

1 **A Little Data goes a Long Way: Automating Seismic Phase Arrival Picking at Nabro**  
2 **Volcano with Transfer Learning**

3

4 **Enter authors here: Sacha Lapins<sup>1</sup>, Berhe Goitom<sup>1</sup>, J-Michael Kendall<sup>2</sup>, Maximilian J.**  
5 **Werner<sup>1</sup>, Katharine V. Cashman<sup>1</sup> and James O. S. Hammond<sup>3</sup>**

6 <sup>1</sup>School of Earth Sciences, University of Bristol, UK.

7 <sup>2</sup>Department of Earth Sciences, University of Oxford, UK.

8 <sup>3</sup>Department of Earth and Planetary Sciences, Birkbeck, University of London, UK.

9

10 Corresponding author: Sacha Lapins (sacha.lapins@bristol.ac.uk)

11

12 **Key Points:**

- 13 • Transfer learning using existing model trained on California earthquake data produces  
14 effective new model for monitoring at Nabro volcano
- 15 • Nabro transfer learning model shows improved S-wave picking resulting in smaller  
16 location errors than even manual phase picks
- 17 • Changing task from classification to segmentation results in more efficient model  
18 processing 14 months of data from 7 stations in 4 hours

19

## 20 **Abstract**

21  
22 Supervised deep learning models have become a popular choice for seismic phase arrival  
23 detection. However, they tend to perform poorly on out-of-distribution data and require large  
24 training sets to aid generalization and prevent overfitting. This presents an issue when using these  
25 models in new monitoring settings. In this work, we develop a deep learning model for automating  
26 phase arrival detection at Nabro volcano using a limited amount of training data (2498 event  
27 waveforms recorded over 35 days) through a process known as transfer learning. We use the  
28 feature extraction layers of an existing, extensively-trained seismic phase picking model to form  
29 the base of a new all-convolutional model, which we call U-GPD. We demonstrate that transfer  
30 learning reduces overfitting and model error relative to training the same model from scratch,  
31 particularly for small training sets (e.g., 500 waveforms). The new U-GPD model achieves greater  
32 classification accuracy and smaller arrival time residuals than off-the-shelf applications of two  
33 existing, extensively-trained baseline models for a test set of 800 event and noise waveforms from  
34 Nabro volcano. When applied to 14 months of continuous Nabro data, the new U-GPD model  
35 detects 31,387 events with at least four P-wave arrivals and one S-wave arrival, which is more  
36 than the original base model (26,808 events) and our existing manual catalogue (2,926 events),  
37 with smaller location errors. The new model is also more efficient when applied as a sliding  
38 window, processing 14 months of data from 7 stations in less than 4 hours on a single GPU.

39

## 40 **Plain Language Summary**

41

42 Seismic monitoring increasingly relies on automated signal processing as the rate of data  
43 acquisition grows. Supervised deep learning models have proven to be effective for detecting and  
44 characterizing seismic events, but training such highly parameterized models generally requires  
45 large amounts of manually labelled data. Once trained, however, these models extract general  
46 seismic waveform features that can be used to train new models with more limited training data.  
47 In this work, we use the generalized knowledge of seismic data from a model trained on millions  
48 of earthquakes in California to train a new model for detecting volcanic earthquakes at Nabro  
49 volcano, Eritrea, a recently active and, prior to its 2011 eruption, poorly monitored volcano. Using

50 a small training set of waveforms, the new model more accurately detects phase arrivals and noise  
51 than off-the-shelf applications of two baseline models. The new model is efficient, processing 14  
52 months of data in less than 4 hours. It is also effective, detecting more volcanic events and showing  
53 improved levels of S-wave arrival picking. The result is smaller event location errors than even  
54 our manual picks. This level of efficiency and consistency highlights the role that machine learning  
55 can play in volcano-seismic monitoring.

56

## 57 **1 Introduction**

58

59 Seismic monitoring plays a fundamental part in mitigating hazards at volcanoes. During  
60 periods of unrest, thousands of earthquakes can occur each day, producing a diverse range of  
61 seismic signals that reflect a multitude of interlinked volcanic processes (e.g., migrating fluids,  
62 fault movement, explosions, rockfalls). These earthquakes are generally recorded by broadband  
63 seismometers, which are highly sensitive to ground motion across a wide range of frequencies and  
64 record signals at high sample rates (typically 100 times or more per second). This level of detail,  
65 however, comes at the cost of generating vast amounts of data. Many seismic networks utilize tens  
66 or even hundreds of seismometers at a given time (e.g., Hansen & Schmandt, 2015), making real-  
67 time manual inspection of these time series practically infeasible. Previous seismic deployments  
68 have also generated extensive legacy datasets that can offer insights into historical volcanic activity  
69 and opportunities to further our understanding of volcanic processes. The main challenge is  
70 therefore to identify and characterize volcanic earthquakes in a robust and timely manner so as to  
71 provide vital clues regarding the state of a volcano and the likelihood or impact of an eruption or  
72 hazard, as well as be able to accurately and efficiently process large existing datasets for further  
73 analysis within a reasonable timeframe.

74

75 Identifying earthquake phase arrivals, particularly the initial primary (P-) and  
76 secondary/shear (S-) wave arrivals, forms the basis of most seismic processing tasks (e.g.,  
77 determining locations, magnitudes and source parameters). Manually identifying these phase  
78 arrivals yields greater accuracy and estimates of arrival time uncertainty than automated  
79 approaches but is extremely time-consuming. Alternatively, most automated approaches are orders

80 of magnitude quicker but typically require clear phase arrivals, existing ‘templates’ of previously  
81 catalogued earthquakes (e.g., Gibbons & Ringdal, 2006; Lengliné et al., 2016; Shelly et al., 2007),  
82 or pre-processing / feature extraction steps calibrated for a small set of earthquake characteristics  
83 (e.g., trigger algorithms based on the ratio of short-term average to long-term average signal  
84 amplitude, STA/LTA; Withers et al., 1998). A challenge for application to volcanology is that  
85 volcanic earthquakes can exhibit widely varying time-frequency characteristics, often with low  
86 amplitudes or obscured phase arrivals, and new phases of unrest can produce previously unseen  
87 seismic signals that differ from existing earthquake templates. Furthermore, methods based on  
88 existing seismic catalogues are unsuitable for new seismic deployments where a catalogue of  
89 events has not been collected.

90  
91 A recently successful approach for seismic phase arrival detection is the use of supervised  
92 deep learning models (e.g., Dokht et al., 2019; Mousavi et al., 2019; Ross et al., 2018; Woollam  
93 et al., 2019; Zhu & Beroza, 2019). These methods are based on convolutional neural networks  
94 (CNN), a variant of classical neural networks that employ convolution operations, as opposed to  
95 matrix multiplication, in at least part of the model. These operations are employed in ‘hidden’  
96 convolutional layers that allow the network to learn a large set of filters to extract useful features  
97 from the input data and map them to a desired output (e.g., to identify phase arrivals in earthquake  
98 waveforms; Fig 1). Typically, multiple convolutional layers are applied in succession and in  
99 combination with other operations, such as non-linear ‘activation’, down-sampling and  
100 normalization, to extract complex patterns from the data using a hierarchy of simpler filter kernels.  
101 These extracted features can then be fed into a standard fully-connected neural network or other  
102 machine learning architecture for classification, segmentation, regression, clustering or inference  
103 (e.g., Mousavi et al., 2019; Ross et al., 2018; van den Ende & Ampuero, 2020). As such, the  
104 ‘convolutional’ part of CNNs act as the model’s feature extraction system. With each successive  
105 convolutional layer, the extracted features move from lower-level, general signal features  
106 (resembling, for example, long/short period wavelets in seismological waveform models; Fig 1A  
107 inset) to more task specific, high-level features (Yosinski et al., 2014). The final ‘classification’  
108 layers of the model map these features to the desired output and can be considered the most task  
109 specific part of the model, empirically tuned to the distribution of the training data (Yosinski et  
110 al., 2014). Such an approach gives these models a strong advantage over traditional algorithms

111 that require considerable manual intervention or rely on a small set of manually determined  
112 characteristics and simple threshold criteria. Supervised deep learning models, however, require  
113 substantial amounts of labelled data during training to generalize to out-of-sample data (the amount  
114 dependent on various factors, such as network architecture, number of network parameters and  
115 training hyperparameters; e.g., D’souza et al., 2020; He et al., 2019; Sun et al., 2017). They can  
116 also suffer significant loss in performance when faced with data that differs in source or  
117 distribution from their training data (e.g., Barbedo, 2018; Zech et al., 2018). This requirement can  
118 place the traditional paradigm of supervised learning (i.e., using a large amount of hand-labelled  
119 data to train a single model for a desired domain or problem) out of reach for many real-world  
120 applications.

121  
122         Transfer learning is based on the idea of knowledge transfer from one task to another (Pan  
123 & Yang, 2010; Zhuang et al., 2020) and can be a powerful tool when we do not have sufficient  
124 labelled data to train a reliable model from scratch. At its simplest, the first  $n$  convolutional layers  
125 and their weights from the feature extraction part of an existing model are copied to the first  $n$   
126 layers of a new model for a related or similar task, with the remaining layers either re-initialized  
127 with randomized weights or replaced (e.g., Razavian et al., 2014; Yosinski et al., 2014). These  
128 tasks need not be near-identical or even superficially related, as long as low-level data  
129 characteristics are shared between tasks (e.g., Efremova et al., 2019; Tran et al., 2020; Zamir et  
130 al., 2018). The intuition is that generalized knowledge of data structure and properties from one  
131 model trained with abundant labelled data (or ‘big data’) can guide a learning algorithm towards a  
132 good solution for a new task with far more limited, or even no, labelled data.

133  
134         In this paper, we evaluate the utility of inductive transfer learning (i.e., when labelled data  
135 are available for both the source and target tasks) for small seismic training sets and produce a  
136 deep learning model that accurately and robustly picks phase arrivals from a deployment at Nabro  
137 volcano in Eritrea, a region with little or no prior seismic monitoring. We leverage the knowledge  
138 acquired from training a model on millions of seismic waveforms recorded by the Southern  
139 California Seismic Network (SCSN), hereby referred to as the GPD model (Generalized seismic  
140 Phase Detection; Ross et al., 2018), and apply it to seismograms from Nabro volcano in Eritrea,  
141 for which we have limited hand-labelled data (manual phase arrival picks) from the first couple of

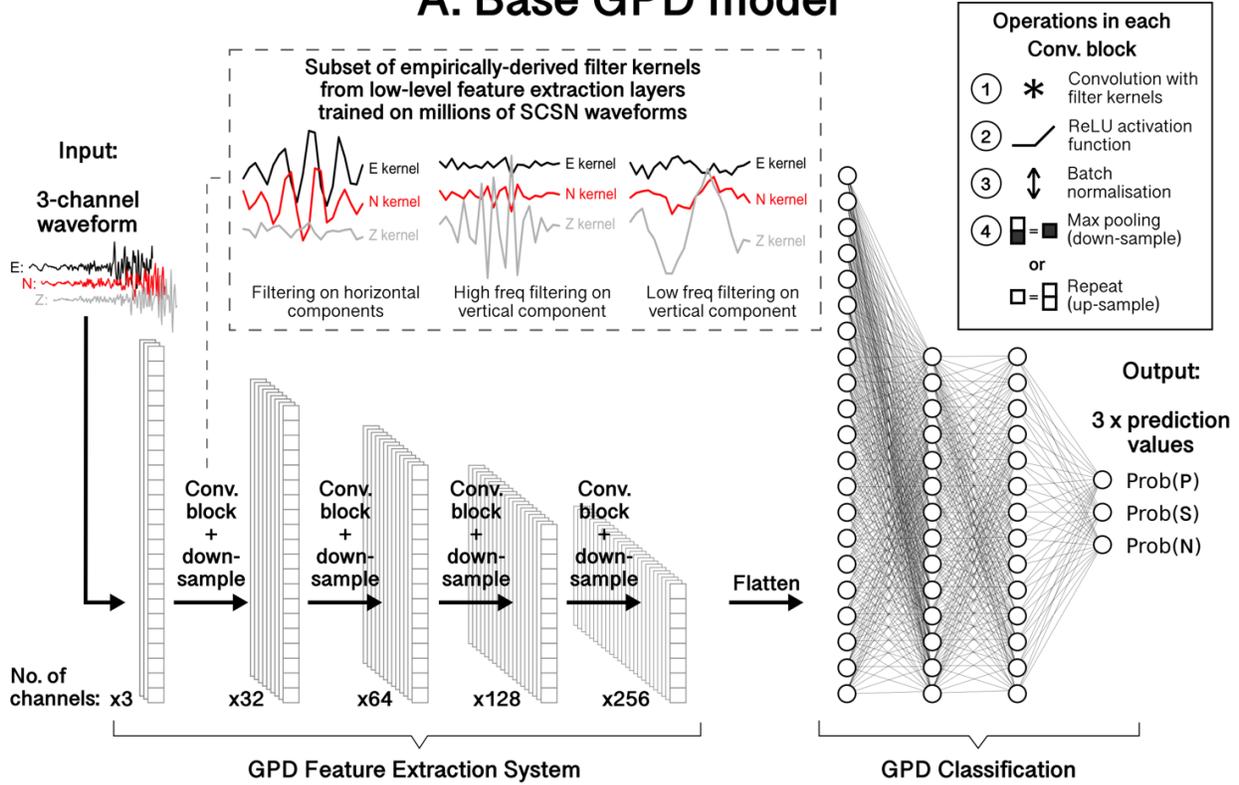
142 months of a 14-month seismic deployment (Goitom, 2017; Hamlyn et al., 2014). The new model  
143 task differs from the original GPD model task in that it is modified from one of *classification*  
144 (assigning a single class label *P-wave*, *S-wave* or *noise* to an entire 4-second waveform; Fig 1A)  
145 to one of *segmentation* (assigning a class label *P-wave*, *S-wave* or *noise* to *each datapoint* within  
146 that 4-second waveform; Fig 1B). We achieve this by replacing the fully-connected uppermost  
147 layers of the original GPD model with further convolutional layers, creating an all-convolutional  
148 model commonly referred to as a U-Net (Ronneberger et al., 2015). We refer to this specific model  
149 design as the U-GPD model, utilizing GPD model weights within a U-Net architecture. The new  
150 data from Nabro volcano also exhibit differences in instrument calibration and sample rates from  
151 the original GPD model training data, as well as differing waveform characteristics between  
152 tectonic and volcanic event types (Lahr et al., 1994; Lapins et al., 2020; McNutt & Roman, 2015).  
153

