
 285 locations $\{\boldsymbol{\mu}_j \in Z\}_{j=1}^K$, i.e., centroids around which other similar signals may be found.

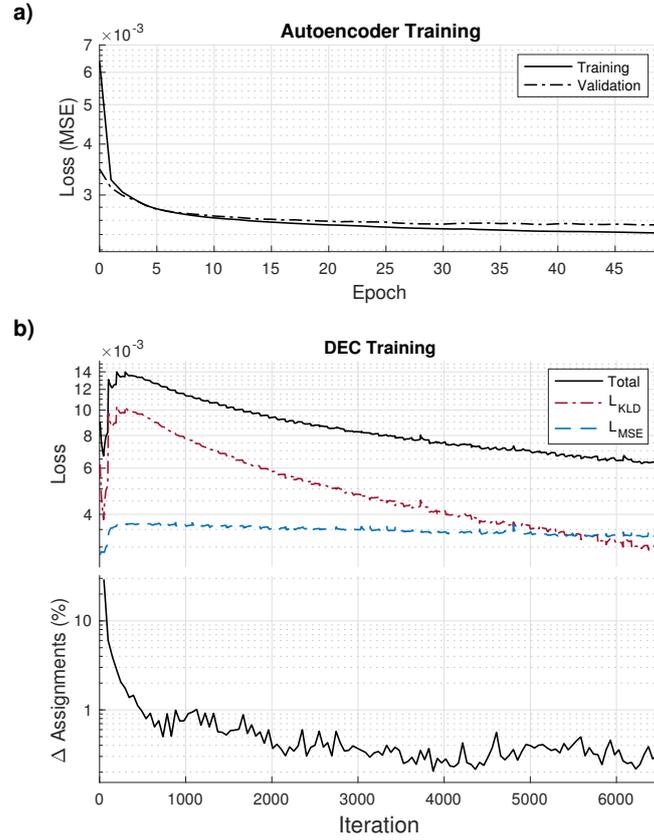


Figure 4. (a) Training and validation losses during autoencoder training. To avoid overfitting the model, training is stopped when the early stopping criterion is met (in this case, at 49 epochs). (b) In the upper plot, loss curves are shown for deep embedded clustering (DEC). In the lower plot, the percentage of samples which undergo class reassignment at each update interval is shown; training is stopped once the change is less than 0.2%

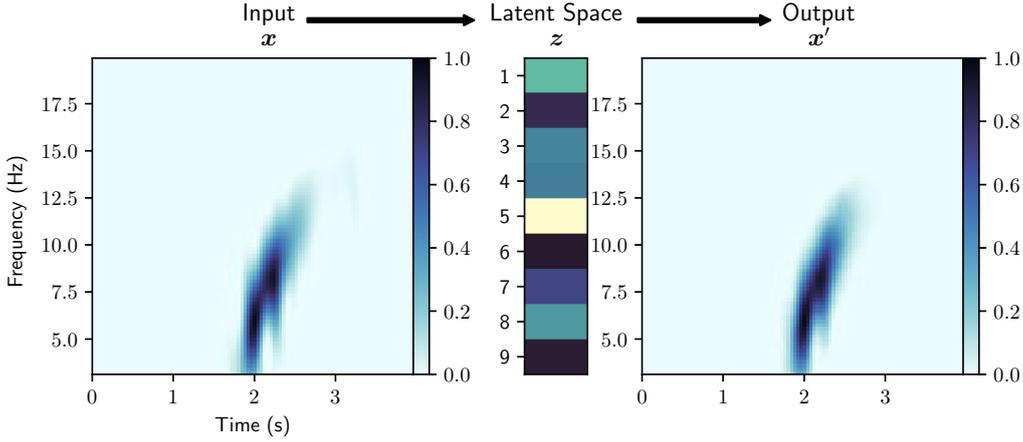


Figure 5. A trained autoencoder takes an input spectrogram \mathbf{x} , encodes it into an embedded, 9-dimensional latent feature vector \mathbf{z} , then reconstructs the input as \mathbf{x}' . The autoencoder preserves features correlated within a given cluster and discards the remaining signal. This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

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

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

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

289 Centroids are initialized with the Gaussian mixture model (GMM) clustering algo-
 290 rithm (Bishop, 2006). Among signal classes, GMM accounts for differences in sample
 291 size and distributions that are expected in natural seismic data by treating the latent
 292 data as a mixture of Gaussian distributions, each with a unique centroid and variance.
 293 The clusters and centroids are then updated using an expectation-maximization (EM)
 294 algorithm. To accelerate EM convergence, k -means clustering (Hartigan & Wong, 1979)
 295 is used to initialize the GMM clustering algorithm.

296 With the centroids initialized, DEC seeks to further improve clustering by using
 297 the difference between the embedded spectrograms and the cluster centroids as an ad-
 298 ditional loss function for updating model parameters. Because the input data is unlabeled,
 299 a self-supervised method is required. We implement the method developed by Xie
 300 et al. (2016), who, drawing from the t-Distributed Stochastic Neighbor Embedding (t-
 301 SNE) algorithm (van der Maaten & Hinton, 2008), propose measuring the difference be-

302 tween a t-Student’s distribution kernel of the embedded spectrograms and an auxiliary
 303 target distribution. A simplified Student’s t-distribution is used to measure the similar-
 304 ity between an embedded spectrogram, \mathbf{z}_i , and the cluster centroids $\boldsymbol{\mu}_j$:

$$q_{ij} = \frac{(1 + \|\mathbf{z}_i - \boldsymbol{\mu}_j\|^2)^{-1}}{\sum_j (1 + \|\mathbf{z}_i - \boldsymbol{\mu}_j\|^2)^{-1}}. \quad (4)$$

305 Equation (4) results in a set of soft class assignments, i.e., the probability that embed-
 306 ded spectrogram i will be assigned to class j , which are used to compute the auxiliary
 307 target distribution, p . The form of p is designed to improve clustering performance, em-
 308 phasize embeddings with high-confidence assignments, and normalize each cluster cen-
 309 troid’s contribution to the loss function so that large clusters minimally distort Z (Xie
 310 et al., 2016):

$$p_{ij} = \frac{q_{ij}^2 / \sum_i q_{ij}}{\sum_j (q_{ij}^2 / \sum_i q_{ij})}. \quad (5)$$

311 The dissimilarity between the distributions given by equations (4) and (5) is measured
 312 using the Kullback-Leibler divergence (Kullback & Leibler, 1951). From the divergence
 313 the clustering layer’s loss function is obtained:

$$L_C = D_{\text{KL}}(P \parallel Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}}. \quad (6)$$

314 In DEC, the clustering layer is attached to the trained autoencoder’s bottleneck,
 315 as shown in Figure 3. During training of the DEC model, the loss functions from equa-
 316 tions (1) and (6) are combined into a total loss function,

$$L = L_{\text{AEC}} + \lambda L_C, \quad (7)$$

317 where λ is a hyperparameter that balances the contributions of the two losses, since they
 318 are of differing magnitudes. λ must be tuned: if it is too large, the clustering loss will
 319 cause model instability and lead to distortion of the latent space, in which case the la-
 320 tent space will no longer represent the salient features of the data. If λ is too small, the
 321 effect on clustering performance will be minimal. We found that $\lambda = 0.05$ yielded op-
 322 timal performance for model training and clustering.

323 Two constituent processes occur simultaneously during DEC model training. First,
 324 the full loss from equation (7) is backpropagated through the DEC model parameters,
 325 which include the autoencoder as well as the cluster centroids. Second, to account for
 326 the cluster centroids changing as training progresses, the distributions q_{ij} and p_{ij} are up-
 327 dated at intervals. The update interval is a hyperparameter that must be tuned. Through

hyperparameter tuning, an update interval of twice per training epoch was found to be optimal for clustering performance, minimizing DEC loss, and training within a reasonable time frame. Training is stopped after the number of samples changing assignments after every update interval reaches less than 0.2% of the total number of samples.

