

1

2

*[Geophysical Research Letters]*

3

Supporting Information for

4

**Characterization of Seismicity from Different Glacial Bed Types: Machine Learning  
Classification of Laboratory Stick-Slip Acoustic Emissions**

5

6

**S. Saltiel<sup>1</sup>, N. Groebner<sup>2</sup>, T. Sawi<sup>3</sup>, C. McCarthy<sup>3</sup>, B.K. Holtzman<sup>3</sup>**

7

<sup>1</sup> Nevada Seismological Laboratory, University of Nevada, Reno, NV, USA

8

<sup>2</sup> Strabo Analytics, Inc, New York, NY, USA.

9

<sup>3</sup> Lamont-Doherty Earth Observatory, The Earth Institute, Columbia University of New  
York, NY, USA

10

11

Corresponding author: Seth Saltiel (ssaltiel@unr.edu)

12

13

**14 Contents of this file**

15

Text S1 to S5

16

17

18

**Additional Supporting Information (Files uploaded separately)**

19

Captions for Movie S1

20

21

**Introduction**

22

Supporting information of additional details on experimental methods and materials, as  
well as data processing. Text S1 includes details of ice, rock, and till sources and  
preparation procedures; apparatus design; and experimental protocols. Text S2 includes  
data cleaning and normalization processing steps. Dataset S1 is a movie of the experiment  
and data stream in real-time, including audible stick-slips that are simultaneous with AEs  
and mechanical stress-drops being recorded. The datasets generated for this study are  
available on figshare.com at doi: 10.6084/m9.figshare.21257730, and Jupyter notebook for  
processing data is available at <https://github.com/StraboAI/IcesAEs>.

29

30

31 **Text S1: Experimental Details**

32 For this study we only used bulk ice samples, frozen slowly from deionized water in  
33 a slightly oversized die, and subsequently cut down to 50 x 50 x 100 mm with a microtome  
34 housed in a cold room (~ - 12 °C). The bulk freezing process results in large, non-uniform  
35 grain size compared to ‘standard ice,’ created using a narrow range of seed ice grain sizes  
36 [Cole 1979]. Saltiel et al., [2021] showed an insignificant frictional difference between the  
37 two types, so we employed bulk ice in this study. The simplified freezing process is much  
38 less time intensive and allows the ultrasonic transducers to be frozen directly into the ice  
39 sample (Figure S1), minimizing travel distance from the ice-bed interface and contact  
40 surfaces which can greatly diminish recorded acoustic amplitudes. The sliding surfaces  
41 were roughened with a no. 100 grit sandpaper using the same procedure as McCarthy et  
42 al., [2017], who determined a roughness average (Ra) of  $7 \pm 1 \mu\text{m}$  using a profilometer  
43 (Mitutoyo SF-210).