154 In the following section, we introduce transfer learning and recent applications in  
155 seismological deep learning. In Sections 3 and 4, we present our proposed transfer learning  
156 method, U-GPD model architecture and seismic data recorded at Nabro volcano. In Section 5, we  
157 present a series of model comparisons. We first use common training metrics to demonstrate that  
158 transfer learning reduces overfitting and model error, particularly for very small training sets (<  
159 1000 waveforms), when compared with a model reinitialized with randomized weights before  
160 training (i.e., trained from scratch with no transfer learning). We then apply these new models to  
161 a test dataset of known P-/S-wave arrivals and sections of noise and compare performance with  
162 off-the-shelf applications of the base GPD model and another extensively-trained phase-picking  
163 model, PhaseNet (Zhu & Beroza, 2019). We find that the U-GPD transfer learning model yields  
164 improved phase arrival identification, particularly for S-waves, and false detection rate at Nabro  
165 volcano. Altering the model task from classification to segmentation also improves pick time  
166 residuals over the base GPD model for these test data. Finally, we apply both our new U-GPD  
167 transfer learning model and the original base GPD model to the full 14-month seismic deployment  
168 at Nabro volcano through a sliding window approach. The new U-GPD model identifies more  
169 useable S-wave arrivals than the base GPD model, yielding smaller subsequent location errors than  
170 even our manual analyst's phase arrival picks. The new model also runs an order of magnitude  
171 faster, processing 14 months of data from 7 broadband seismometers in less than 4 hours on a  
172 single GPU. Our findings indicate that transfer learning can be extremely useful for volcano

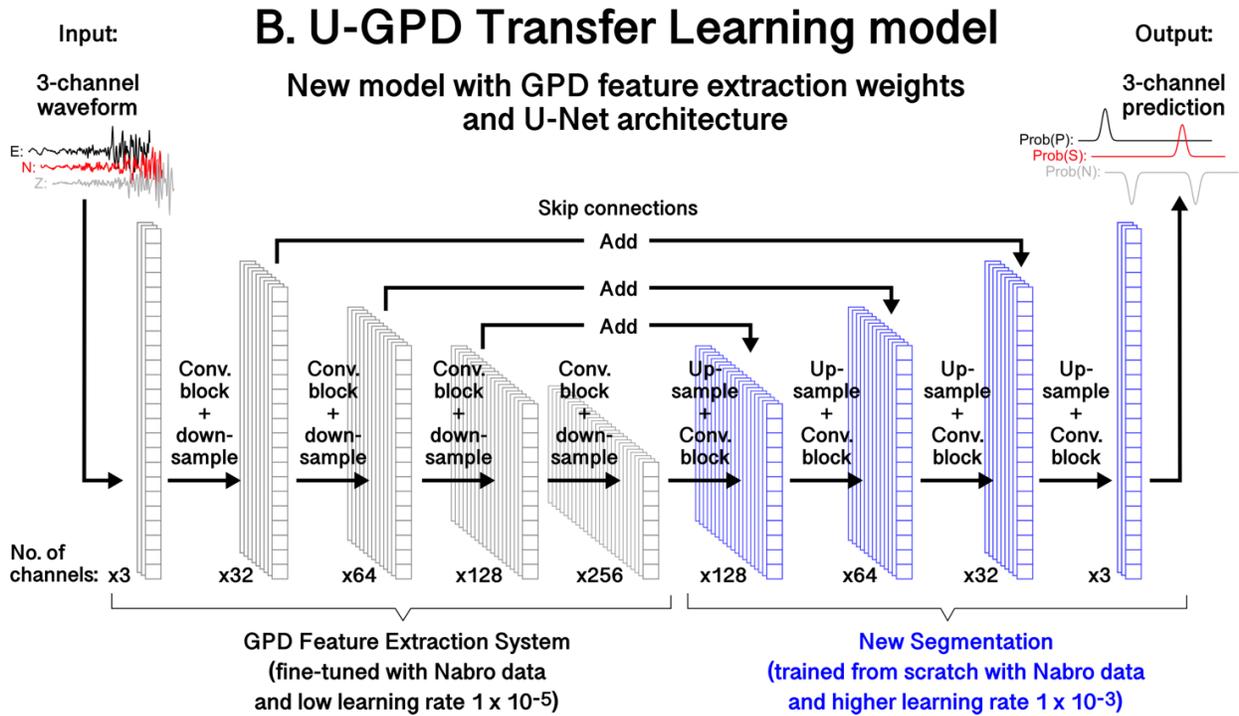
173 seismic monitoring, even with limited computing resources and data. We conclude this paper with  
174 a discussion of our findings, methodology and practical considerations of transfer learning in  
175 Section 6. All data and code used throughout this paper are made fully and publicly available (see  
176 *Data Availability Statement*).

177

## A. Base GPD model



## B. U-GPD Transfer Learning model



178

179

180 **Figure 1. a)** Model architecture for Generalized seismic Phase Detection (GPD) CNN model (Ross  
181 et al., 2018). Model can be considered as two parts: a feature extraction system (convolutional  
182 layers) and classification part (fully connected layers). GPD model outputs 3 x prediction values  
183 (probability of P, S or noise) for an entire 400-sample 3-component waveform (i.e., output  
184 dimensions: 1 x 3). Examples of filter kernels (dashed line inset) from lowest convolutional layer  
185 that extract generalized seismic waveform features determined through model training on  
186 extensive SCSN dataset. These indicate that the GPD model has learnt to extract different features  
187 from vertical and horizontal components. **b)** Proposed transfer learning model architecture (“U-  
188 GPD”). GPD model feature extraction system is copied to new model and fine-tuned with new  
189 Nabro data and low learning rate. Low learning rate ensures that useful features are not  
190 ‘unlearned’. New convolutional layers replace the GPD classification layers and are trained using  
191 new Nabro data and higher learning rate. Model outputs 3 x prediction values for each datapoint  
192 in 400-sample 3-component waveform (i.e., output dimensions: 400 x 3).

193

## 194 **2 Transfer Learning**

195

196 There are many approaches to transfer learning (see Pan & Yang, 2010; Zhuang et al., 2020  
197 for comprehensive surveys), including using ‘off-the-shelf’ feature extraction systems from  
198 existing state-of-the-art CNNs (e.g., Maqsood et al., 2019; Razavian et al., 2014), learning domain-  
199 invariant or global representations across multiple tasks (e.g., Glorot et al., 2011; Li et al., 2014;  
200 Tzeng et al., 2015; Zhuang et al., 2015), applying pre-processing steps to make input data  
201 representations more similar between datasets (e.g., Daumé, 2007; Sun et al., 2016) and the use of  
202 domain-adversarial models (e.g., Ganin et al., 2016). Here we employ the first of these approaches  
203 for P- and S-wave arrival time picking at Nabro volcano, utilizing pre-trained filters from an  
204 existing, extensively trained CNN model (the GPD model; Ross et al., 2018) to train a new model  
205 with different output dimension and task type (see *Section 3.1, U-GPD Model Architecture*). Other  
206 seismological studies that have employed transfer learning in this way have used pre-trained filters  
207 from models designed for non-seismological tasks, such as image recognition. For example, filters  
208 trained to recognize photographic images or handwritten characters have been used to detect  
209 earthquakes and classify volcano-seismic event types from spectrograms (Huot et al., 2018; Lara  
210 et al., 2020; Titos et al., 2020) and interpret seismic facies (Dramsch & Lühje, 2018).

211  
212  
213  
214  
215  
216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241

Some studies have chosen to fine-tune entire seismic deep learning models, essentially updating the models with new data (or equivalently ‘pre-training’ the models with larger datasets, depending on perspective). El Zini et al. (2020) pre-train an autoencoder with abundant unlabeled data to learn compressed data representations of 2D seismic images. These model weights then serve as a starting point for a model that segments seismic images, with weights fine-tuned using limited labelled training data. This approach was shown to outperform the transfer of weights from image recognition models and training a model from scratch. Bueno et al. (2020) fine-tune a Bayesian neural network (BNN) to improve classification of volcano-seismic event characteristics between datasets and time periods. They show that this approach increases model accuracy and reduces epistemic uncertainty when applied to new volcanic systems or phases of activity. With a similar aim but different approach to the work of this paper, Chai et al. (2020) utilize pre-trained weights from another existing phase arrival detection model, PhaseNet (Zhu & Beroza, 2019), to pick phase arrivals from hydraulic fracturing experiments. They use the entirety of the PhaseNet model and its pre-trained weights as a starting point for training and then fine-tune all model weights equally using just 3,500 seismograms. They present improved results over the original PhaseNet model, which was trained using 700,000 seismograms of regional Californian seismicity, when applied to higher sample rate data (2 kHz) from a very different setting (i.e., hydraulic fracturing). Whilst these studies show that fine-tuning entire models can be an effective strategy, poor hyperparameter choices (model learning rate, number of training epochs, etc.) can inadvertently retrain the model (also known as ‘catastrophic forgetting’; e.g., Kirkpatrick et al., 2017) or lead to settling on a non-global minimum within the parameter space, reopening the potential for overfitting when the number of model parameters is large and the training dataset is small (El Zini et al., 2020; Yosinski et al., 2014). The work in this paper differs from that of Chai et al. (2020) in that only the weights from the feature extraction part (i.e., the first ‘half’) of the GPD model are transferred to our new U-GPD model. These weights are fine-tuned using a much lower learning rate (weight update step size) to retain useful learned knowledge from the original model but optimize cohesion with the rest of the new model, which is redesigned to reduce the total number of trainable parameters, among other optimizations (see Section 3.1, Model Architecture), and initialized with randomized weights (Fig 1).

## 242 **3 Proposed Model**

### 243 **3.1 U-GPD Model Architecture**

244

245 As outlined briefly above, we utilize pre-trained parameters from the convolutional layers  
246 of the GPD model as a starting point for our U-GPD transfer learning model. The original GPD  
247 model was trained using 4.5 million hand-labelled seismograms (1.5 million of each class *P*, *S* and  
248 *noise*) recorded by the Southern California Seismic Network (SCSN) between the years 2000 and  
249 2017. These training data were all 400-sample (4 sec) 3-component waveforms, high-pass filtered  
250 above 2 Hz and (re)sampled at 100 Hz. All events had epicentral distances less than 100 km and  
251 magnitudes between -0.81 and 5.7 *M* (various magnitude scales). The GPD model was chosen as  
252 a base for our transfer learning model as these data characteristics are comparable to those  
253 observed and recorded by volcano observatories. Furthermore, the short input length of 4 seconds  
254 (400 samples at 100 Hz sample frequency) means there is less chance of erroneously labelling or  
255 missing relatively small magnitude or overlapping phase arrivals. Finally, the GPD model's  
256 'sequential' architecture, with each layer being solely connected to the layers directly before and  
257 after, also means the model is more interpretable and makes it easier to isolate its feature extraction  
258 system.