4.3 Selecting Optimal Number of Clusters

Determining the optimal number of clusters, K , is a major challenge in unsupervised machine learning. Although there are statistical methods available for choosing the optimal number of clusters (Rousseeuw, 1987; Tibshirani et al., 2001), in this study we treat K as a hyperparameter, iterating the DEC workflow over a range of values for K and evaluating the results to choose the best value. Results are evaluated both quantitatively and qualitatively. Quantitative evaluation is performed for each class by examining cumulative distribution functions and probability density functions as functions of distance to each class centroid, $d_{i,j}$ (equation (3)). The qualitative approach is to visually inspect the similarity of the latent feature vectors \mathbf{z}_i to their respective class centroids $\boldsymbol{\mu}_j$, and to see if the spectrograms and seismograms assigned to each class likewise exhibit similarity. In general, the formation of two or more similar classes may indicate that too many classes were initialized, and the data can be grouped into a single class in post-processing. Too much variance among the spectrograms within a class may indicate the need for one or more additional classes. We found that $K = 8$ was the optimal number of classes for the RIS data set.

5 Results

5.1 Clustering Performance

Clustering with DEC results in two distinct phases: first, the GMM clustering 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

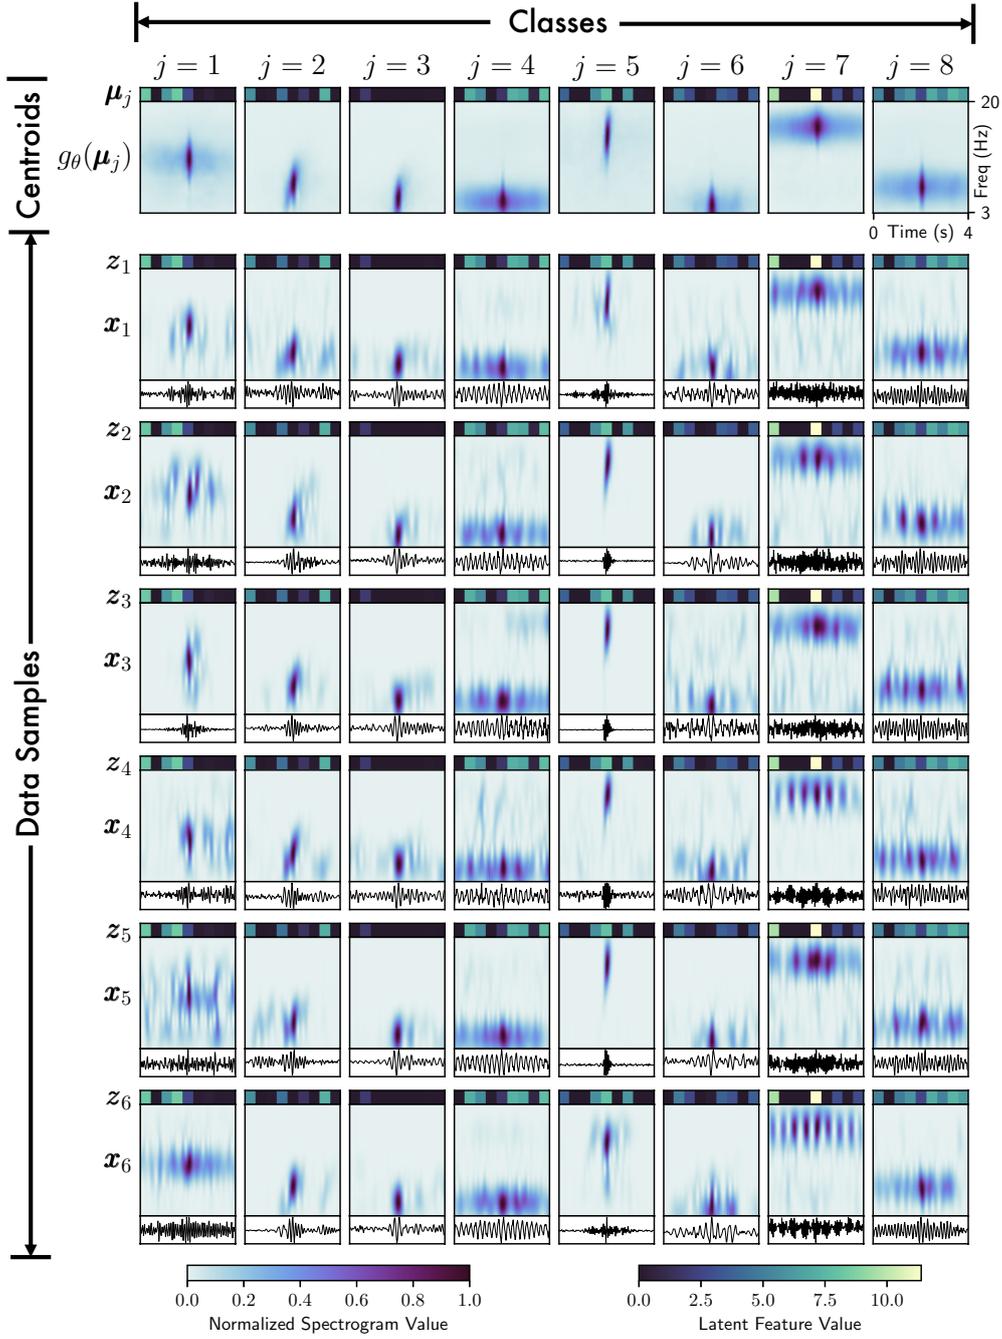


Figure 6. Within a given class j , the cluster centroids μ_j are similar to the latent feature data z_i . Though the centroids are not members of the data set, their reconstructions $g_\theta(\mu_j)$ exhibit similar characteristics to the spectrograms x_i assigned to each class. Seismograms plotted above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

358 a member of the data set, because the centroid represents the salient features of its class,
 359 its reconstruction is expected to resemble the spectrograms \mathbf{x}_i assigned to its class. Sub-
 360 sequent rows show the latent feature vectors \mathbf{z}_i , spectrograms \mathbf{x}_i , and associated seis-
 361 mograms of the data samples assigned to the respective classes.

362 For each class, latent feature vectors \mathbf{z}_i exhibit similar values to the class centroid
 363 $\boldsymbol{\mu}_j$, indicating that DEC has successfully grouped similar latent data samples into a class,
 364 and that the centroid is representative of the data in its class. The spectrograms in each
 365 class are likewise similar to each other and to the centroid reconstruction $g_{\theta}(\boldsymbol{\mu}_j)$, con-
 366 firming that the latent features embedded in the centroids are representative of the spec-
 367 trograms in the class. Finally, the similarity in the latent space and the time-frequency
 368 domain extends to the time domain, where seismograms in each class are similar to one
 369 another.

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

377 To determine to what extent DEC further improves clustering over GMM cluster-
 378 ing, t-SNE is used to visualize the 9-dimensional latent space in two dimensions (van der
 379 Maaten & Hinton, 2008). t-SNE can illuminate possible clusters within data in an un-
 380 supervised manner by displaying data in geometrically separated clusters. In Figure 7,
 381 t-SNE results of the latent feature space clustered with GMM show that the data are
 382 largely contiguous with few exceptions. Applying the labels assigned by GMM cluster-
 383 ing to the data points shows that, while there is some geometric separation between the
 384 clusters, the embedding is characterized by overlapping and dispersed class members,
 385 indicating poor separation in the latent space and potentially incorrect assignment of
 386 samples to classes. Contrast this with Figure 7, in which t-SNE results at the conclu-
 387 sion of DEC show both geometric separation as well as homogeneous class assignments.