44



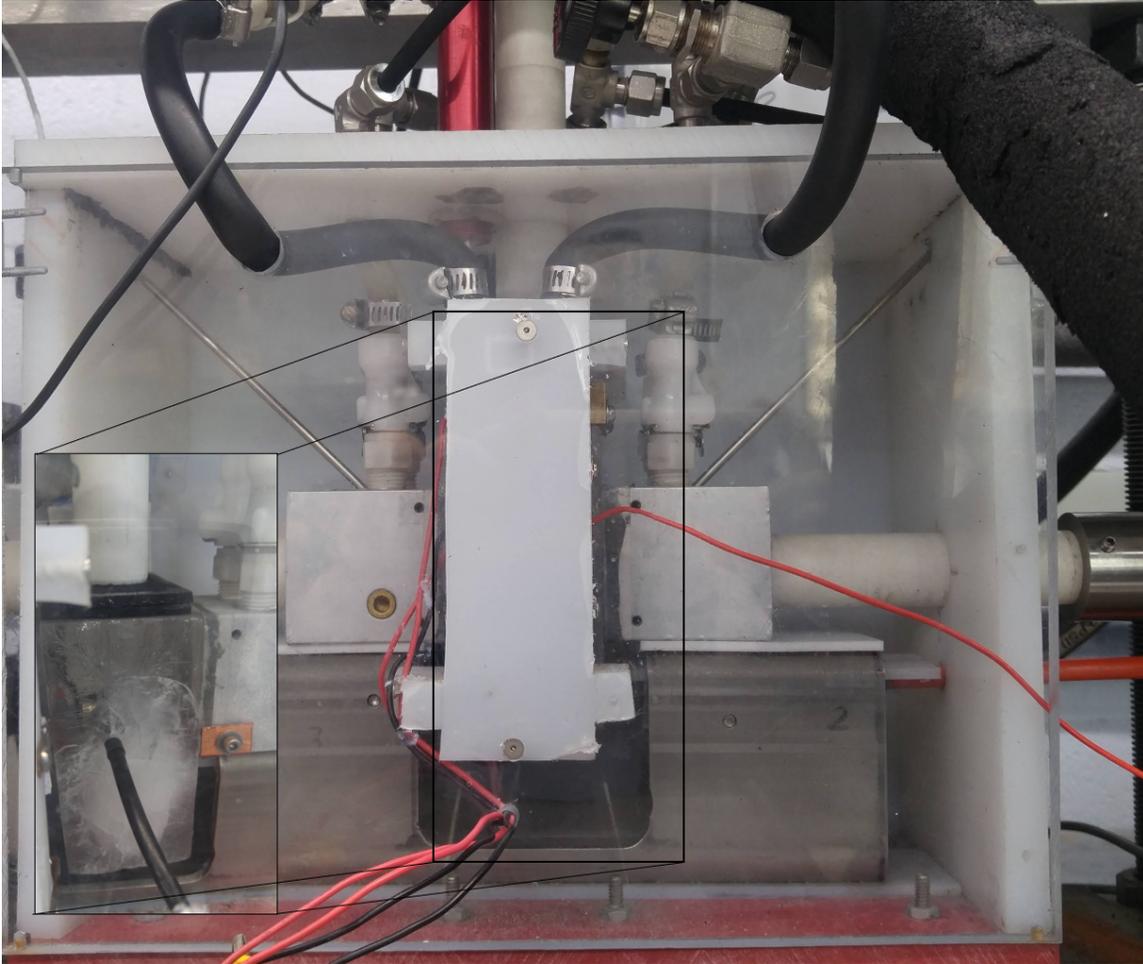
**Figure S1:** Bulk ice with an ultrasonic transducer (AE sensor) frozen into it. The bulk freezing process allows the suspension of the sensor in the deionized water during slow freezing. The sensor is oriented to face the sides of the block, where the ice-bed interface, source of AEs, will be when loaded into the apparatus.

65 As in Saltiel et al., [2021], we control temperature with Peltier thermoelectric coolers  
66 in front and behind the ice block, as well as circulation of chiller fluid through the side  
67 blocks where both temperature and flow rate of chiller fluid were actively controlled to  
68 reach the desired temperature. Resistance Temperature Detectors (RTDs) ported directly  
69 behind the till or rock monitor the temperature as close to the sliding interfaces as possible.  
70 Unlike in Saltiel et al., [2021], we preformed experiments with both stable and changing  
71 temperature to explore the effect on stick-slip instability, stress-drops, and resulting AEs.

72 Actively chilled aluminum side blocks were employed with either frozen till or rock  
73 attached to their ice-facing sides (Figures 1a, S2). All till experiments used a sample  
74 collected from the Matanuska glacier in south-central Alaska and were prepared using the  
75 same procedure described in Saltiel et al., [2021]. For rock beds, we employed Barre  
76 granite quarried from Barre Township, Vermont, that was cut into two 10 x 50 x 50 mm

77 slabs. A hole was drilled into the back side of the rock with the size and orientation of the  
78 side blocks' RTD port, to embed the RTD and measure the temperature directly behind the  
79 ice-rock interface. These slabs were then epoxied onto the aluminum side blocks and  
80 roughened using no. 100 grit sandpaper.

81



82

83 **Figure S2:** Photo of apparatus fully loaded. Since the peltier coolers cover the ice block, a  
84 photo without the cover is inset in the bottom left corner showing the central ice block at  
85 the end of an experiment, at the end of its full displacement.

86

87 All experiments were undertaken at  $\sim 50$  kPa of normal stress and a load point velocity  
88 of  $100 \mu\text{m/s}$  (just over 3 km/yr) for the entire displacement of 40 mm. This relatively high  
89 load point velocity was chosen because previous work has shown that stability decreases  
90 with slip velocity [Zoet et al., 2013, Saltiel et al., 2021]. Since the load point Linear  
91 Variable Differential Transformer (LVDT) only has 20 mm of stroke, the load point was  
92 stopped halfway through each experiment and then LVDT was reset to complete the rest  
93 of the experimental displacement. In this way, every experiment included a hold of about  
94 60 seconds during which the shear stress relaxed and then reloaded, usually resulting in the  
95 largest stress-drop and AE of each experiment.