259

260 During model training, we fine-tune these pre-trained parameters using a very small  
261 learning rate ( $1 \times 10^{-5}$ ), rather than keep them fixed (e.g., Yosinski et al., 2014). The aim of this  
262 fine-tuning step is to modify any highly specific features from the source domain (particularly in  
263 the higher-level feature extraction layers) and overcome optimization difficulties arising from  
264 splitting the GPD convolutional layers from co-adapted classification layers (Yosinski et al., 2014)  
265 without unlearning the important generalized waveform features we wish to exploit. We then  
266 replace the GPD model's fully-connected layers (i.e., the task-specific classification part of the  
267 model) with further convolutional layers and up-sampling operations, combined with ReLU  
268 activation function (Nair & Hinton, 2010) and batch normalization (Ioffe & Szegedy, 2015), to  
269 produce a model output with the same dimensions as model input (400 samples x 3 channels; Fig  
270 1). Each of the three output channels represents the model's prediction (or 'probability') of a P-  
271 wave arrival, S-wave arrival or neither (hereby referred to as *noise*), respectively, at each datapoint

272 in the waveform. This all-convolutional approach has been adopted by other phase arrival picking  
273 models (e.g., Woollam et al., 2019; Zhu & Beroza, 2019) and has several distinct advantages when  
274 applied to seismic phase arrival detection: i) it provides less ambiguous labelling of phase arrivals  
275 when compared to the original GPD model's approach of assigning a single class prediction ( $P$ ,  $S$   
276 or *noise*) to an entire 400-sample 3-channel waveform; ii) convolutional layers tend to have fewer  
277 parameters than fully connected neural network layers so less training data is required to avoid  
278 overfitting; iii) by producing a model with input and output traces of same dimension, we require  
279 less overlap when applied as a rolling window method, producing a model that runs orders of  
280 magnitude faster on continuous sections of data. These new convolutional layers are initialized  
281 with completely randomized weights and trained with a higher learning rate ( $1 \times 10^{-3}$ ) than the  
282 pre-trained GPD weights. We use dilated filter kernels in these new convolutional layers (e.g., van  
283 den Oord et al., 2016; Yu & Koltun, 2016) to increase the size of the model's receptive field (or  
284 'field of view') and aggregate multi-scale context. Finally, the new layers are subjected to spatial  
285 dropout (Tompson et al., 2015), where 30% of the feature maps (output of filter operations) in  
286 each convolutional layer are effectively dropped (set to zero) at the start of each training epoch.  
287 This step promotes independence between the features the model extracts and prevents overfitting  
288 (Tompson et al., 2015). Precise details of U-GPD model dimensions and hyperparameters are  
289 provided in *Supplementary Materials*.

290

291 The overall network architecture outlined above is sometimes referred to as a U-Net  
292 (Ronneberger et al., 2015). With each step through the network, the input data are progressively  
293 downsampled with an increasing number of features extracted, creating a contracting network path  
294 that is forced to sacrifice detail and learn a more compressed, general representation of the input  
295 waveform to discriminate between classes ( $P$ ,  $S$  or *noise*). The model then follows a symmetrically  
296 expanding path, where the data are progressively upsampled and the number of features reduced,  
297 to regain precise temporal or spatial detail and return an output with equal dimension to the model  
298 input (Ronneberger et al., 2015). Skip connections (addition operators), which act as direct, one-  
299 way pathways between layers in the contracting and expanding sides of the model (Fig 1B), are  
300 used to retain precise waveform details that may be lost through this contraction/expansion process  
301 and have been shown to greatly improve the likelihood of model parameters settling on the global  
302 minimum during training (Li et al., 2017).

303

### 304 **3.2 Phase Arrival Labels and Model Hyperparameters**

305

306 Each 3-component waveform in our training dataset has a corresponding 3-channel ‘mask’  
307 that provides a ground truth label (*P*, *S* or *noise*) for each waveform datapoint. During training, the  
308 model aims to minimize the difference between its predictions and these ground truth labels.  
309 Labels are presented as binary values (0’s or 1’s), with P-wave arrivals indicated by a +/- 0.14 sec  
310 boxcar function, centered on the manually picked P-wave arrival time, and S-wave arrivals  
311 indicated by +/- 0.19 sec boxcar function, also centered on the manually picked S-wave arrival  
312 time. These boxcar widths provide a good balance between phase arrival detection rate and arrival  
313 time precision and compensate for human error in the ground truth labels. Previous studies have  
314 used Gaussian-style probability masks, with values ranging between 0 and 1, for labelling phase  
315 arrivals (e.g., Woollam et al., 2019; Zhu & Beroza, 2019). We find that label accuracy on our test  
316 data (e.g., Fig 5) and event location error distributions from the full deployment (e.g., Fig 8C &  
317 D) are near-identical when using either approach but training with boxcar masks produces a model  
318 that detects ~ 10% more events when run over continuous data.

319

320 As with the original GPD model, our new U-GPD model was trained using a categorical  
321 cross entropy loss function (see Supplementary Materials) and the Adam optimization algorithm  
322 (Kingma & Ba, 2014). The model weights that produced lowest loss value on the validation dataset  
323 during training were selected as our final model weights. Other loss functions that address the  
324 imbalance between arrival and noise labels (as the majority of labels in any given waveform are  
325 not a phase arrival), such as a focal loss function that effectively adds weighting parameters to  
326 cross entropy loss (Lin et al., 2017), were trialed but yielded no improvement in model  
327 performance.

328

## 329 **4 Data**

330

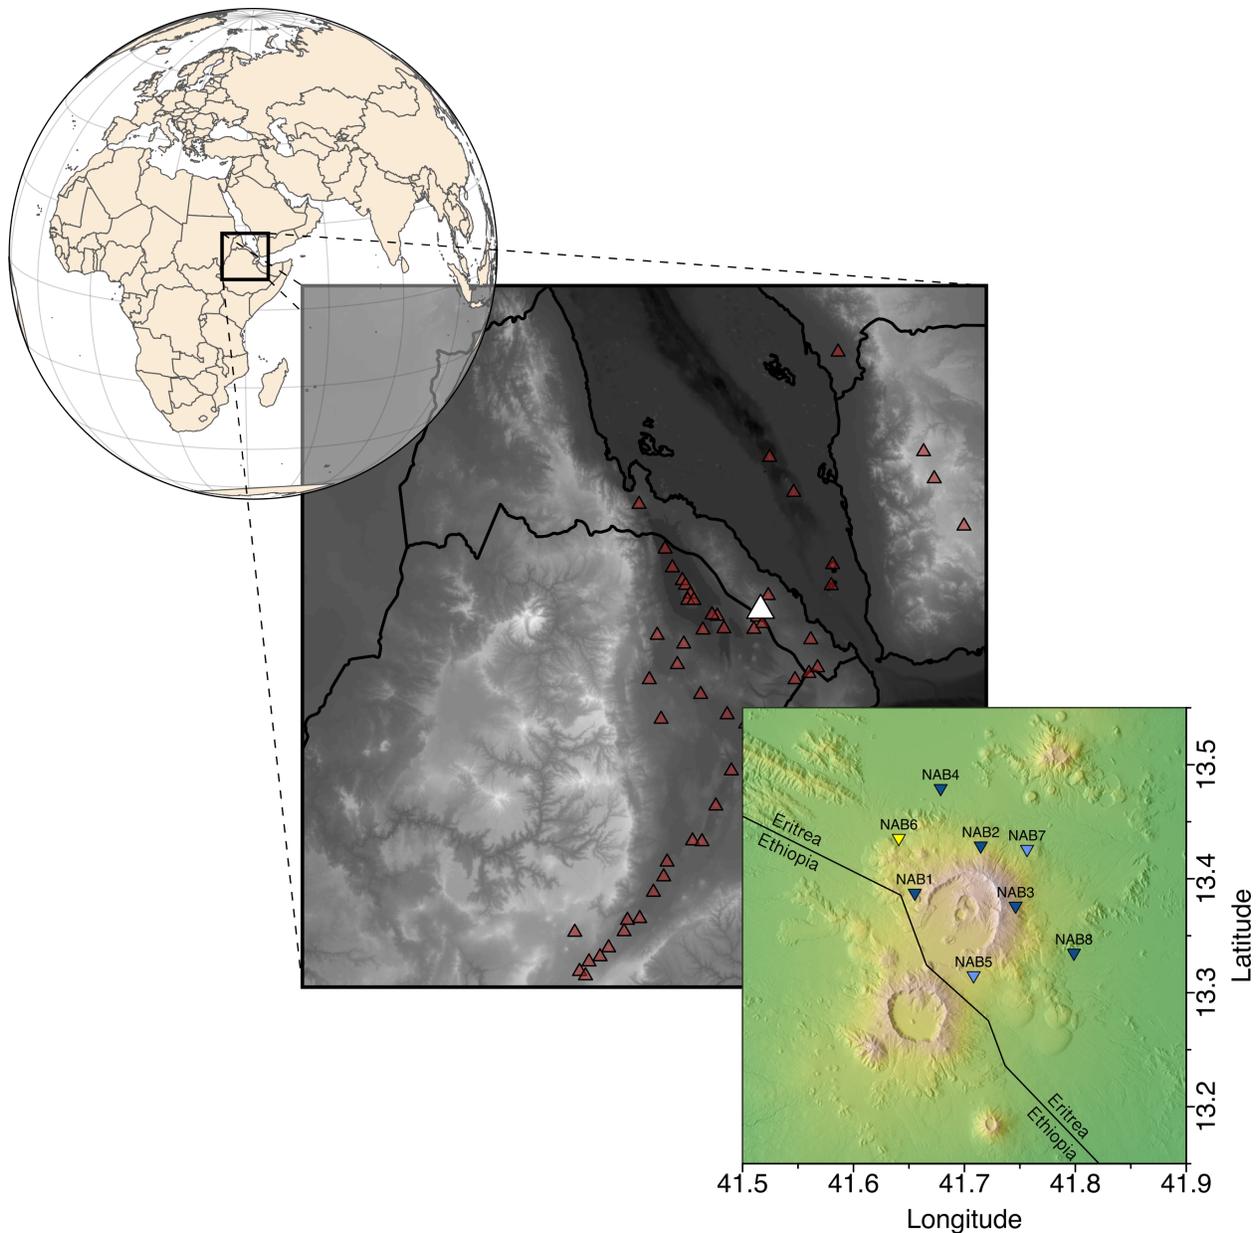
331 Nabro volcano is one of two calderas that form the Bidu Volcanic Massif on the Eritrea-  
332 Ethiopia international border (Fig 2). Located in the Afar region at the northern end of the Main

333 Ethiopian Rift, it erupted unexpectedly for the first time in recorded history on 12<sup>th</sup> June, 2011,  
334 disrupting continental aviation and initiating a significant humanitarian crisis (Bojanowski, 2011;  
335 Donovan et al., 2018; Goitom et al., 2015). At the time, there were no seismic or other monitoring  
336 networks operating in Eritrea but earthquakes were felt around the volcano several hours and days  
337 prior to eruption, prompting evacuation (Goitom et al., 2015). This seismicity is the first of note  
338 in global catalogues for the region (Goitom et al., 2015). Despite this fortuitous warning, at least  
339 seven people were tragically killed and about 12,000 were displaced (Bojanowski, 2011; Goitom  
340 et al., 2015; Hamlyn et al., 2014). The eruption is particularly notable for the vast amount of SO<sub>2</sub>  
341 emitted into the atmosphere, one of the largest eruptive SO<sub>2</sub> masses globally since the eruption of  
342 Mount Pinatubo in 1991 (Fromm et al., 2014; Goitom et al., 2015; Theys et al., 2013), and the  
343 comparative rarity of recorded historical eruptions in the region (Goitom et al., 2015; Hamlyn et  
344 al., 2014).

345

346 In August, 2011, approximately two months after the eruption began, eight 3-component  
347 broadband seismometers (5 x Guralp CMG-6T, 3 x Guralp CMG-40T; Fig 2) were deployed  
348 around the volcano to monitor ongoing activity (Hamlyn et al., 2014). These stations remained  
349 operational for 14 months until October, 2012. The first two months of data were collected at a  
350 sample rate of 100 Hz before dataloggers were switched to a sample rate of 50 Hz for the remainder  
351 of the deployment to maximize data recovery while minimizing service runs. Manual phase arrival  
352 picking conducted on the first four months of data (2011-08-30 to 2011-12-31; Goitom, 2017;  
353 Hamlyn et al., 2014) identified a total of 2926 events, from which the first 35 days of data (all 100  
354 Hz sample rate) were quality checked and used for training and validating our transfer learning  
355 model. Five subsequent days of data (2 x 100 Hz days, 3 x 50 Hz days) were selected and quality  
356 checked to serve as test data. The reason to exclude 50 Hz data from model training is to emulate  
357 data availability in the early stages of this seismic deployment and demonstrate that changes in  
358 sample rate can be overcome without compiling new training datasets through a process known as  
359 data augmentation. The raw data for all datasets (training, validation and testing) were self-  
360 normalized, with linear trend removed, and left unfiltered.