388 While t-SNE offers an intuitively visual way to look for clusters in data, results are
 389 sometimes difficult to interpret and are impossible to reproduce exactly due to the in-

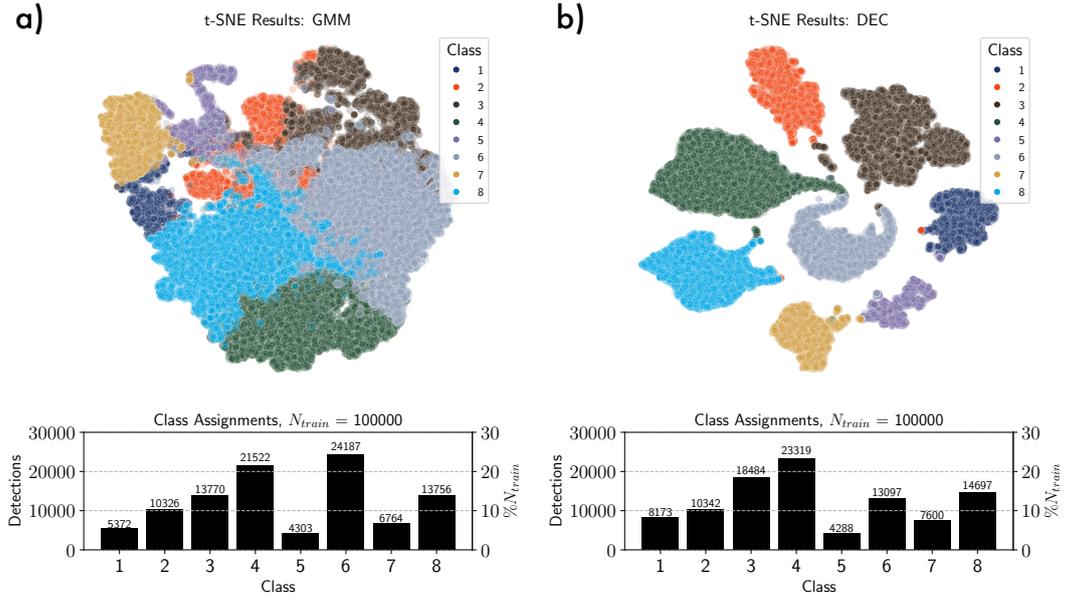


Figure 7. Visualization of the 9-dimensional latent data space is shown in two dimensions using the t-Student Stochastic Neighbor Embedding (t-SNE) plot. (a) Clustering and label assignment is performed with Gaussian mixture model clustering and exhibits limited separation within the data and overlapping 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.

390 herent randomness of the algorithm. Running t-SNE iteratively and with the same ran-
 391 dom seed can mitigate these limitations, but examination of the effects of DEC on the
 392 densities of the clusters provides a more concrete visualization. In Figure 8a, the cumu-
 393 lative distribution functions (CDF) for each class are shown as functions of distance to
 394 the centroid (equation (3)). For each class, the latent data move substantially closer to
 395 their assigned centroid, as evidenced by the decreased mean and variance of the CDF.
 396 Of interest to the ability for DEC to distinguish between clusters is the relation of each
 397 cluster to the others. In Figure 8b, the probability density functions (PDF) of all clus-
 398 ters are shown as functions of distance to each centroid. Before DEC, though GMM clus-
 399 tering results in the PDF of each class being closest to its centroid, there is significant
 400 overlap with other clusters, and the clusters themselves are not particularly dense. Af-
 401 ter DEC, the PDF of each class is closer to its centroid, denser, and farther removed from
 402 the other clusters. Thus, DEC effectively separates each cluster from the others, allow-
 403 ing for better distinction between clusters in the latent space.

404 The effects of DEC become readily apparent when the latent feature vectors are
 405 stacked and sorted according to their distance from each centroid, as shown in Figure 9.
 406 By sorting the latent space by sample index i such that $d_{i+1,j} > d_{i,j}$, cluster separa-
 407 tion can be visualized directly in the latent space. Before DEC, centroids are initialized
 408 with the GMM clustering algorithm without modification to the latent data. Closest to
 409 each class centroid, the latent feature vectors are similar in appearance to the centroid,
 410 but transition continuously to different patterns as the sorted index i increases. The con-
 411 trast with the latent space after DEC is stark: because DEC moves latent data assigned
 412 to a particular class closer to the centroid, the effect is that the latent feature vectors
 413 take on similar values, and therefore appearance, to the centroid. The result is that the
 414 latent space appears more sharply segmented after DEC, with the samples closest to the
 415 centroid of nearly uniform appearance to the centroid itself. For reference, the relative
 416 location of the other class centroids are marked with white vertical lines. Before DEC,
 417 the latent feature vectors belonging to the other classes are not readily apparent, whereas
 418 after DEC, most of the other centroid locations are associated with their distinctive la-
 419 tent feature vectors.

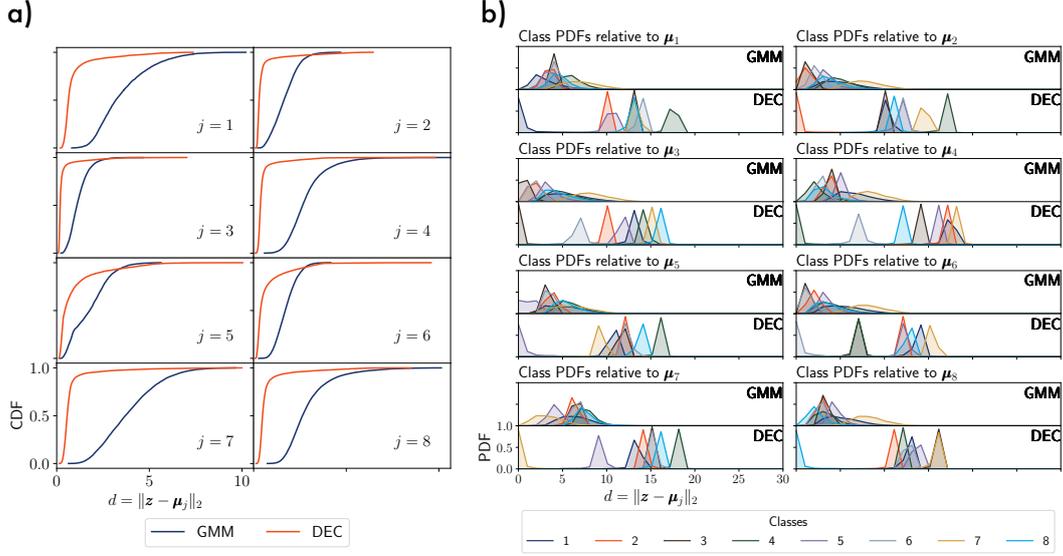


Figure 8. (a) Within each class, deep embedded clustering (DEC) reduces the mean distance of the assigned latent data to the centroid. The variance of the distance also decreases. As a result, the cumulative distribution functions shift to the left and have a steeper slope. (b) The effects of DEC are 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.