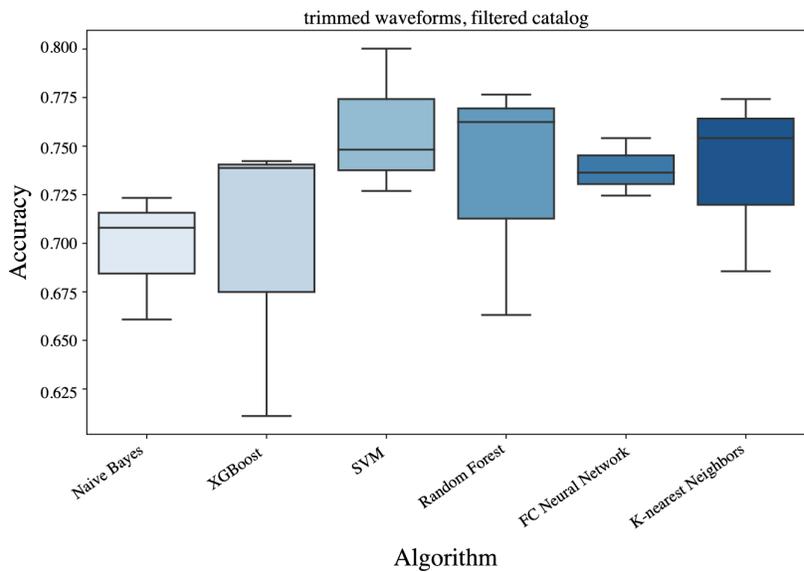
95

96 **Text S2: Data Cleaning, Trimming, and Normalization**

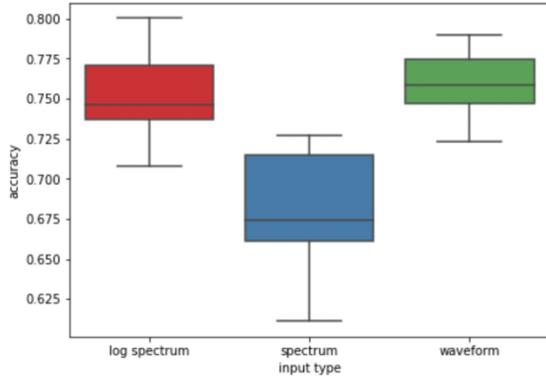
97 We implemented a data cleaning, trimming, and normalization approach based on  
98 that implemented by Nolte and Pyrak-Nolte [2022]. First, waveforms were trimmed to a  
99 total of 1200 samples, including 400 samples before the trigger point, giving a total window  
100 of 15 microseconds. Waveforms were then normalized by the sum of the squared  
101 amplitudes of the first 400 samples after the trigger, multiplied by a cosine taper. Zero and  
102 large amplitude waveforms were removed, defined as having a sum of the first 400  
103 normalized samples greater than 15. This threshold was found to give the best catalog of  
104 non-noise events without removing too many. 325 events were then removed that a have  
105 high amplitude low frequency noise component. Finally, the waveforms were realigned to  
106 the first maximum peak after the trigger, which refined alignment by a few samples in most  
107 cases. From this catalog of normalized, filtered, and aligned 1200-sample waveforms, we  
108 used a trial-and-error approach to determine how much of the pre- and post-trigger  
109 waveforms to use for training the models and found a total length of 150 samples, with 45  
110 before the trigger, was optimal. This subsample of the waveforms emphasizes the first  
111 arrivals of each AE, which are more dependent on source effects, while ignoring the coda,  
112 which depend more on path effects. Although, as we will show in the next section, the  
113 original, unprocessed catalog was able to produce as high prediction accuracies, the  
114 processed waveforms were clearer to interpret, the main point of this study.  
115

116 **Text S3: Results from Suite of ML Classification Algorithms**

117 We systematically tested of a suite of ML classification algorithms, the original,  
118 full catalog and that created by the trimming and cleaning processing steps described  
119 above, using both waveforms and spectra. Figures S3 – S6 show the distributions of  
120 prediction accuracies for each of these combinations of algorithms and catalogs.  
121