361



362

363

364 **Figure 2.** Regional topographic map (90 m CGIAR Shuttle Radar Topography Mission and  
 365 GEBCO bathymetry model, grey-scale map center) and seismic deployment (30 m ALOS Digital  
 366 Surface Model, color map bottom right) around Nabro volcano. Red triangles (center map) indicate  
 367 Holocene volcanoes (Global Volcanism Program, 2013) with Nabro volcano highlighted in white.  
 368 Inverted blue triangles (bottom right map) indicate operational broadband seismic stations  
 369 deployed around Nabro volcano from August 2011 to October 2012 (station NAB6, inverted  
 370 yellow triangle, was flooded shortly after deployment and not operational). Training and validation  
 371 data were taken from dark blue stations only (NAB1, NAB2, NAB3, NAB4 and NAB8).

372

373 A total of 2921 waveforms with labelled P- and S-wave arrivals from 978 events (2011-  
374 08-30 to 2011-10-03) and five stations were used as training and validation data (only five stations  
375 were consistently operational during this time; dark blue stations in Fig 2 bottom right map).  
376 Training and validation data were grouped and divided so that no event appeared in both datasets  
377 to avoid data leakage (the model being trained on event data that also appears in validation or  
378 testing). 857 events (2498 waveforms) were used for model training and 121 events (423  
379 waveforms) were used for model validation, a training-validation split of approximately 85%-15%.  
380 624 sections of noise (20 secs length) were manually identified across all five stations (2011-08-  
381 31 to 2011-09-27), with 500 sections (2500 waveforms) and 85 sections (425 waveforms) used for  
382 model training and validation, respectively. Two noise waveforms were randomly dropped from  
383 each dataset so that the training and validation noise data comprise 2498 and 423 waveforms,  
384 respectively, to match the number of event waveforms.

385

386 A separate test dataset of 400 event waveforms with labelled P- and S-wave arrivals (132  
387 events) and 400 noise waveforms (80 sections of noise) was also produced for subsequent model  
388 testing. These data come from a different time period than those used for training and validation  
389 data, with 200 waveforms from a period where data were recorded at 100 Hz sample rate (2011-  
390 10-04 and 2011-10-05) and 200 waveforms from a period with 50 Hz sample rate (2011-10-14,  
391 2011-10-15 and 2011-11-27) for each category. All training, validation and test data were  
392 manually identified and quality checked.

393

394 The success of U-Net architectures relies on an effective data augmentation strategy when  
395 working with smaller datasets (Ronneberger et al., 2015). This allows the network to learn  
396 invariance to certain changes in input signal without them needing to appear in the annotated  
397 dataset. Here we outline a data augmentation strategy that improves performance of our U-GPD  
398 transfer learning model (Supplementary Materials). First, as all stations were switched from 100  
399 Hz sample frequency to 50 Hz sample frequency part way through the seismic deployment, we  
400 randomly select subsets of the training data (all originally sampled at 100 Hz) to be decimated to  
401 50 Hz sample frequency throughout training. Each training sample (i.e., each 3-component  
402 waveform) has a probability of 0.5 of being selected for decimation before each training epoch,

403 with an anti-aliasing, low-pass finite impulse response (FIR) filter applied and linear phase shift  
404 removed. Second, we randomly time-shift our P- and S-wave arrivals relative to the model input  
405 ‘window’, so that our waveforms differ slightly from epoch to epoch and the model must learn  
406 signal features that indicate arrivals rather than where they occur within the input window (i.e.,  
407 arrivals don’t need to occur in the center of the window for the model to detect them). With our  
408 noise data, a random 400-sample window is chosen at each training epoch from our 20-second  
409 noise sections, introducing more waveform variety between training epochs.

410

411 All data processing and model training/testing were performed in Python using the ObsPy  
412 (Beyreuther et al., 2010; Krischer et al., 2015; Megies et al., 2011), TensorFlow (Abadi et al.,  
413 2015; <https://tensorflow.org>) and Keras (Chollet et al., 2015; <https://keras.io>) libraries.

414

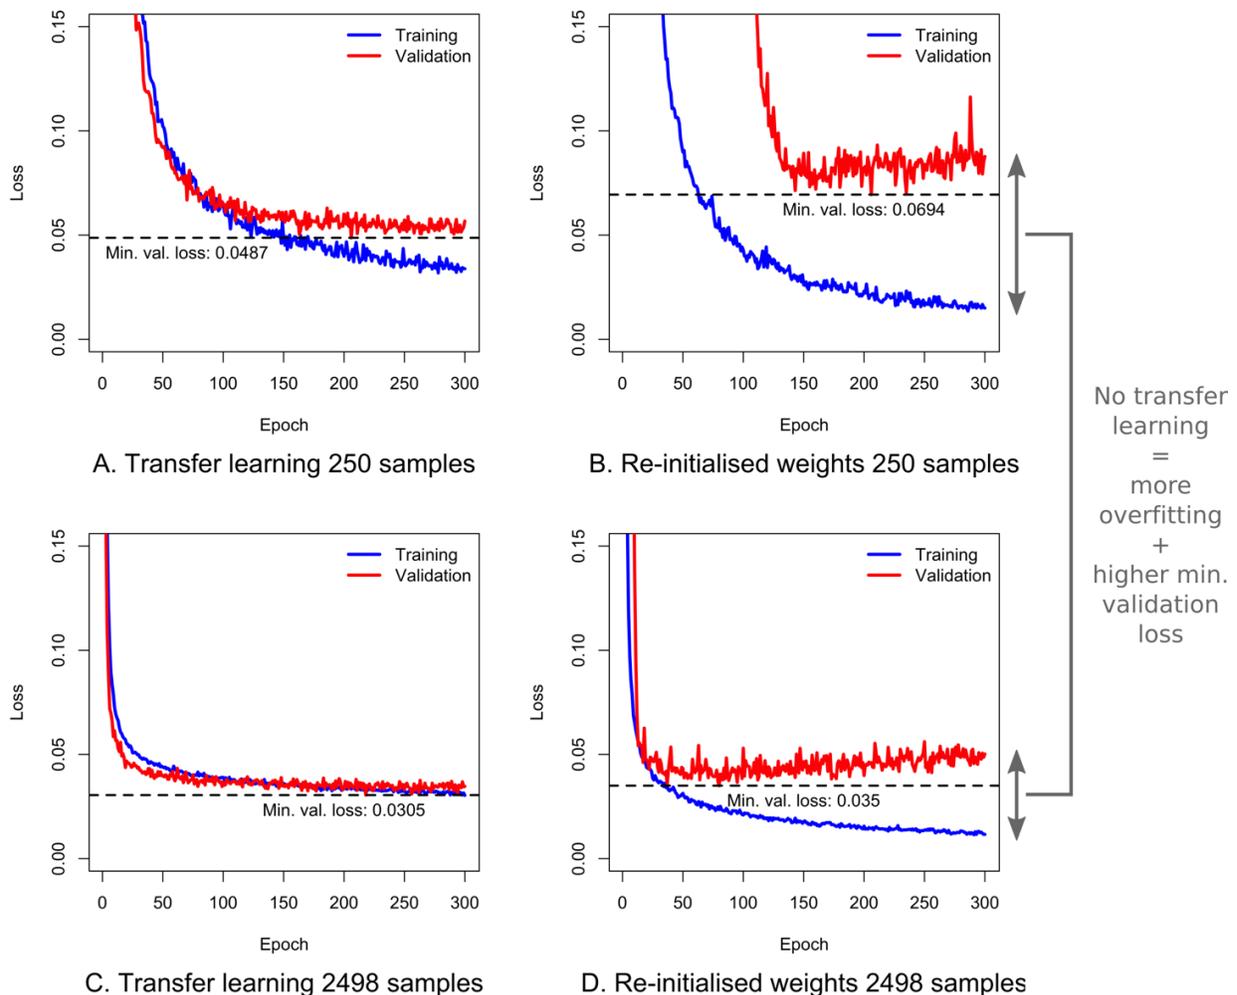
## 415 **5 Results**

### 416 **5.1 Training Metrics (Transfer Learning vs No Transfer Learning)**

417

418 To examine the impact of transfer learning and determine how much training data is  
419 required to produce an effective model, we use varying sized subsets of the training data  
420 throughout model training (i.e., 250, 500, 750, ..., 2000, 2250 and 2498 training samples). Figure  
421 3 compares how model loss (measure of distance between model predictions and ground truth  
422 labels) on training and validation data evolves throughout training between our transfer learning  
423 model and the same model with completely re-initialized weights (i.e., with no transfer learning)  
424 for our smallest and largest subsets of training data (250 and 2498 training samples, respectively).  
425 The learning rate is set to be equal ( $1 \times 10^{-3}$ ) across the whole re-initialized model as we are no  
426 longer fine-tuning existing knowledge. All other hyperparameters, including dropout rate, are kept  
427 the same. The models trained without transfer learning (Fig 3B and D) show a much greater degree  
428 of overfitting: the model loss on the training data continues to decrease with more training while  
429 the loss on validation data (data that the model does not use during training) hits an inflection point  
430 and starts increasing, reflecting that the model is ‘memorizing’ the precise features of the training  
431 data at the cost of generalization (Shorten & Khoshgoftaar, 2019). By contrast, the validation loss  
432 continues to decrease for the models trained with transfer learning (Fig 3A and C). Furthermore,

433 the minimum validation loss achieved by the transfer learning models for each training dataset size  
 434 is lower than when transfer learning is not employed (Fig 3 horizontal dashed lines). Such  
 435 diagnostics indicate that transfer learning is successfully preventing overfitting to the training data  
 436 and will likely produce a model that generalizes better to non-training data (Shorten &  
 437 Khoshgoftaar, 2019). The greatly improved performance on validation data using the smallest  
 438 subset of training data (Fig 3A and B) shows that transfer learning is particularly useful for  
 439 reducing overfitting and model loss when training data are very limited, but this advantage is  
 440 progressively diminished with increasing training dataset size (Figs 3 and 4).  
 441



442

443

444 **Figure 3.** Model loss vs. training epoch number. **a)** Transfer learning model and 250 training  
 445 samples of each class (P, S or neither). **b)** Model trained without transfer learning (i.e., initially

446 randomized weights) and 250 training samples of each class. **c)** Transfer learning model and full  
447 training dataset (2498 training samples of each class). **d)** Model trained without transfer learning  
448 (i.e., initially randomized weights) and full training dataset. Blue curve shows model loss for  
449 training data, red curve shows model loss for validation data (not seen during training). A lower  
450 model loss on training data (blue) than validation data (red) means the model shows signs of  
451 overfitting. The degree of overfitting (gap between blue and red curves) is much greater for the  
452 models without transfer learning (**b** and **d**) with validation loss hitting an inflection point then  
453 increasing whilst training loss continues to decrease. The transfer learning models also achieve a  
454 smaller minimum validation loss (horizontal dashed line) for each training set size.