Latent space sorted by $d_{i,j} = \|\mathbf{z}_i - \boldsymbol{\mu}_j\|_2 \mid d_{i+1,j} > d_{i,j}$

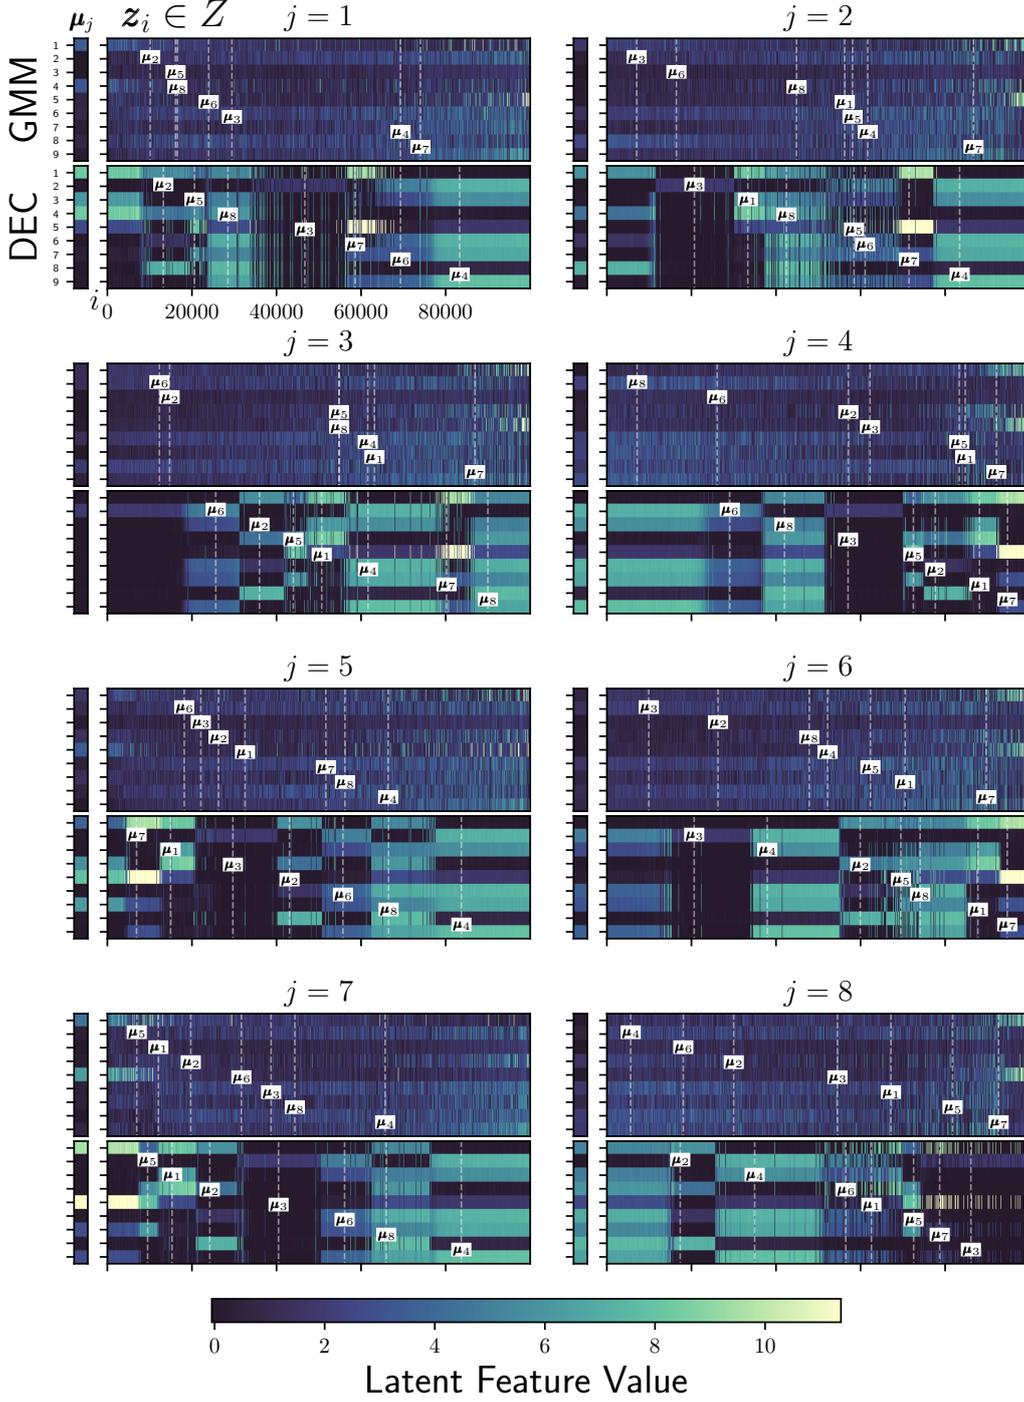


Figure 9. For each class j , latent data samples \mathbf{z}_i are shown stacked according to their distance $\|\mathbf{z}_i - \boldsymbol{\mu}_j\|$ from the centroid $\boldsymbol{\mu}_j$ (shown to the left). Distance of the other cluster centroids relative to the selected class j are indicated with vertical 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.

5.2 DEC Methodology Considerations

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

A second key strength of DEC is that clustering improvement and model optimization occur simultaneously. The outcome is denser, more separated clusters. This is a desirable effect in distance-based clustering, but it introduces a vulnerability: the success of the results may depend on the quality of the initial centroids. This challenge can be suitably mitigated by initializing the centroids with the GMM clustering algorithm, which accommodates clusters with unbalanced populations, elongated shapes, and differing variances. Additionally, Xie et al. (2016) demonstrated that DEC is robust against imbalances in class sample sizes, including size disparities up to a factor of 10. We assess that further refinement of centroid initialization is not necessary to achieve a workflow capable of performing satisfactory data exploration.

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

6 Discussion: Glaciological Implications

The full RIS array data set contains 427,798 seismic detections. A summary of the data set statistics and class characteristics (Table 3) 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 1, 4, and 7 have pronounced differences (more than 10%) between the number of detections occurring in the summers versus the winters, while differences for classes 2, 3, 5, and 8 are less pronounced (between 5% and 10%). Class 6 appears to have little difference (less than 5%) between austral summers and winters. Inter-annual comparisons for each season show that classes 1, 2, 5, and 8 experienced an increase in activity in the 2016 austral summer over the 2015 austral summer, with classes 5 and 8 exhibiting the largest changes. Classes 5 and 8 also increase by factors of three and two, respectively, in the 2016 austral winter over the 2015 austral winter. These trends can be investigated in more detail from 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 and detected seismicity associated with ocean gravity waves impacting the shelf front that cause fracturing (icequakes) and calving (Chen et al., 2019). Furthermore, seismicity at these stations during the 2016 austral summer is higher than the same period in 2015, indicative of the impact of El Niño on Antarctic ice shelf fronts (Nicolas et al., 2017). The remaining DR stations and stations RS01 through RS07 exhibit the opposite pattern: austral summers are relatively quiet, with increased detection frequencies in the austral winters. The four most seismically active stations were located near grounding zones: station RS09 (41,615 detections) on the eastern flank of Roosevelt Island; station RS11 (25,884 detections) on the Shirase Coast; station RS08 (18,655 detections) on the western flank of Roosevelt Island; and station RS17 (18,653 detections) on Steershead Ice Rise. All of these stations exhibited persistent seismicity throughout the two deployment years, with the exception of station RS17, which was offline for several weeks from August to September 2016. These

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	Std. dev.	Max.		
		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