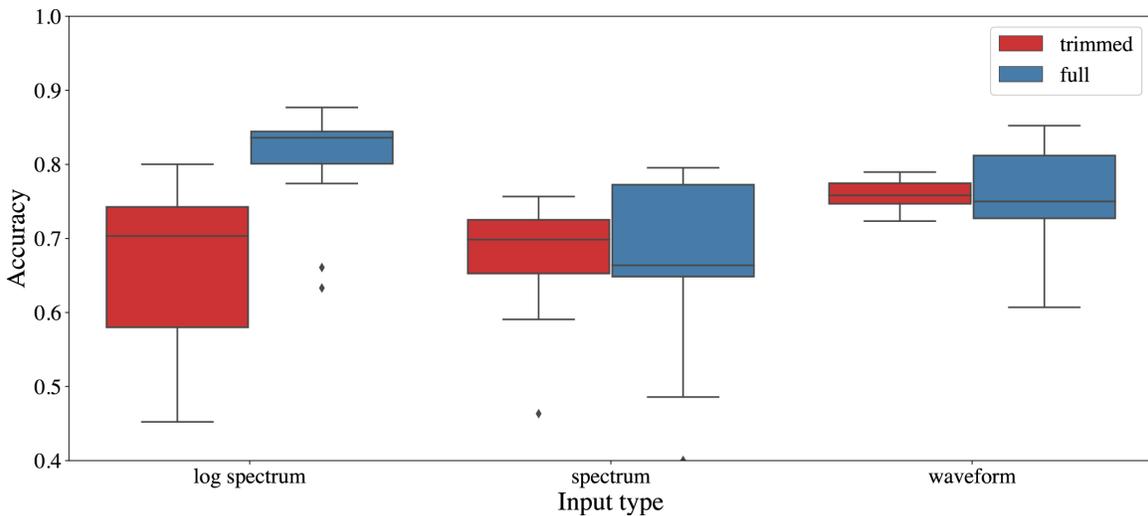


**Figure S3:** Whisker plot showing the distribution of prediction accuracies using the processed waveform catalog for each algorithm tested. Random Forest Classifier shows the highest mean accuracy of the all the algorithms which give a distribution, and, most importantly, provides feature importance for interpretation, so we focus on those results.



**Figure S4:** Whisker plot showing the distribution of prediction accuracies for each input data type tested, using the processed catalog. Waveforms show the tightest distribution and highest mean. Spectra are not very accurate, because the low frequency power dominates the spectral power and thus contains little information (see S4 below). But  $\log_{10}$  of the spectrum retain the high frequency information and accuracy can be as high as the predictions using waveforms.

148  
149



150  
151  
152  
153  
154  
155  
156  
157

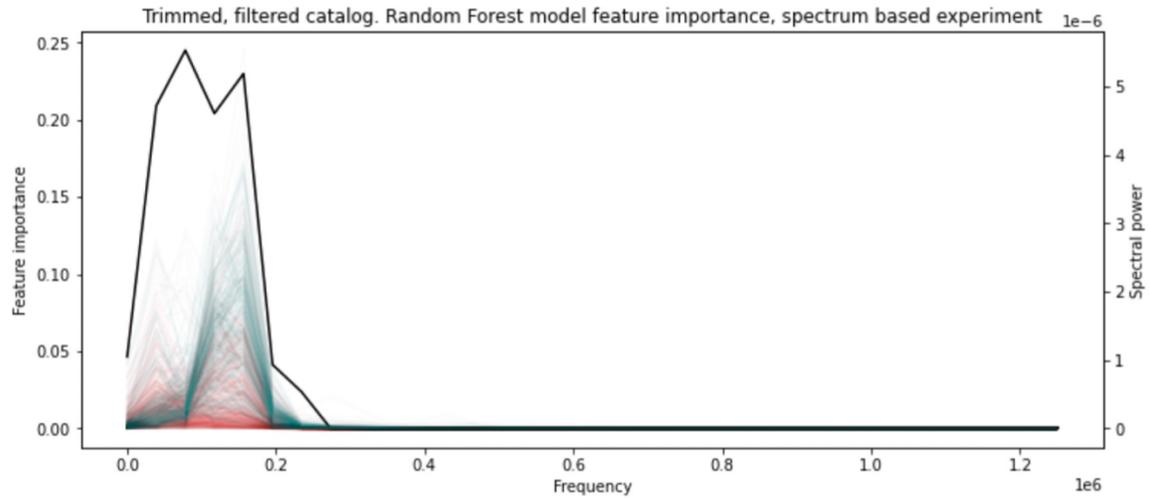
**Figure S5:** Whisker plot showing the distribution of prediction accuracies for the original, 'full', catalog of events vs. the processed, 'trimmed', catalog, using the processing steps described in text S2. Although the full catalog is able to give as good, or sometimes better predictions accuracies, which is not surprising since it contains more information, we focus our analysis on the processed, 'trimmed', catalog since the results are easier to interpret, the main focus of this study.