455

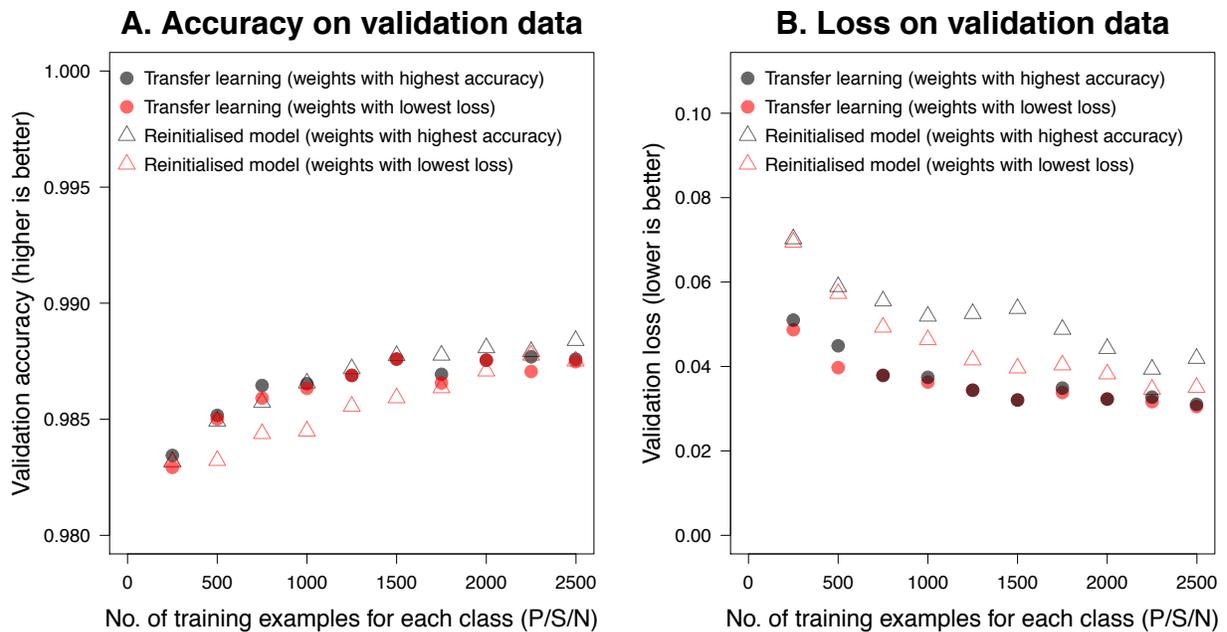
456 Figure 4 shows the highest model accuracy (the proportion of labels the model classifies  
457 correctly) and lowest model loss achieved by our transfer learning and re-initialized models on  
458 validation data when trained using each subset size of training data. The transfer learning model  
459 achieves lower model loss regardless of training dataset size (Fig 4B). As training dataset size  
460 increases, the difference between the lowest loss achieved by the two models (gap between red  
461 circles and red triangles, Fig 4B) decreases and the advantages of transfer learning diminish.  
462 Generally, loss is considered a more robust metric than accuracy for model performance on future  
463 data as it measures the distance between model predictions and ground truth labels, whereas  
464 accuracy simply measures a binary true/false score. However, accuracy still provides useful  
465 information regarding model performance. In particular, the transfer learning model shows a stable  
466 relationship between maximizing model accuracy and minimizing model loss (gap between black  
467 and red circles is very small for all training subset sizes), where the training strategy of minimizing  
468 model loss appears to achieve the same goal as maximizing model accuracy, again a sign of  
469 reduced overfitting. The re-initialized model (black and red triangles), on the other hand, shows a  
470 much less stable relationship in this regard, with diverging training scores (Fig 4) indicating that  
471 high model accuracy comes at the cost of higher model loss and low model loss comes at the cost  
472 of lower model accuracy for these small training set sizes when transfer learning is not employed.  
473 The increased model loss for model weights with highest model accuracy (black triangles) also  
474 suggests that the model has become overconfident in its predictions (it has large errors on the small  
475 proportion of labels it gets wrong) and is therefore likely to perform worse on out-of-distribution

476 data, with more false or missed phase arrival detections (e.g., a phase arrival being labelled as  
 477 noise with high model confidence, or vice versa).

478

479 Model performance between the two approaches (transfer learning vs re-initialization)  
 480 converges as training set size increases, indicating that the need for transfer learning decreases  
 481 with increased training set size, as expected. In fact, model performance with transfer learning  
 482 appears to plateau, or possibly even degrade, at training subset sizes of more than 1500 samples.  
 483 This suggests that, with enough training data, transfer learning could potentially inhibit the model's  
 484 ability to learn useful features in the new data that are absent in the original GPD training data.  
 485 This apparent variance in performance may also simply be a result of the stochasticity arising from  
 486 training using randomized weights in the new part of our transfer learning model.

487



488

489

490 **Figure 4.** Model accuracy (a) and loss (b) for various subsets of training data. Open red circles  
 491 are transfer learning model weights from epoch that achieves lowest validation loss (e.g., dashed  
 492 horizontal lines in Fig 3), open black circles are transfer learning model weights from epoch that  
 493 achieves highest validation accuracy, solid red triangles are re-initialized model (no transfer  
 494 learning) weights from epoch that achieves lowest validation loss, and solid black triangles are re-  
 495 initialized model weights from epoch that achieves highest validation accuracy.

496

497 **5.2 Test Dataset (Known Arrival Times)**

498

499 Following model training, we test the above models (i.e., new model with and without  
500 transfer learning) and two baseline models (GPD and PhaseNet) using the test dataset outlined in  
501 Section 4. We examine the proportion of correct class predictions (Fig 5) and the residuals between  
502 model and manually determined phase arrival pick times (Fig 6). Due to differences in model task  
503 types (classification vs segmentation), we apply all models as sliding windows over 1000-sample  
504 waveforms (note that the PhaseNet model takes a 3000-sample waveform as input so we examine  
505 only the middle 1000 samples for this model). To account for human picking error in collating our  
506 test set, we define a true positive for each phase arrival type (P or S) as the model prediction  
507 exceeding a given threshold value for that arrival type within 0.5 secs of the manually determined  
508 arrival, such that predicted arrival times very close to the manually determined arrival time are  
509 considered accurate. A true positive for sections of noise is defined as no phase arrival prediction  
510 exceeding a given threshold value at any point within that section of data. The test data are pre-  
511 processed as per the training data for each model (i.e., GPD model tested on 2 Hz high-pass filtered  
512 data and all other models, including PhaseNet, tested on raw data; all detrended and self-  
513 normalized).

514

515 The GPD model is tested using four different threshold values (Fig 5A – D) as this value  
516 strongly controls the number of false or missed phase arrival detections generated by this model.  
517 When the threshold is set to be whichever class label (P, S or N) has the highest predicted value  
518 for a given waveform, nearly all P- and S-wave arrivals are detected by the GPD model (99.75 %  
519 and 95 % detection rate, respectively; Fig 5A). However, this threshold criterion makes the GPD  
520 model extremely prone to false phase arrival detections in sections of noise, with 44 % of 1000-  
521 sample noise waveforms in our test dataset containing at least one false phase arrival detection  
522 (Fig 5A, bottom right square) and many of our 1000-sample event waveforms containing multiple  
523 phase arrival triggers (not shown). When this threshold criterion is applied to continuous sections  
524 of data from Nabro, the number of false phase arrival detections overwhelmingly outweighs the  
525 number of true phase arrival detections and becomes unmanageable in terms of correctly

526 associating phases, identifying true events and processing the data within computational memory  
527 constraints.

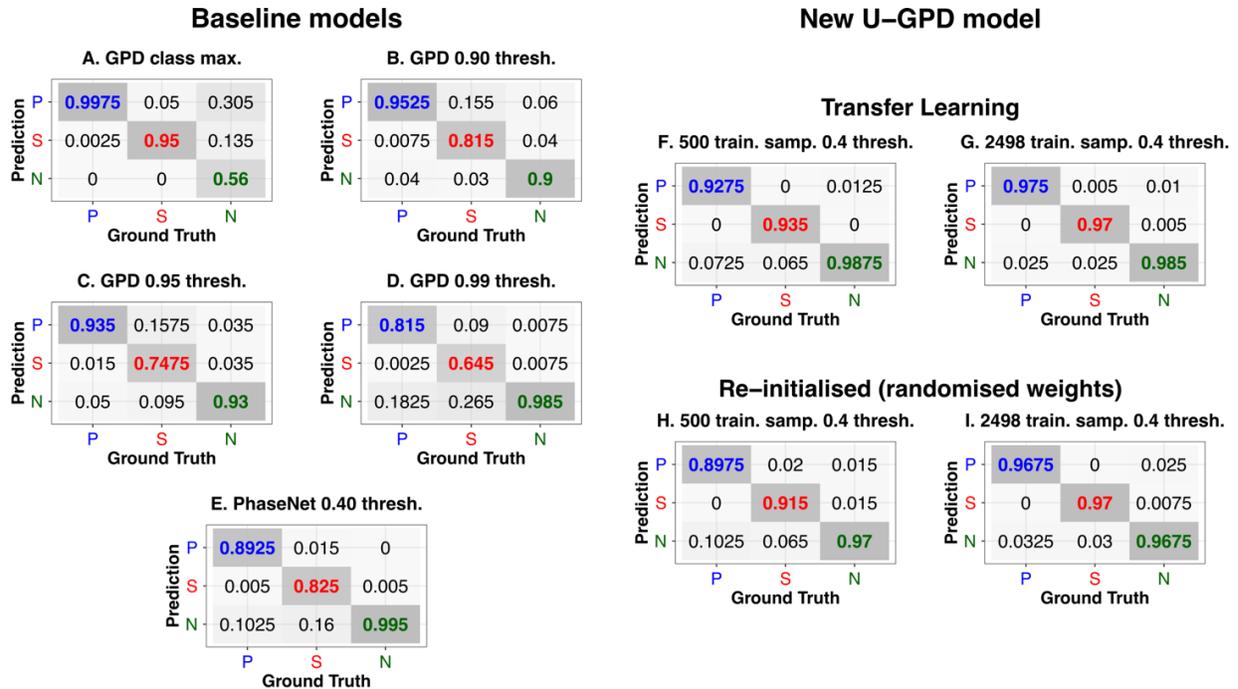
528

529         One way to lower the number of false phase arrival detections is to use a higher threshold  
530 value for P- and S-wave predictions. Figure 5B shows the GPD model's performance on our test  
531 data using a 0.9 threshold value (i.e., a P or S prediction 'probability' must exceed 0.9 to be  
532 included). The number of false detections in sections of noise is greatly reduced (down from 44 %  
533 of waveforms to 10 % of waveforms) but at the cost of reduced true phase arrival detections (~  
534 95% and ~82% of P- and S-wave arrivals, respectively). Part of this performance dip is  
535 undoubtedly due to the difference in sample rates between one half of the test data (50 Hz) and the  
536 GPD model's training data (all 100 Hz). When the threshold value is increased further (i.e., P or S  
537 prediction must exceed 0.95 or 0.99; Fig 5C and D), the GPD model yields even fewer false phase  
538 arrival detections in noise sections but at the cost of fewer P- and S-wave arrivals.

539

540         Figure 5E shows the performance of the PhaseNet model on our test dataset. This model is  
541 included as it adopts the same U-Net segmentation approach as our new model and is trained on  
542 data from a variety of instrument types, although the training data is still exclusively from southern  
543 California. The PhaseNet model is much less prone to false phase arrival detections than the GPD  
544 model (Fig 5E, bottom right square); as such, a much lower threshold value (0.4) can be used to  
545 maximize the number of true phase arrival detections. This model accurately identifies ~ 89% and  
546 ~ 83 % of P- and S-wave arrivals in our test dataset, which is better than the GPD model with a  
547 threshold value that achieves a similar false detection rate (e.g., Fig 5D), but detects fewer phase  
548 arrivals than our transfer learning and reinitialized models trained with Nabro data (Fig 5F – I).

549



550

551

552 **Figure 5.** Confusion matrices for base GPD model (a – d), PhaseNet model (e), U-GPD transfer  
 553 learning model (f, 500 training samples, and g, 2498 training samples) and re-initialized model (h,  
 554 500 training samples, and i, 2498 training samples). Values in matrices are proportion of ground  
 555 truth phase arrivals assigned by each model to a given class (values of 1 along diagonal from top  
 556 left to bottom right means all phase arrivals and sections of noise correctly identified).