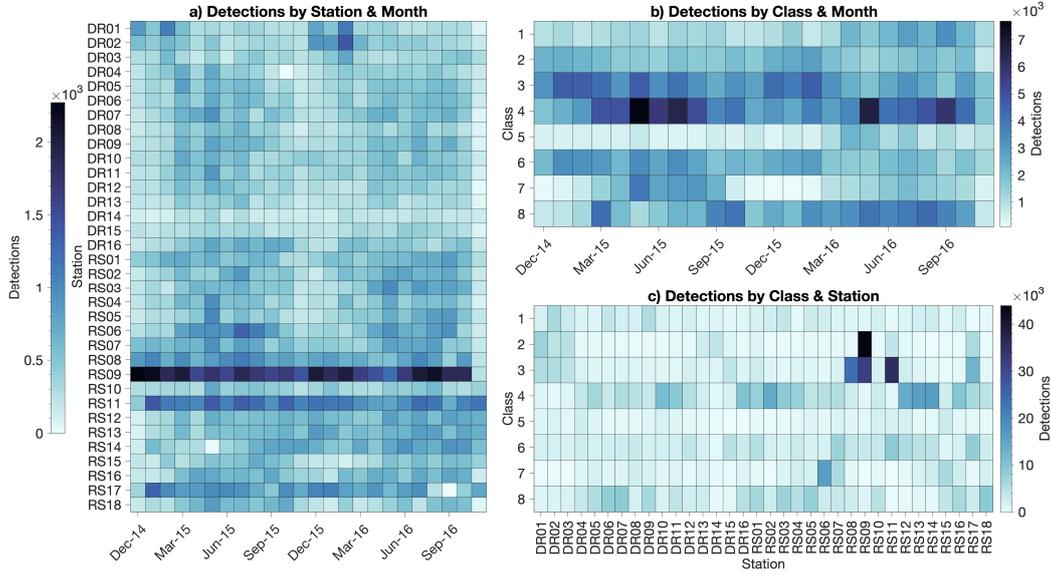


Figure 10. (a) The frequency of detections comprising the Ross Ice Shelf data set is shown by station and month. Deep embedded clustering (DEC) provides a further breakdown by (b) class and month for all stations, and (c) class and station.

484 stations are on either fully or partially grounded ice, suggesting that the seismicity re-
485 sults from interactions of basal ice with the solid earth.

486 Some classes of signal detections exhibit temporal patterns that are visible in Fig-
487 ure 10b. Classes 1 and 5 have increased detection frequencies in the austral winter of 2016,
488 while classes 2 and 3 have increased detections in the austral summers. Class 6 also ex-
489 hibits increased detections in the austral summers, though detections in the austral win-
490 ter of 2015 are also high. Meanwhile, classes 4, 7, and 8 exhibit low seismicity during the
491 austral summers. A further dimension to the analysis is shown in Figure 10c, which shows
492 the distribution of classes by station. The most frequently occurring class in the data
493 set is class 4, whose waveform and spectral variation resemble seismic tremor (Figure 6).
494 Class 4 signals occur across the array with peak activity in the austral winters when lo-
495 cal storms are more intense, suggesting meteorological forcing. Class 3 is the prominent
496 signal type at stations near grounding zones (RS08, RS09, RS11, and RS17), as is class
497 2 at RS09.

498 An important caveat for the detection statistics shown in Table 3 and Figure 10
499 arises from the physics governing seismic propagation. For a given amplitude, low fre-
500 quency seismic energy propagates farther than high frequency seismic energy. We thus
501 expect the seismometers in the RIS array to detect low-frequency signals originating far-
502 ther away than high-frequency signals. For example, from Figure 6, class 6 is similar to
503 classes 2 and 3, with the notable difference in that class 6 lacks energy above 8 Hz, and
504 has much smaller amplitude than classes 2 and 3. Thus, class 6 may be generated by a
505 similar source mechanism as classes 2 and 3, but has a longer propagation path.

506 Factoring in signal amplitude also affects the range at which seismic energy is de-
507 tected. From Table 3, class 7 has an average spectral peak at 16.1 Hz, the highest of the
508 classes, with a total of 32,276 detections, the second lowest of the classes. Similarly, class
509 5 has the second-highest average spectral peak at 14.8 Hz, with the least amount of de-
510 tectations among the classes. These two classes are nevertheless distinct from each other
511 in amplitude and waveform type: from Table 3, class 7 has a mean amplitude of 7.0 nm/s^2 ,
512 while class 5 has a mean amplitude of 408.5 nm/s^2 . From Figure 6, class 7 consists of
513 continuous signals, while class 5 signals are impulsive and likely result from fracturing.
514 Detection statistics are further affected by signal-to-noise ratios at the seismometers and

515 by limitations of the automated seismic event detector, such as the inability to separate
516 signals from different classes that are received nearly simultaneously.

517 Consideration should also be given to determining if classes are duplicates of the
518 same seismic source mechanism. Seismic surface waves in the ice undergo dispersion as
519 they propagate, which DEC may interpret as separate signal classes. Propagation mod-
520 eling can be used to calculate expected dispersion relations to confirm if this is the case.
521 Such distinctions could be useful in identifying common propagation paths or provid-
522 ing source range discrimination.

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

530 **6.1 Seasonal seismicity at the RIS front**

531 Approximately 2 km from the RIS front on Nascent Iceberg, station DR02 exhibits
532 a seasonal pattern of seismicity associated with changes in air temperature and sea ice
533 concentration in the Ross Sea. During the austral winter, sea ice coverage on the Ross
534 Sea reaches nearly 100%, damping ocean swell. During the austral summer, sea ice con-
535 centration (Figure 11a) decreases to approximately 25%, permitting ocean gravity waves
536 to directly impact the ice shelf front and cause iceberg calving. Additionally, warmer air
537 temperatures (Figure 11b) may promote calving with associated increased icequake ac-
538 tivity (Chen et al., 2019).

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

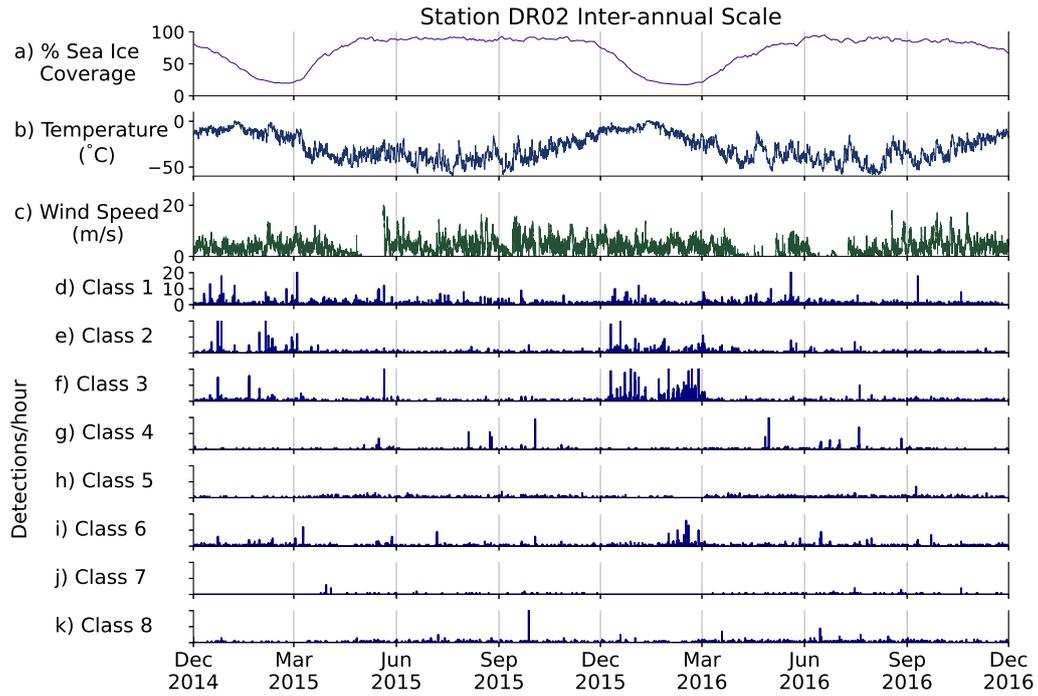


Figure 11. Two years of (a) sea ice coverage on the Ross Sea, (b) temperature and (c) wind speed at Gill automated weather station (approximately 223 km south of DR02), and (d-k) icequake detection statistics for each signal class. Classes 2, 3, and 6 exhibit increased seismicity during the austral summers.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

546 tween 10-21 January 2016 (Nicolas et al., 2017; Chaput et al., 2018), which can lead to
547 hydrofracture and contribute to ice shelf disintegration (Hubbard et al., 2016; Alley et
548 al., 2018).