158 **Text S4: Predictions using Spectrum vs Log Spectrum**

159 We first undertook our analysis using spectrum, to test the predictive power of  
160 spectral information. But since the low frequency power dominates, using straight spectral  
161 power greatly diminishes the amount of data available (Figure S6a), and thus the  
162 predictions are relatively poor (Figure S4). By taking the log of the spectrum the higher  
163 frequency information is useful (Figure S6b) and predictions are more accurate.

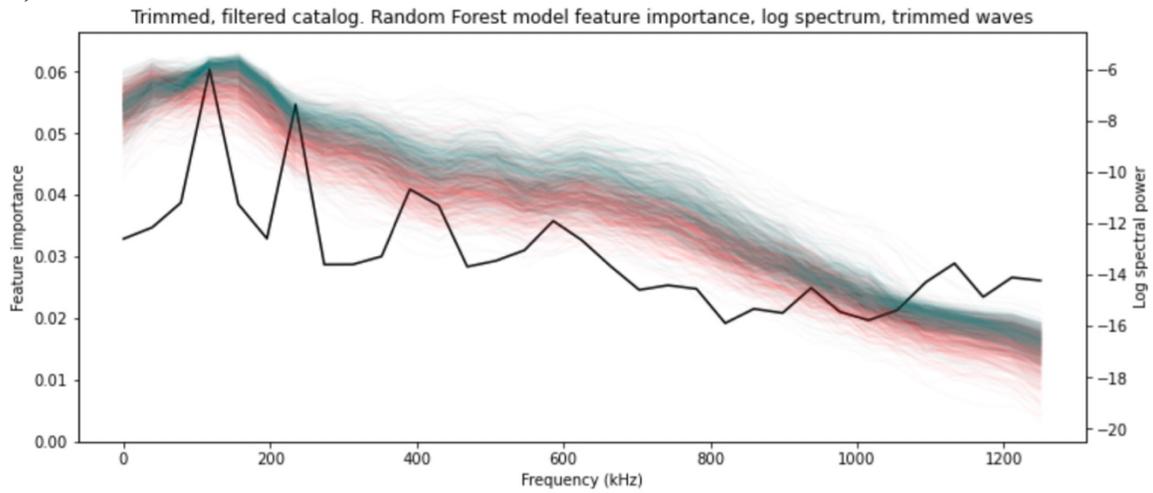
164  
165  
166  
167  
168  
169  
170

171 a)



172  
173

b)