557

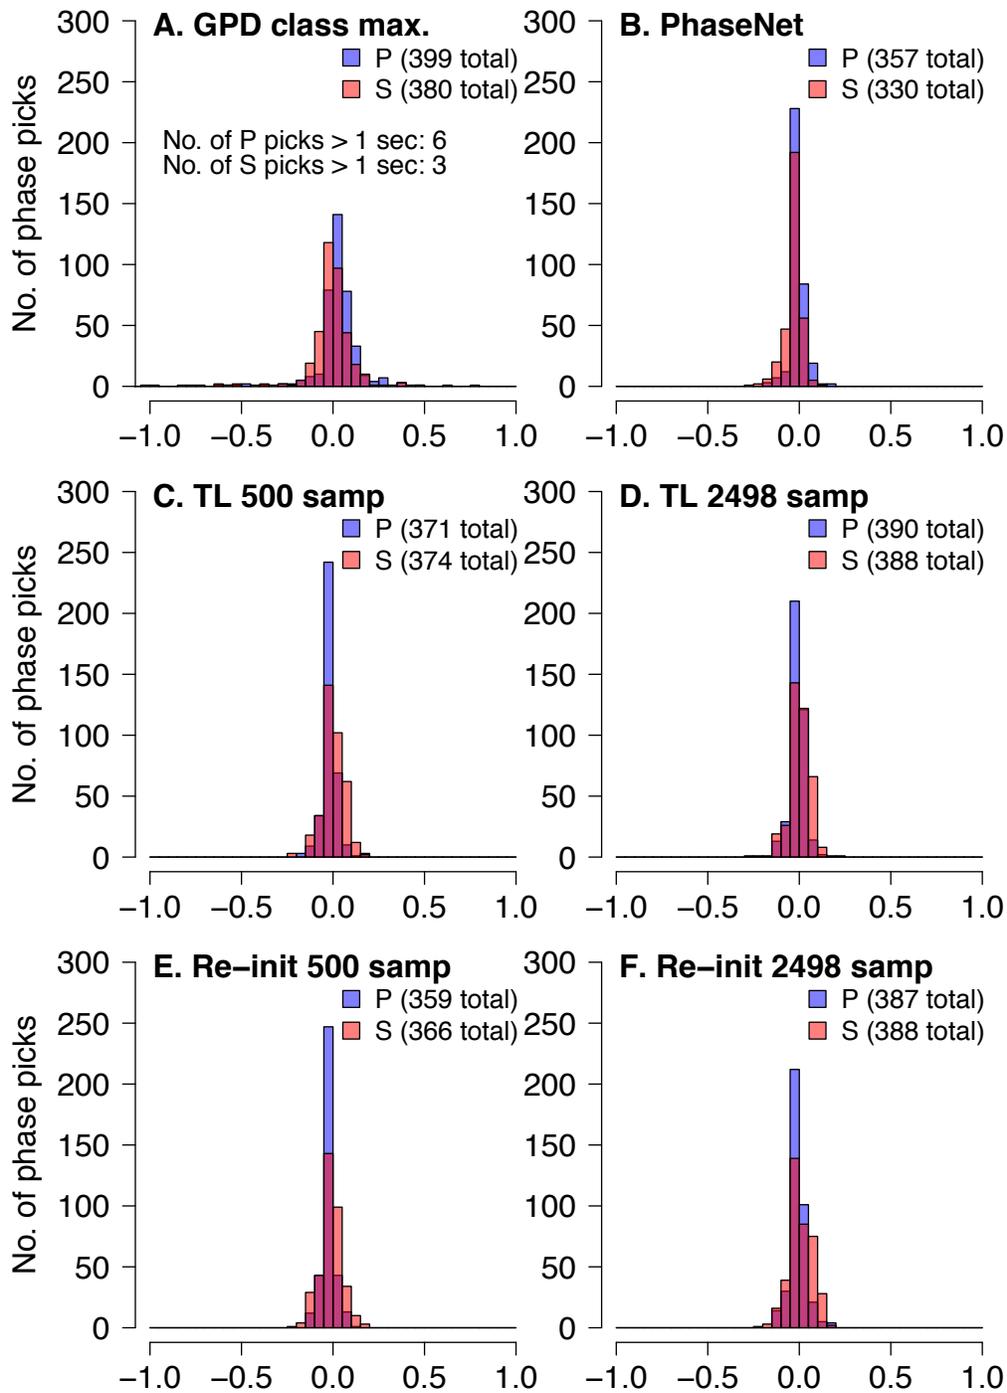
558 When trained using a subset of just 500 training samples for each class (P/S/N) and  
 559 evaluated using a prediction threshold value of 0.4, the transfer learning approach correctly detects  
 560 ~ 93% and ~94% of P- and S-wave arrivals with very few false phase arrival detections in sections  
 561 of noise (~ 1 %; Fig 5F), a clear improvement over our model trained with re-initialized weights  
 562 and the same training subset (Fig 5H). When our full training dataset is used (2498 samples for  
 563 each class), model performance converges between transfer learning (Fig 5G) and re-initialization  
 564 (Fig 5I), with a similar number of correctly identified phase arrivals and false detections in noise,  
 565 although the transfer learning model still performs marginally better, particularly on sections of  
 566 noise. In essence, the transfer learning model strikes a better balance between high phase arrival  
 567 detection rate (~ 97 – 98% for each phase arrival type; Fig 5G, top left and center squares) and low  
 568 false detection rates in sections of noise (~ 1%; Fig 5G, bottom right square) on our test data from

569 Nabro volcano than any of the existing baseline models (Fig 5A – E) or training a model from  
570 scratch (Fig 5I).

571

572         Figure 6 shows the residuals for each model between their predicted phase arrival times  
573 and the original manual pick times for these test waveforms. Predicted phase arrival times were  
574 determined using a simple trigger algorithm (e.g., Withers et al., 1998) on each model’s probability  
575 time series with the time series index that yields maximum predicted value chosen as the pick time  
576 for a given phase arrival type. The models that employ semantic segmentation (i.e., PhaseNet, our  
577 U-GPD transfer learning model and our re-initialized model; Fig 6B – F) show comparable pick  
578 time precision (root mean square deviation [RMSD] of 0.036, 0.038 and 0.044 seconds,  
579 respectively, for each model’s P-wave predictions and RMSD of 0.053, 0.053 and 0.065 seconds,  
580 respectively, for each model’s S-wave predictions), with no predicted arrival times more than 0.3  
581 seconds before or after the manually determined pick time, regardless of training dataset size. The  
582 GPD model (Fig 6A), by comparison, has a more diffuse range of phase arrival pick times (RMSD  
583 of 0.217 seconds for P-waves and 0.188 seconds for S-waves), with some model picks made more  
584 than 1 second before or after the manually determined arrival time. This is almost certainly a result  
585 of its more ambiguous class labelling (Fig 1) and the broad phase arrival probability peaks it  
586 generates.

587



Residuals vs. manual phase picks (secs)

588

589

590 **Figure 6.** Model phase pick residuals vs. manual phase picks for base GPD model (a), PhaseNet

591 model (b), U-GPD transfer learning model (c, 500 training samples, and d, 2498 training samples),

592 and reinitialized model (e, 500 training samples, and f, 2498 training samples). The models based  
593 on semantic segmentation (b – f) yield smaller phase pick residuals.

594

### 595 **5.3 Full 14-Month Deployment (Unknown Arrival Times)**

596

597 Whilst evaluating model performance on individual, manually scrutinized waveforms is  
598 useful for benchmarking and yielding estimates of model efficacy, the model’s performance in a  
599 ‘real-world’ setting is ultimately of most importance to seismic analysts. Evaluating such  
600 performance is inherently more challenging, however, as the number of events in long sections of  
601 monitoring data and their respective phase arrival times are unknown, and other considerations,  
602 such as computational time and resources (e.g., memory requirements and availability of optimized  
603 hardware), affect model feasibility as a monitoring tool.

604

605 In this section, we present results of our best performing model in the prior section (U-  
606 GPD transfer learning model trained with full training dataset of 2498 samples of each class) and  
607 the original base GPD model when run over the full 14-month Nabro seismic deployment (Fig 7).  
608 As with the test dataset in Section 5.2, phase arrivals are detected at individual stations through a  
609 simple trigger algorithm, where an arrival is detected if the probability assigned to that class label  
610 (P or S) exceeds a given threshold (e.g., 0.4 for our U-GPD transfer learning model). The phase  
611 arrival time is determined as the waveform sample with the highest probability for that phase.

612

613 The U-GPD transfer learning model was applied to the data as a sliding window with 50  
614 % overlap (i.e., applied at ‘time shifts’ of 200 samples) over 24-hour sections of data from each  
615 individual station. The model takes 5 seconds to process 24 hours of 3-component data at 100 Hz  
616 sample rate (or 3 seconds per day at 50 Hz sample rate) on a single graphics processing unit (GPU;  
617 NVIDIA GeForce RTX 2080 Ti), a rate many orders of magnitude faster than ‘real-time’ even  
618 when run on hundreds of stations. To avoid poor predictions due to window edge effects, only the  
619 middle 200 sample predictions out of 400 from each window are used to predict phase arrivals and  
620 are concatenated to produce one long continuous prediction trace without overlap or gaps and with  
621 the same sample rate as that of the input signal (i.e., 100 or 50 Hz). With all other processing steps  
622 (e.g., software initialization, data read/write, signal windowing, running trigger algorithm, etc.),

623 the U-GPD transfer learning model picks phase arrivals at all 7 available stations from the full 14-  
624 month deployment in less than 4 hours using a single GPU (greatly reduced when parallelized over  
625 multiple GPUs), indicating that it could easily be used within real-time monitoring constraints.

626

627         Conversely, as the GPD model produces only one class prediction per window (Fig 1A),  
628 we apply this model with much greater overlap (97.5 %; every 10 samples of data) and with  
629 varying threshold values (0.9, 0.95 and 0.99) for phase arrival detection triggering. This generates  
630 a prediction trace with a much coarser sample rate than the original input signal (i.e., from 100 or  
631 50 Hz to 10 or 5 Hz, respectively) and takes 26 seconds per 24 hours' 3-component data at 100 Hz  
632 sample rate (or 15 seconds per day at 50 Hz sample rate) on the same NVIDIA GPU, approximately  
633 a five-fold increase in computational time with a tenth of the temporal detail. With all other  
634 processing steps, the GPD model took almost 50 hours to run over the full 14-month deployment  
635 using a single GPU, more than a ten-fold increase in computational time over the transfer learning  
636 model, due to more (pre-)processing required (e.g., more signal windows generated and  
637 subsequent processing). Assuming a linear increase in computational time, running the model as a  
638 sliding window over every sample of data would take  $\sim 260$  seconds per 24 hours' 3-component  
639 data at 100 Hz sample rate and  $\sim 500$  hours (nearly 3 weeks) for the full 14-month deployment and  
640 7 stations. While this is still faster than real-time, these timescales for a single or limited number  
641 of station(s) could become limiting when applied at hundreds of stations, particularly without high  
642 performance computing resources.

643

### 644 **5.3.1 Phase Association Method**

645

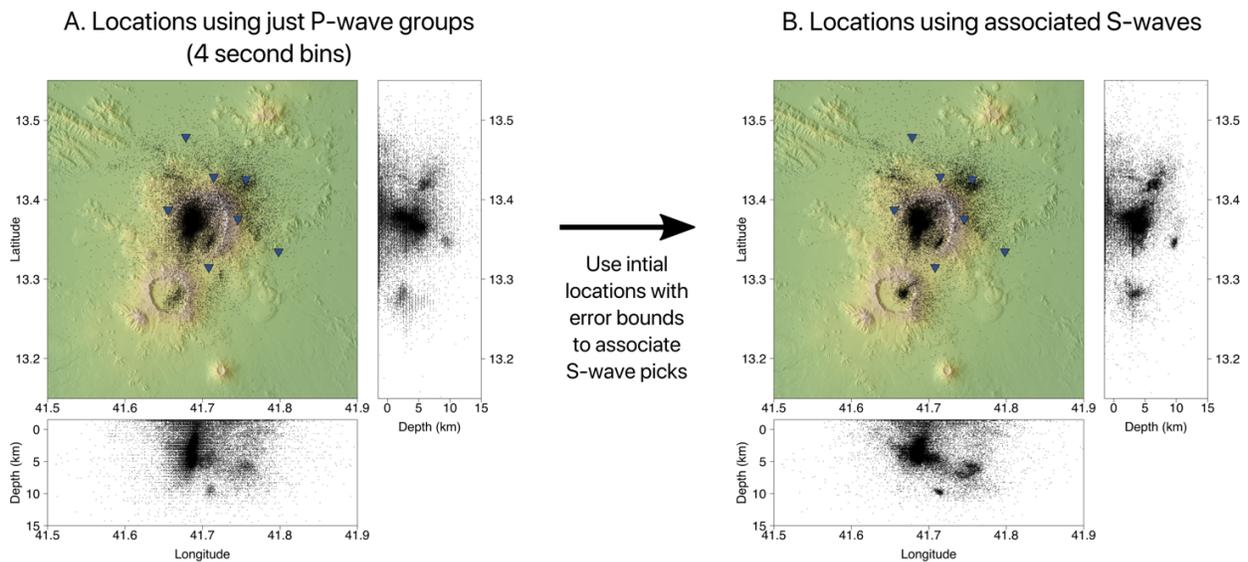
646         Both models detect P- and S-wave phase arrivals but do not associate them to the same  
647 event. To assess the number of locatable events detected, we group P-wave phase arrival triggers  
648 into 4-second bins and keep only bins with arrivals detected at four or more stations. This bin size  
649 was chosen to encompass the maximum plausible travel time between any two stations. If multiple  
650 arrivals were detected at the same station within a 4-second bin, the detection threshold was  
651 increased for all arrivals in that particular bin to retain only the highest probability phase picks. If  
652 any of these bins now had arrivals at less than four stations, as a result of removing lower

653 probability phase picks, they were discarded as there would be too few stations to constrain event  
 654 location. If there were still multiple arrivals present at any given station, only the arrivals with  
 655 highest probability for each station were kept. Finally, if phase arrival bins intersected (a subset of  
 656 one bin was contained in another), the bin with highest mean probability was kept. This association  
 657 method is clearly quite crude, and only works for small, local arrays, but allows a broad evaluation  
 658 of model performance at detecting phase arrivals.

659

660 We associate S-wave arrivals to their corresponding P-wave arrivals by first locating events  
 661 using NonLinLoc (e.g., Lomax et al., 2000), a widely used software package for probabilistic  
 662 earthquake location, using the P-wave arrival bins outlined above and a simple 1D velocity model  
 663 (Fig 7A). The difference between P-wave arrival and event origin times were used to predict which  
 664 S-wave arrival detections should be associated with each P-wave arrival using a  $V_p/V_s$  ratio of  
 665 1.76 and S-wave travel time error of 0.25 (25%). S-wave arrival triggers that lay within this error  
 666 bound for each detected P-wave arrival were associated to that event. S-wave arrivals at stations  
 667 without a detected P-wave arrival were not included. All events were then located again in  
 668 NonLinLoc using all included phase arrivals (Fig 7B).