549 Although class 1 has elevated activity during the summers, it maintains activity
550 throughout the winter months, suggesting that gravity wave activity is not the dominant
551 forcing. The persistence of class 1 signals, which often consist of impulse trains, suggests
552 they may be caused by icequakes resulting from the motion of the ice shelf itself, as the
553 ice flow velocity in the vicinity of station DR02 is among the highest observed on the
554 RIS (Klein et al., 2020). Classes 4, 5, 7, and 8 (Figure 11g,h,j,k) are more active dur-
555 ing the coldest periods of the year (April-September), suggesting that these signals may
556 be associated with extremely cold temperatures or strong wind events. Cold-weather en-
557 hanced seismicity occurs at a rift approximately 140 km south of the ice front (Olinger
558 et al., 2019). Alternatively, from Table 3, these classes are lower amplitude than those
559 most active during the austral summer, which suggests that these detections may be masked
560 by higher amplitude signals associated with classes 2 and 3. Across all classes, discrete
561 instances of high seismicity occur that do not correspond to environmental forcing. Such
562 instances may indicate the occurrence of fracturing ice (icequakes) or events associated
563 with crevasse expansion.

564 **6.2 Diurnal seismicity on Roosevelt Island**

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

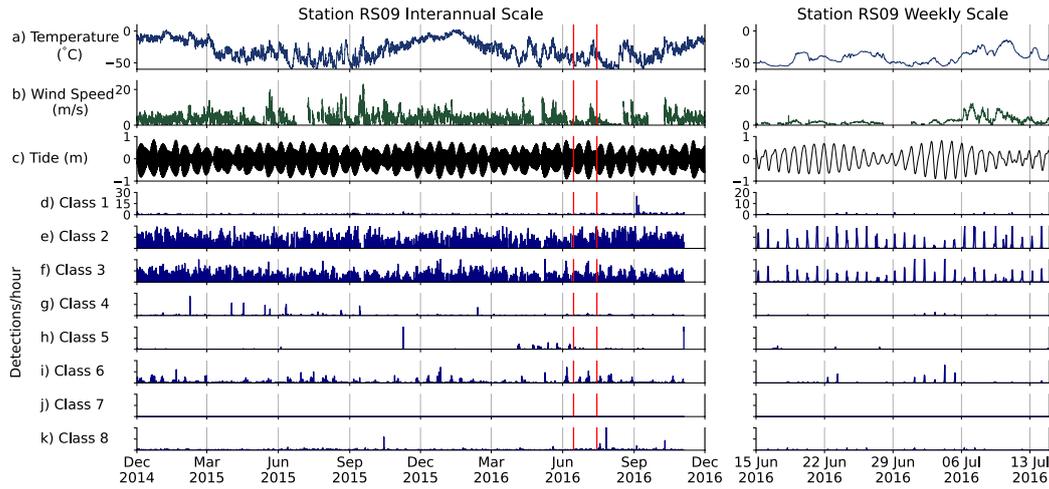


Figure 12. Two years of (a) temperature and (b) wind speed at Margaret automated weather station (MGT, approximately 122 km southwest of RS09, Figure 1), (c) model-derived tides calculated at station DR10, and (d-k) icequake detection statistics for each signal class. Inter-annual timescale is shown 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.

577 Of particular interest at station RS09 is evidence of seismicity associated with the
 578 diurnal tide (Figure 12). On an inter-annual timescale, class 6 exhibits a periodic mod-
 579 ulation of seismicity which tends to correlate with spring tides. Variability over fortnight
 580 tidal cycles is shown between 15 June 2016 and 15 July 2016. This weekly timescale shows
 581 that classes 2 and 3, the dominant signal classes, correlate with diurnal tides. Even some
 582 relatively non-active classes (1, 4, and 8) show signs of diurnal seismicity. These results
 583 are consistent with a previous study that found more than 95% of detections at RS09
 584 were from tidally induced swarms of icequakes that occur throughout the year (Cole, 2020).

585 Other stations located at grounding zones exhibit similar patterns of seismicity, though
 586 to a lesser extent than RS09. Station RS11, located east of RS09 on the Shirase Coast,
 587 exhibits patterns of seismicity similar to RS09. These similarities indicate that ice shelf
 588 seismicity at grounding zones is associated with similar ice shelf processes. RS08, on the
 589 western flank of Roosevelt Island, and RS17, at Steershead Ice Rise, also exhibit diur-
 590 nal seismicity, suggesting a dynamic diurnal process common to the grounding zones.
 591 These patterns of seismicity indicate that the interaction of the ice shelf with the solid

592 earth at grounding zones is modulated by tides. Among the four stations at grounding
593 zones, class 3 signals are the most common. With a mean peak frequency of 5.3 Hz and
594 a mean amplitude of 187 nm/s², these signals are among the strongest detected on the
595 array.

596 **7 Conclusions**

597 Application of deep embedded clustering (DEC) to the Ross Ice Shelf (RIS) array
598 data set identified eight classes of impulsive signals, with linkage of three of the classes
599 to tidal variability near grounding zones. Additionally, stations near the RIS front showed
600 increased icequake activity during the 2016 El Niño austral summer. The highest seis-
601 micity was observed at grounding zones, particularly along the eastern flank of Roosevelt
602 Island.

603 DEC is an effective way to explore large seismic data sets, particularly in its abil-
604 ity to identify dominant types of seismicity. The results provided by DEC, when con-
605 textualized with non-seismic environmental data, can assist in the identification or cor-
606 relation of seismic source mechanisms, as demonstrated with the RIS environmental data.
607 Additionally, DEC can be readily tailored to investigate different aspects of the same or
608 new data sets. Combined with its effectiveness at clustering seismic detections, this flex-
609 ibility suggests that DEC can be incorporated into existing seismic workflows in order
610 to speed up exploratory data analysis.

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

618 **Acknowledgments**

619 This work was funded by the Office of Naval Research through the National Defense Sci-
620 ence and Engineering Graduate Fellowship Program, and by National Science Founda-
621 tion (NSF) grant PLR 1246151, with support for Bromirski by NSF 1744856. Seismic

622 data from network XH (D. Wiens & Bromirski, 2014) were downloaded through IRIS
 623 Web Services (<https://service.iris.edu/irisws/>). Seismic data were processed us-
 624 ing Obspy software (Beyreuther et al., 2010). Figures were generated in MATLAB ([https://](https://www.mathworks.com)
 625 www.mathworks.com) and with Matplotlib (<https://matplotlib.org>). The DEC model
 626 was produced using PyTorch (<https://pytorch.org>). Antarctica elevation data, ground-
 627 ing line, and coast line were obtained from Bedmachine (Morlighem et al., 2017) and plot-
 628 ted using Antarctic Mapping Tools for MATLAB (Greene et al., 2017). Surface temper-
 629 atures were obtained from AMRC, SSEC, University of Wisconsin–Madison ([https://](https://amrc.ssec.wisc.edu)
 630 amrc.ssec.wisc.edu). Tide data were generated by the CATS2008 model (Padman et
 631 al., 2002). Ross Sea ice coverage was obtained from NASA NSIDC (Cavalieri et al., 1996,
 632 updated yearly). Code for this workflow is available at [https://github.com/NeptuneProjects/](https://github.com/NeptuneProjects/RISClusterPT)
 633 [RISClusterPT](https://github.com/NeptuneProjects/RISClusterPT).