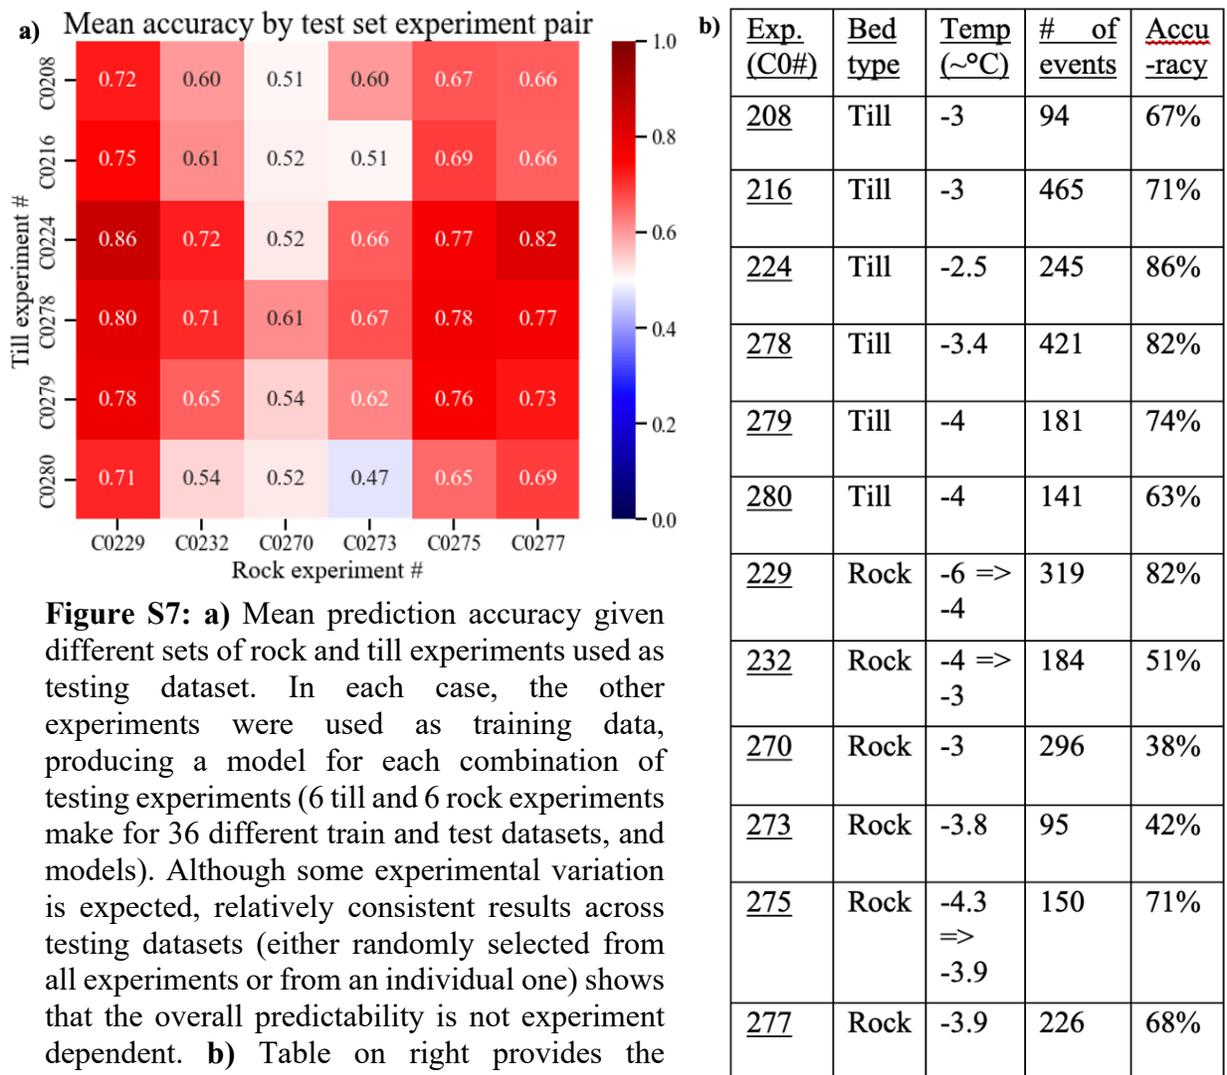


174  
175  
176  
177  
178

**Figure S6:** a) Spectrum from every till (teal) and rock (red) event, and the feature importance used to make Random Forest Classifier model predictions. Most spectral power is below 200 kHz, b) by taking the log spectrum, the higher frequency information is useable and prediction accuracy is improved.