669



670

671

672 **Figure 7.** U-GPD transfer learning model event locations (total no. of events = 33,950) using  
 673 automated phase association strategy. **a)** P-wave phase arrival triggers are grouped into 4 second

674 bins and these groupings are used to obtain initial event hypocenters and origin times. **b)** S-wave  
675 phase arrival triggers are associated to P-waves in **(a)** using initial origin times, a  $V_p/V_s$  ratio of  
676 1.76 and a travel-time error of 25 %. Events are then located again using all included P-wave and  
677 S-wave arrivals.

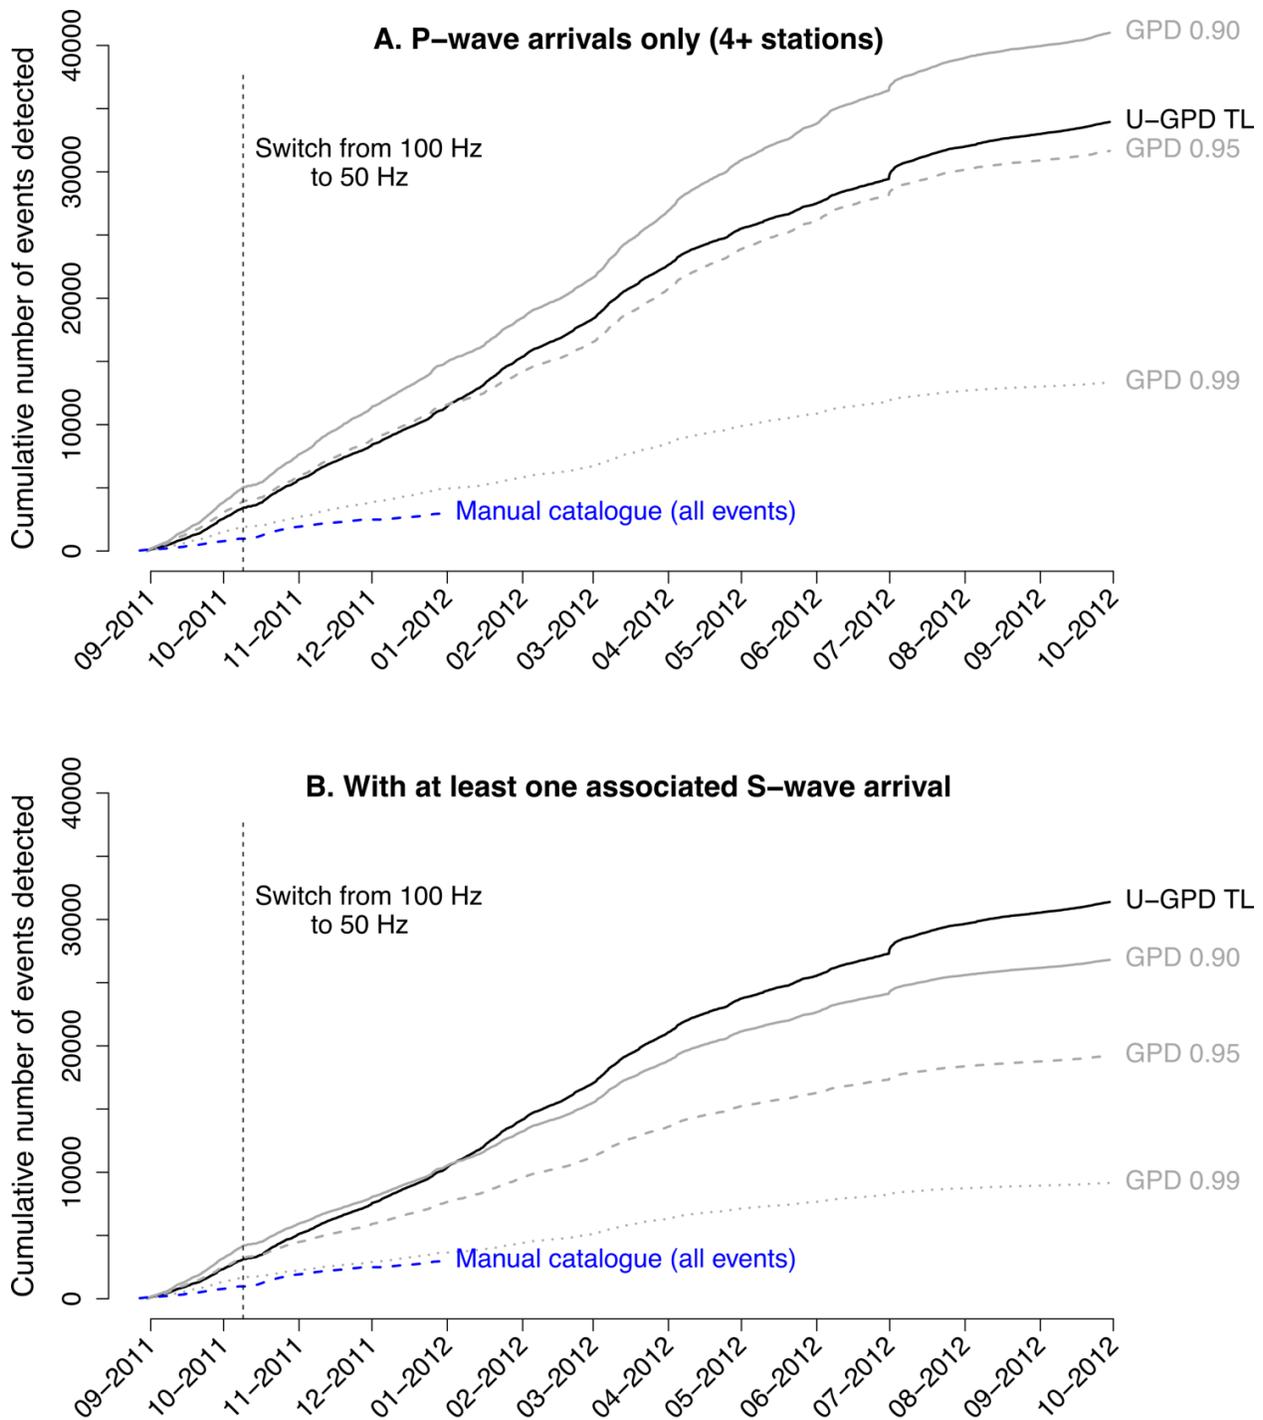
678

### 679 **5.3.2 Detected Events and Location Errors**

680

681 Figure 8 shows the cumulative number of events detected by the U-GPD transfer learning  
682 model (threshold value of 0.4; black solid line) and the original GPD model (threshold values of  
683 0.9, 0.95 and 0.99; grey lines). The cumulative number of events from an existing manual  
684 catalogue for this deployment (Goitom, 2017; Hamlyn et al., 2014), some of which provided the  
685 transfer learning model training data, is also given for reference. When only P-wave arrivals are  
686 used (Fig 8A), the GPD model with detection threshold of 0.9 appears to detect the most events  
687 (total no. of events detected by GPD model = 41,007; total no. of events detected by transfer  
688 learning model = 33,950). A threshold of 0.95 also detects more events than the transfer learning  
689 model until shortly after the switch in instrument sample rates from 100 Hz to 50 Hz. However,  
690 when we consider events with at least one associated S-wave arrival, the transfer learning model  
691 detects more events overall (Fig 8B; no. of events detected by transfer learning model = 31,387;  
692 no. of events detected by GPD model with 0.9 threshold = 26,808). This is consistent with the  
693 results from our test dataset in Section 5.2, with the proportion of S-wave arrivals accurately  
694 detected by the GPD model at these threshold values much lower than the proportion of P-wave  
695 arrivals detected (Fig 5B – D). Furthermore, 6 % of noise waveforms and 16% of S-wave arrivals  
696 from our test data were mislabeled by the GPD model (0.9 threshold value) as P-wave arrivals (Fig  
697 5B), a higher rate of false detections or labels than the transfer learning model (1 % of noise  
698 sections and 0.5% of S-waves, respectively; Fig 5G). This means that a higher proportion of the  
699 P-wave groupings detected by this model with 0.9 threshold value are likely to include mislabeled  
700 S-waves or false arrivals.

701



702

703

704 **Figure 8.** Cumulative number of events detected by GPD model (various thresholds, grey lines)

705 and transfer learning model trained on full Nabro dataset (2498 samples of each class, 0.4

706 threshold, black line). Blue dashed line is existing manual catalogue (Goitom, 2017). All training

707 / validation waveforms are from dates before switch in sample frequency (vertical dashed line). **a)**

708 Cumulative number of events detected using P-wave arrivals only (see main text for event binning  
709 procedure). **b)** Cumulative number of events with at least one associated S-wave arrival.

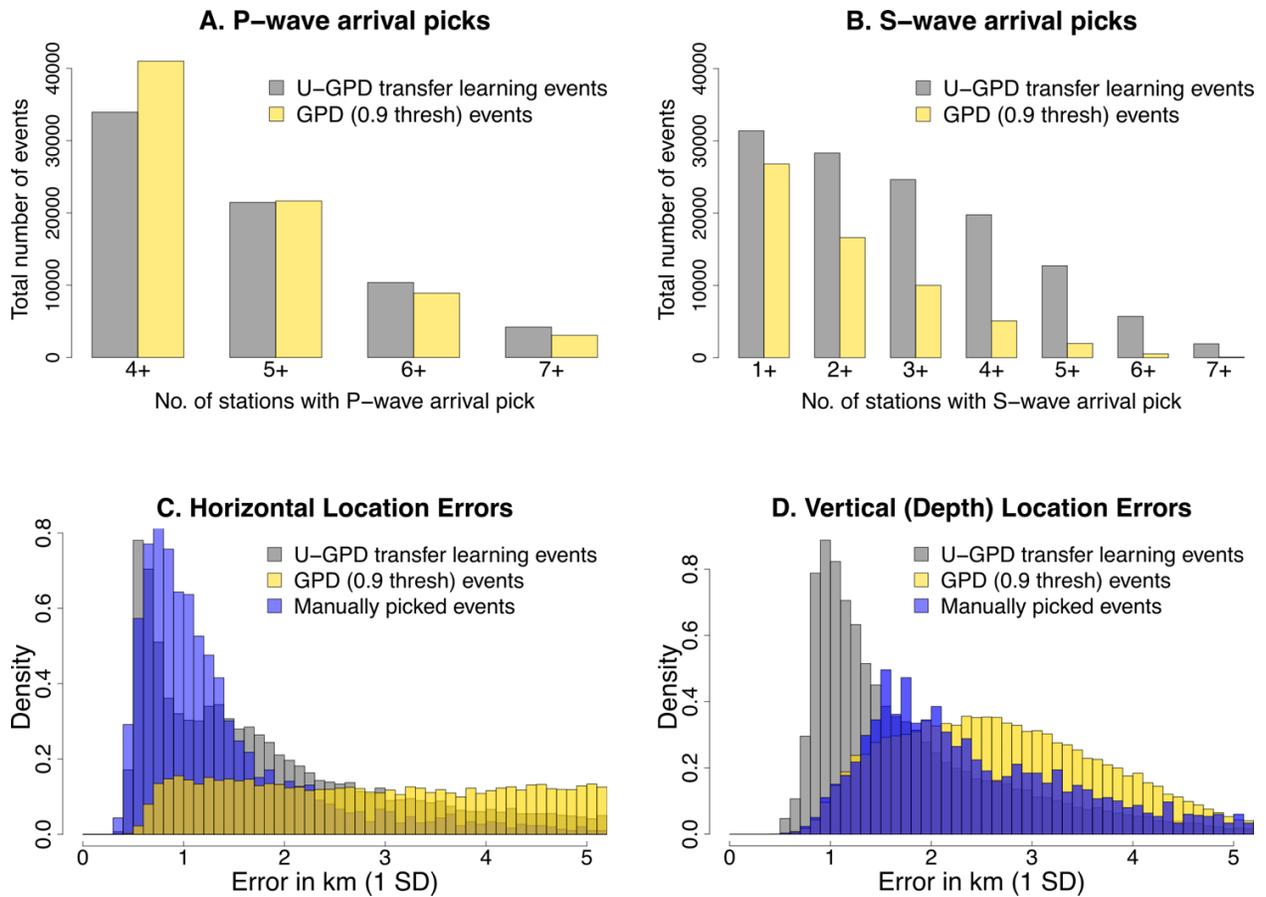
710

711 To scrutinize these results further, we examine the number of stations with P- and S-wave  
712 arrival detections per event (Fig 9A – B). In general, the events detected and picked by the U-GPD  
713 transfer learning model include more stations and considerably more S-wave arrivals than those  
714 picked by the GPD model, although the number detected by the GPD model may have been  
715 reduced by using a coarser prediction trace (every 10 samples, a requirement to reduce model run  
716 time to a reasonable timeframe). This increase in the number of stations and S-wave arrivals per  
717 event will constrain event locations, as seen in the location errors derived from the models' phase  
718 arrival picks (Fig 9C – D).