634 References

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

- 654 fields observed on the Ross Ice Shelf. *Journal of Glaciology*, 1–17. doi:
 655 10.1017/jog.2020.83
- 656 Barcheck, C. G., Tulaczyk, S., Schwartz, S. Y., Walter, J. I., & Winberry, J. P.
 657 (2018, March). Implications of basal micro-earthquakes and tremor for ice
 658 stream mechanics: Stick-slip basal sliding and till erosion. *Earth and Planetary
 659 Science Letters*, 486, 54–60. doi: 10.1016/j.epsl.2017.12.046
- 660 Beaucé, E., Frank, W. B., & Romanenko, A. (2018, January). Fast Matched Fil-
 661 ter (FMF): An Efficient Seismic Matched-Filter Search for Both CPU and
 662 GPU Architectures. *Seismological Research Letters*, 89(1), 165–172. doi:
 663 10.1785/0220170181
- 664 Bellman, R. E. (1961). *Adaptive Control Processes: A Guided Tour*. Rand Corpora-
 665 tion.
- 666 Bergen, K. J., & Beroza, G. C. (2018, June). Detecting earthquakes over a seismic
 667 network using single-station similarity measures. *Geophysical Journal Interna-
 668 tional*, 213(3), 1984–1998. doi: 10.1093/gji/ggy100
- 669 Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann,
 670 J. (2010, May). ObsPy: A Python Toolbox for Seismology. *Seismological
 671 Research Letters*, 81(3), 530–533. doi: 10.1785/gssrl.81.3.530
- 672 Bianco, M. J., & Gerstoft, P. (2018, December). Travel Time Tomography With
 673 Adaptive Dictionaries. *IEEE Transactions on Computational Imaging*, 4(4),
 674 499–511. doi: 10.1109/TCI.2018.2862644
- 675 Bianco, M. J., Gerstoft, P., Olsen, K. B., & Lin, F.-C. (2019, December). High-
 676 resolution seismic tomography of Long Beach, CA using machine learning. *Sci-
 677 entific Reports*, 9(1), 14987. doi: 10.1038/s41598-019-50381-z
- 678 Bindschadler, R. A., King, M. A., Alley, R. B., Anandakrishnan, S., & Padman, L.
 679 (2003, August). Tidally Controlled Stick-Slip Discharge of a West Antarctic
 680 Ice Stream. *Science*, 301(5636), 1087–1089. doi: 10.1126/science.1087231
- 681 Bindschadler, R. A., Vornberger, P. L., King, M. A., & Padman, L. (2003). Tidally
 682 driven stick-slip motion in the mouth of Whillans Ice Stream, Antarctica. *An-
 683 nals of Glaciology*, 36, 263–272. doi: 10.3189/172756403781816284
- 684 Bishop, C. (2006). *Pattern Recognition and Machine Learning* (First ed.). Springer-
 685 Verlag New York.
- 686 Bromirski, P. D., Chen, Z., Stephen, R. A., Gerstoft, P., Arcas, D., Diez, A., . . . Ny-

- 687 blade, A. (2017, July). Tsunami and infragravity waves impacting Antarctic
688 ice shelves. *Journal of Geophysical Research: Oceans*, *122*(7), 5786–5801. doi:
689 10.1002/2017JC012913
- 690 Bromirski, P. D., Diez, A., Gerstoft, P., Stephen, R. A., Bolmer, T., Wiens, D. A.,
691 ... Nyblade, A. (2015, September). Ross ice shelf vibrations. *Geophysical*
692 *Research Letters*, *42*(18), 7589–7597. doi: 10.1002/2015GL065284
- 693 Bromirski, P. D., & Stephen, R. A. (2012). Response of the Ross Ice Shelf, Antarc-
694 tica, to ocean gravity-wave forcing. *Annals of Glaciology*, *53*(60), 163–172. doi:
695 10.3189/2012AoG60A058
- 696 Cavalieri, D. J., Parkinson, C. L., Gloersen, P., & Zwally, H. J. (1996, updated
697 yearly). *Sea Ice Concentrations from Nimbus-7 SMMR and DMSP SSM/I-*
698 *SSMIS Passive Microwave Data, Version 1*. Boulder, Colorado USA. NASA
699 National Snow and Ice Data Center Distributed Active Archive Center.
- 700 Chamarczuk, M., Nishitsuji, Y., Malinowski, M., & Draganov, D. (2020, January).
701 Unsupervised Learning Used in Automatic Detection and Classification of
702 Ambient-Noise Recordings from a Large-N Array. *Seismological Research*
703 *Letters*, *91*(1), 370–389. doi: 10.1785/0220190063
- 704 Chamberlain, C. J., Hopp, C. J., Boese, C. M., Warren-Smith, E., Chambers, D.,
705 Chu, S. X., ... Townend, J. (2018, January). EQcorrscan: Repeating and
706 Near-Repeating Earthquake Detection and Analysis in Python. *Seismological*
707 *Research Letters*, *89*(1), 173–181. doi: 10.1785/0220170151
- 708 Chaput, J., Aster, R. C., McGrath, D., Baker, M., Anthony, R. E., Gerstoft, P., ...
709 Stevens, L. A. (2018, October). Near-Surface Environmentally Forced Changes
710 in the Ross Ice Shelf Observed With Ambient Seismic Noise. *Geophysical*
711 *Research Letters*, *45*(20). doi: 10.1029/2018GL079665
- 712 Chazan, S. E., Gannot, S., & Goldberger, J. (2019, March). Deep Clustering Based
713 on a Mixture of Autoencoders. *arXiv:1812.06535 [cs, stat]*.
- 714 Chen, Z., Bromirski, P. D., Gerstoft, P., Stephen, R. A., Lee, W. S., Yun, S., ...
715 Nyblade, A. A. (2019, August). Ross Ice Shelf Icequakes Associated With
716 Ocean Gravity Wave Activity. *Geophysical Research Letters*, *46*(15), 8893–
717 8902. doi: 10.1029/2019GL084123
- 718 Chen, Z., Bromirski, P. D., Gerstoft, P., Stephen, R. A., Wiens, D. A., Aster, R. C.,
719 & Nyblade, A. A. (2018, October). Ocean-excited plate waves in the Ross and