179 **Text S5: Testing Experimental Differences**

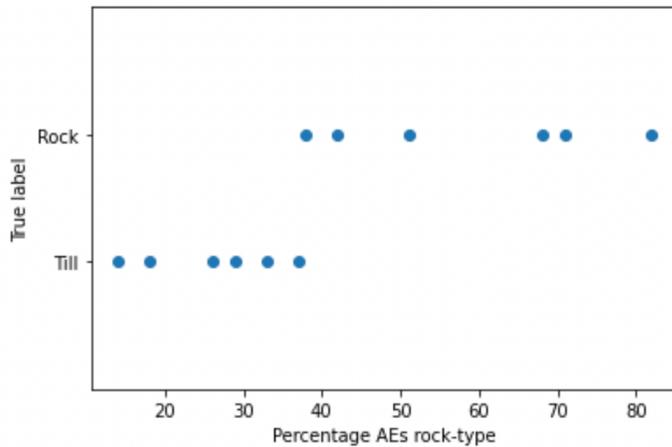
180 To ensure that the prediction is not based on some aspect of the waveform specific  
 181 to the ice sample or other uncontrolled aspect of the experiment and not the bed type which  
 182 we are testing for, we also tested each experiment independently, not allowing the  
 183 algorithm to train on data from the same experiment as the testing. We divide the data into  
 184 training and test sets based on experiment, i.e., for a given model training run the  
 185 waveforms from 5 till and 5 rock experiments are used for the training set, and the  
 186 remaining 1 till and 1 rock experiment are used for testing. By separating training and test  
 187 sets by experiment, any experiment-dependent features of the waveforms would be  
 188 irrelevant for classification. The prediction accuracy is summarized by a 6 till by 6 rock  
 189 experiments matrix, giving the accuracy for 36 models with each combination used as the  
 190 testing data (Figure S7).



191 **Figure S7: a)** Mean prediction accuracy given  
 192 different sets of rock and till experiments used as  
 193 testing dataset. In each case, the other  
 194 experiments were used as training data,  
 195 producing a model for each combination of  
 196 testing experiments (6 till and 6 rock experiments  
 197 make for 36 different train and test datasets, and  
 198 models). Although some experimental variation  
 199 is expected, relatively consistent results across  
 200 testing datasets (either randomly selected from  
 201 all experiments or from an individual one) shows  
 202 that the overall predictability is not experiment  
 203 dependent. **b)** Table on right provides the  
 204 temperature range, number of events, and  
 205 accuracy for each individual experiment.

206 This prediction accuracy calculates how often the model could correctly classify individual  
 207 waveforms as coming from till or rock beds, but we envision a tool whereby a collection

208 of seismic events recorded from a given location would be analyzed to determine the  
 209 probability it came from a till- or rock-bedded section of a glacier. So, the more relevant  
 210 accuracy is if a single experiment can be accurately predicted to be till or rock, and how  
 211 many events would be needed to make such a prediction accurate. Since its clear from  
 212 Figure 5 that there are overlapping ‘till-like’ rock events and visa-versa, the direct  
 213 prediction does not have to be used for the overall population prediction. For example, we  
 214 find that all the experiments can be correctly predicted if 37.5% ‘rock-like’ events, or  
 215 62.5% ‘till-like’ events, is used as the cut-off for overall prediction (Figure S8). Our data  
 216 shows a sharp cut off at these values, so it likely would not remain a perfect classifier with  
 217 more experiments, but it does suggest how predictions might be made given the  
 218 overlapping event populations.



**Figure S8:** Each experiments percentage of events predicted as rock, which we label as ‘rock-like’ events. The till and rock experiments perfectly separate if more than 37.5% of the events are predicted as rock.

230 Since there are rock experiments with more ‘till-like’ events than ‘rock-like’ events, it is  
 231 possible that the model is ‘defaulting’ to till since there are slightly more till than rock  
 232 events overall. We do not believe this is the case, given the significant overlap in the  
 233 characteristics of rock and till events (Figure 5). While the rock stress-drops have a tighter  
 234 distribution (Figure 4c), these stress drops do not follow a simple relationship with  
 235 recurrence interval, as would be expected with a single healing rate and as seen with the  
 236 till experiments (Figure 4d). Although there is not enough data to fully constrain, Figure  
 237 4d suggests that some rock experiments sit on the till healing relation (stress-drops of about  
 238 25 kPa per second of recurrence interval), while others have lower healing rates. This may  
 239 explain the imbalance in prediction accuracy, why there are more ‘till-like’ rock AEs than  
 240 ‘rock-like’ till AEs. Some experiments near the cut-off, such as 270, would be very difficult  
 241 to predict correctly. 270 is one of the rock experiments with a high healing rate (~22 kPa/s),  
 242 which might contribute to its having more ‘till-like’, misclassified events.

243 **Movie S1:** Movie of experiment and AE recording in real-time. Audible stick-slips and  
 244 mechanical stress-drop data (not shown) both simultaneously occur with the recorded AEs.  
 245 Some events appear to have two arrivals, probably one from each ice interface, since they  
 246 have different path lengths they arrive at the sensor at slightly different times even if they  
 247 occur at the same time. In these cases, the processing steps from text S2 remove the later  
 248 arrival.