- 720 Pine Island Glacier ice shelves. *Journal of Glaciology*, *64*(247), 730–744. doi:
721 10.1017/jog.2018.66
- 722 Cole, H. M. (2020). *Tidally Induced Seismicity at the Grounded Margins of the*
723 *Ross Ice Shelf, Antarctica* (Master’s Thesis). Colorado State University, Fort
724 Collins, Colorado.
- 725 De Angelis, H., & Skvarca, P. (2003, March). Glacier Surge After Ice Shelf Collapse.
726 *Science*, *299*(5612), 1560–1562. doi: 10.1126/science.1077987
- 727 Diez, A., Bromirski, P., Gerstoft, P., Stephen, R., Anthony, R., Aster, R., . . . Wiens,
728 D. (2016, May). Ice shelf structure derived from dispersion curve analysis
729 of ambient seismic noise, Ross Ice Shelf, Antarctica. *Geophysical Journal*
730 *International*, *205*(2), 785–795. doi: 10.1093/gji/ggw036
- 731 Dupont, T. K., & Alley, R. B. (2005). Assessment of the importance of ice-shelf but-
732 tressing to ice-sheet flow. *Geophysical Research Letters*, *32*(4). doi: 10.1029/
733 2004GL022024
- 734 Fürst, J. J., Durand, G., Gillet-Chaulet, F., Tavard, L., Rankl, M., Braun, M., &
735 Gagliardini, O. (2016, May). The safety band of Antarctic ice shelves. *Nature*
736 *Climate Change*, *6*(5), 479–482. doi: 10.1038/nclimate2912
- 737 Gibbons, S. J., & Ringdal, F. (2006, April). The detection of low magnitude seis-
738 mic events using array-based waveform correlation. *Geophysical Journal Inter-*
739 *national*, *165*(1), 149–166. doi: 10.1111/j.1365-246X.2006.02865.x
- 740 Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- 741 Greene, C. A., Gwyther, D. E., & Blankenship, D. D. (2017, July). Antarctic Map-
742 ping Tools for Matlab. *Computers & Geosciences*, *104*, 151–157. doi: 10.1016/
743 j.cageo.2016.08.003
- 744 Hartigan, J. A., & Wong, M. A. (1979). Algorithm AS 136: A K-Means Clustering
745 Algorithm. *Applied Statistics*, *28*(1), 100. doi: 10.2307/2346830
- 746 Hell, M. C., Cornelle, B. D., Gille, S. T., Miller, A. J., & Bromirski, P. D. (2019,
747 November). Identifying Ocean Swell Generation Events from Ross Ice Shelf
748 Seismic Data. *Journal of Atmospheric and Oceanic Technology*, *36*(11), 2171–
749 2189. doi: 10.1175/JTECH-D-19-0093.1
- 750 Hinton, G. E. (2006, July). Reducing the Dimensionality of Data with Neural Net-
751 works. *Science*, *313*(5786), 504–507. doi: 10.1126/science.1127647
- 752 Holtzman, B. K., Paté, A., Paisley, J., Waldhauser, F., & Repetto, D. (2018,

- 753 May). Machine learning reveals cyclic changes in seismic source spec-
 754 tra in Geysers geothermal field. *Science Advances*, 4(5), eaao2929. doi:
 755 10.1126/sciadv.aao2929
- 756 Hotovec-Ellis, A. J., & Jeffries, C. (2016, April). *Near Real-time Detection, Cluster-*
 757 *ing, and Analysis of Repeating Earthquakes: Application to Mount St. Helens*
 758 *and Redoubt Volcanoes* [Invited]. Reno, NV, USA.
- 759 Hubbard, B., Luckman, A., Ashmore, D. W., Bevan, S., Kulesa, B.,
 760 Kuipers Munneke, P., . . . Rutt, I. (2016, September). Massive subsurface
 761 ice formed by refreezing of ice-shelf melt ponds. *Nature Communications*, 7(1),
 762 11897. doi: 10.1038/ncomms11897
- 763 Johnson, C. W., Ben-Zion, Y., Meng, H., & Vernon, F. (2020, August). Identifying
 764 Different Classes of Seismic Noise Signals Using Unsupervised Learning. *Geo-*
 765 *physical Research Letters*, 47(15). doi: 10.1029/2020GL088353
- 766 Johnson, C. W., Meng, H., Vernon, F., & Ben-Zion, Y. (2019, August). Characteris-
 767 tics of Ground Motion Generated by Wind Interaction With Trees, Structures,
 768 and Other Surface Obstacles. *Journal of Geophysical Research: Solid Earth*,
 769 124(8), 8519–8539. doi: 10.1029/2018JB017151
- 770 Kingma, D. P., & Ba, J. (2017, January). Adam: A Method for Stochastic Opti-
 771 mization. *arXiv:1412.6980 [cs]*.
- 772 Klein, E., Mosbeux, C., Bromirski, P. D., Padman, L., Bock, Y., Springer, S. R., &
 773 Fricker, H. A. (2020, October). Annual cycle in flow of Ross Ice Shelf, Antarc-
 774 tica: Contribution of variable basal melting. *Journal of Glaciology*, 66(259),
 775 861–875. doi: 10.1017/jog.2020.61
- 776 Kong, Q., Trugman, D. T., Ross, Z. E., Bianco, M. J., Meade, B. J., & Gerstoft, P.
 777 (2019, January). Machine Learning in Seismology: Turning Data into Insights.
 778 *Seismological Research Letters*, 90(1), 3–14. doi: 10.1785/0220180259
- 779 Kullback, S., & Leibler, R. A. (1951, March). On Information and Sufficiency.
 780 *The Annals of Mathematical Statistics*, 22(1), 79–86. doi: 10.1214/aoms/
 781 1177729694
- 782 MacAyeal, D. R., Banwell, A. F., Okal, E. A., Lin, J., Willis, I. C., Goodsell, B.,
 783 & MacDonald, G. J. (2019, September). Diurnal seismicity cycle linked to
 784 subsurface melting on an ice shelf. *Annals of Glaciology*, 60(79), 137–157. doi:
 785 10.1017/aog.2018.29

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

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

- 852 Zwally, H. J. (2020, June). Pervasive ice sheet mass loss reflects compet-
 853 ing ocean and atmosphere processes. *Science*, *368*(6496), 1239–1242. doi:
 854 10.1126/science.aaz5845
- 855 Snover, D., Johnson, C. W., Bianco, M. J., & Gerstoft, P. (2020). Deep clustering
 856 to identify sources of urban seismic noise in Long Beach, CA. *Geophysical Re-*
 857 *search Letters*, *33*.
- 858 Steinbach, M., Ertöz, L., & Kumar, V. (2004). The Challenges of Clustering High
 859 Dimensional Data. In L. T. Wille (Ed.), *New Directions in Statistical Physics:*
 860 *Econophysics, Bioinformatics, and Pattern Recognition* (pp. 273–309). Berlin,
 861 Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/978-3-662-08968-2_16
- 862 Swindell, H., & Snell, N. (1977). *Station processor automatic signal detection sys-*
 863 *tem, phase I: Final Report, Station Processor Software Development, Texas*
 864 *Instruments Report No. ALEX (01)-FR-77-01*. (Tech. Rep.). Texas Instru-
 865 ments.
- 866 Telesca, L., & Chelidze, T. (2018, November). Visibility Graph Analysis of Seismic-
 867 ity around Enguri High Arch Dam, Caucasus. *Bulletin of the Seismological So-*
 868 *ciety of America*, *108*(5B), 3141–3147. doi: 10.1785/0120170370
- 869 Thoma, M., Jenkins, A., Holland, D., & Jacobs, S. (2008). Modelling circumpolar
 870 deep water intrusions on the amundsen sea continental shelf, antarctica. *Geo-*
 871 *physical Research Letters*, *35*(18). doi: 10.1029/2008GL034939
- 872 Tibshirani, R., Walther, G., & Hastie, T. (2001, May). Estimating the number
 873 of clusters in a data set via the gap statistic. *Journal of the Royal Statistical*
 874 *Society: Series B (Statistical Methodology)*, *63*(2), 411–423. doi: 10.1111/1467
 875 -9868.00293
- 876 Trugman, D. T., & Shearer, P. M. (2017, March). GrowClust: A Hierarchical Clus-
 877 tering Algorithm for Relative Earthquake Relocation, with Application to the
 878 Spanish Springs and Sheldon, Nevada, Earthquake Sequences. *Seismological*
 879 *Research Letters*, *88*(2A), 379–391. doi: 10.1785/0220160188
- 880 van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of*
 881 *Machine Learning Research*.
- 882 Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., & Manzagol, P.-A. (2010, Decem-
 883 ber). Stacked Denoising Autoencoders: Learning Useful Representations in a
 884 Deep Network with a Local Denoising Criterion. *Journal of Machine Learning*

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