719

720 The horizontal errors (Fig 9C; standard deviation of estimated Gaussian error distributions)  
721 for the locations produced using the transfer learning model pick times are comparable to the  
722 existing manually picked events. Furthermore, vertical (depth) errors are much improved over the  
723 manual catalogue (Fig 9D), likely reflecting more consistency in S-wave arrival picking than that  
724 of a manual analyst. The GPD model, by comparison, produces a more diffuse range of horizontal  
725 and vertical errors, which is likely to be a combination of coarser prediction trace, poorer pick  
726 precision (Fig 6A), lack of S-wave arrivals (Fig 9B) and false/mislabeled P-wave arrival detections  
727 (Fig 5B). This interpretation is further supported when we look at the number of event locations  
728 lying within the array (i.e., event locations lying within the convex hull of station coordinates) for  
729 each model: NonLinLoc locates more events within the array using the transfer learning picks ( $n$   
730  $= 23,859$ ) than using the GPD model with 0.9 threshold value ( $n = 22,826$ ). While we expect many  
731 events to occur outside of the array (e.g., at neighboring faults or volcanic centres), this metric  
732 shows that a much larger proportion of event locations detected by the GPD model lie away from  
733 the volcanic edifice, which may reflect poorer pick precision, false/mislabeled arrivals or coarser  
734 prediction trace, but may also reflect the event types (i.e., regional tectonic) that the original model  
735 was trained on.

736



737

738

739 **Figure 9. a)** Number of P-wave arrival picks per event for transfer learning model (grey) and base  
 740 GPD model (gold). **b)** Number of S-wave arrival picks per event. **c)** Histogram of Gaussian  
 741 horizontal location errors (1 standard deviation) for events picked by transfer learning model (grey)  
 742 and base GPD model (gold), and those in the existing manual catalogue (blue). **d)** Histogram of  
 743 Gaussian vertical (depth) location errors (1 standard deviation).

744

## 745 **6 Discussion**

746

747 Transfer learning using existing seismological deep learning models can be a highly  
 748 effective strategy to automate phase arrival picking in settings with little or no prior monitoring.  
 749 We demonstrate that, with a limited number of hand-labelled waveforms (on the order of hundreds  
 750 to low thousands) and a few minutes of training time, one can produce a consistent and effective

751 deep learning model for phase arrival detection that requires no other manual intervention or tuning  
752 and can process years of data in a matter of hours.

753

754 For small training datasets, the use of pre-existing, generalized CNN filters greatly reduces  
755 model overfitting (i.e., model parameters ‘memorizing’ the training data) when compared with  
756 training a model from scratch (Fig 3) and yields a more stable relationship between maximizing  
757 model accuracy and minimizing model error (Fig 4). Furthermore, when combined with a good  
758 data augmentation strategy, transfer learning can also address the issue of processing data when  
759 instrument sample rates differ from those used to train existing models. When applied to data from  
760 Nabro volcano, augmenting our training set with decimated waveforms greatly improves model  
761 performance on lower sample rate data (Supplementary Materials). As such, hand-labelled training  
762 data from the first 35 days of the deployment (all 100 Hz sample rate) were sufficient to detect  
763 phase arrivals throughout the duration of the deployment, even after instrument sample rates were  
764 switched to 50 Hz (Fig 8). Without this data augmentation step, model performance on lower  
765 sample rate data declines dramatically (Supplementary Materials). This shows that where sample  
766 rates are altered or new instruments added during a seismic deployment, data augmentation can  
767 overcome the cost of collecting further hand-labelled data and allow models to be adapted cheaply  
768 and quickly throughout the deployment.

769

770 The introduction of new, task-specific data and the change in model task from one of  
771 classification to one of segmentation also improves our U-GPD model pick time precision (Fig 6),  
772 the number of stations per detected event (Fig 9A), the number of S-wave arrivals detected (Figs  
773 5 and 9B) and computational efficiency over the original base GPD model, as well as potentially  
774 reducing the number of false/mislabeled P-wave detections (Fig 5) and increasing the number of  
775 identified events that relate directly to volcanic activity (evidenced by the increased number of  
776 events located within the array). Without manual intervention or sophisticated phase association,  
777 phase arrival picks from the U-GPD transfer learning model produce locations with smaller depth  
778 errors than the base GPD model and even manually determined phase arrival times (Fig 9D). This  
779 is likely a result of more consistent picking and labelling, particularly for S-wave arrivals, which  
780 is difficult even for manual analysts to perform consistently, and suggests that very few of the  
781 events detected are false.

782

783           Given the greatly improved computational time over the base GPD model, the small  
784 number of training events required and the use of a high-level, user-focused programming library  
785 (Keras), this approach is well within the reach of volcano observatories and research groups.  
786 Previous studies that analyze the pre-, syn- and post-eruptive periods at Nabro volcano have relied  
787 on manually-produced seismic catalogues comprising hundreds of events (e.g., Goitom et al.,  
788 2015; Hamlyn et al., 2014; the latter locating 658 events over 38 days, a rate of < 18 events per  
789 day). Our U-GPD transfer learning model yields a seismic catalogue that is order of magnitudes  
790 larger (33,950 events over 396 days, a rate of > 85 events per day; Figs 7 and 8), with smaller  
791 location errors (Fig 9), in a matter of hours. Furthermore, as the model processes 1D waveform  
792 data, as opposed to 2D spectrogram images in some other existing models (e.g., Dokht et al., 2019;  
793 Lara et al., 2020; Titos et al., 2020), it runs quickly on high resolution data without using a GPU  
794 optimized for deep learning frameworks (32 secs per 24 hours of 100 Hz data on an Intel Core i7  
795 desktop CPU) and so could easily be deployed for real-time monitoring with limited computing  
796 resources or at much larger arrays. The methods and computational times in this paper have relied  
797 on standard, generic libraries (ObsPy, TensorFlow and Keras); the use of more optimized,  
798 compiled code or higher-performance / lower-level languages (e.g., Julia and C) could greatly  
799 improve computational times further.

800

801           Beyond phase arrival picking, the generalized waveform features extracted by existing,  
802 extensively trained models, such as the GPD model (Fig 1A), could serve as a useful feature  
803 extraction system for models designed for other waveform processing tasks. For example,  
804 information regarding frequency content and orientation of seismic energy extracted by the GPD  
805 model (Fig 1A inset) could reasonably provide useful features for a new model designed to  
806 automatically classify volcano seismic event types (e.g., Bueno et al., 2020; Hibert et al., 2017;  
807 Lara et al., 2020), particularly when available annotated datasets are small or unbalanced.  
808 However, with larger datasets, there is the potential for transfer learning to inhibit learning of new,  
809 useful features, particularly if the source and target tasks or data distributions differ considerably.

810

811           The number of seismological studies to date that employ transfer learning is relatively low  
812 (e.g., Bueno et al., 2020; Chai et al., 2020; El Zini et al., 2020; Huot et al., 2018; Titos et al., 2020).

813 This is undoubtedly, in part, due to the lack of extensively trained, well-documented, publicly  
814 available seismological models. However, the number is likely to grow as more extensive datasets  
815 and models are developed and released into the public domain. We credit the availability of the  
816 GPD model in the public domain and use of a popular, user-focused machine learning framework  
817 (Keras) as the foundation of the work presented in this paper. Such availability facilitates  
818 adaptation and experimentation; development of other publicly available models and extensive  
819 datasets would aid progress in the field of seismological machine learning.

820

821 Whilst the application of transfer learning can overcome the perception that deep learning  
822 models require a ‘large upfront cost’ in terms of data and computational resources, the  
823 development and benchmarking of large-scale, extensive models and datasets are still imperative  
824 to push the field of seismological machine learning forwards and extend applications to all aspects  
825 of seismic processing and inference. However, it is hoped that applications such as the one  
826 presented in this paper will motivate the initial investment in the development of such models so  
827 that the cost of producing effective task-specific models (e.g., through transfer learning) is  
828 progressively reduced.

829

830

### 831 **Acknowledgments**

832

833 The seismic data were collected with funding from the Natural Environment Research Council  
834 (NERC) project NE/J012297/1 (“Mechanisms and implications of the 2011 eruption of Nabro  
835 volcano, Eritrea”). The UK seismic instruments and data management facilities were provided  
836 under loan number 976 by SEIS-UK at the University of Leicester. The facilities of SEIS-UK are  
837 supported by NERC under Agreement R8/H10/64. Author SL was supported by a GW4+ Doctoral  
838 Training Partnership studentship from the Natural Environment Research Council (NERC)  
839 [NE/L002434/1]. Author BG was funded by the Engineering and Physical Sciences Research  
840 Council (EPSRC) and the School of Earth Sciences at the University of Bristol. Author MJW was  
841 funded by UKRI GCRF EP/P028233/1 (“PREPARE”) and NERC NE/R017956/1  
842 (“EQUIPT4RISK”). Author JMK was funded by NERC grant NE/R018006/1. Author KVC

843 was supported by the AXA Research Fund. We gratefully acknowledge support from the sponsors  
 844 of the Bristol University Microseismicity ProjectS (BUMPS) and the NERC Centre for the  
 845 Observation and Modelling of Earthquakes, volcanoes and Tectonics (COMET). We also  
 846 gratefully acknowledge the cooperation we received from the Eritrea Institute of Technology,  
 847 Eritrean government, Southern and Northern Red Sea Administrations, local sub-zones and village  
 848 administrations. We thank the Department of Mines, Ministry of Energy and Mines for their  
 849 continued support throughout the project. Special thanks go to Zerai Berhe, Mebrahtu Fisseha,  
 850 Michael Eyob, Ahmed Mohammed, Kibrom Nerayo, Asresehey Ogbatsien, Andemichael  
 851 Solomon and Isaac Tuum. We thank Alem Kibreab and Prof. Ghebrebrhan Ogubazghi for their  
 852 vital help in facilitating the fieldwork.

853  
 854

## 855 **Data Availability Statement**

856

857 All seismic data from the Nabro Urgency Array are publicly available through IRIS Data Services  
 858 (<http://service.iris.edu/fdsnws/dataselect/1/>). See [http://www.fdsn.org/networks/detail/4H\\_2011/](http://www.fdsn.org/networks/detail/4H_2011/) for further  
 859 details on data access and availability. Model training, validation and test sets / metadata are archived and  
 860 available through Zenodo (Lapins et al., 2021; <https://doi.org/10.5281/zenodo.4498549>). Full code to  
 861 reproduce our U-GPD transfer learning model, perform model training, run the U-GPD model over  
 862 continuous sections of data and use model picks to locate events in NonLinLoc (Lomax et al., 2000) are  
 863 available at <https://github.com/sachalapins/U-GPD>, with the release (v1.0.0) associated with this paper also  
 864 archived and available through Zenodo (Lapins, 2021; <https://doi.org/10.5281/zenodo.4558121>).

865

866

## 867 **References**

868

- 869 Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., ... Zheng, X. (2015). TensorFlow: Large-Scale  
 870 Machine Learning on Heterogeneous Distributed Systems (software available from  
 871 <https://www.tensorflow.org>).
- 872 Barbedo, J. G. A. (2018). Impact of dataset size and variety on the effectiveness of deep learning and transfer  
 873 learning for plant disease classification. *Computers and Electronics in Agriculture*, 153, 46–53.  
 874 <https://doi.org/10.1016/j.compag.2018.08.013>
- 875 Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann, J. (2010). ObsPy: A python toolbox

- 876 for seismology. *Seismological Research Letters*, 81(3), 530–533. <https://doi.org/10.1785/gssrl.81.3.530>
- 877 Bojanowski, A. (2011). Volcano mix-up. *Nature Geoscience*, 4(8), 495. <https://doi.org/10.1038/ngeo1222>
- 878 Bueno, A., Benitez, C., De Angelis, S., Diaz Moreno, A., & Ibanez, J. M. (2020). Volcano-Seismic Transfer
- 879 Learning and Uncertainty Quantification with Bayesian Neural Networks. *IEEE Transactions on Geoscience*
- 880 *and Remote Sensing*, 58(2), 892–902. <https://doi.org/10.1109/TGRS.2019.2941494>
- 881 Chai, C., Maceira, M., Santos-Villalobos, H. J., Venkatakrishnan, S. V., Schoenball, M., Zhu, W., ... Thurber, C.
- 882 (2020). Using a Deep Neural Network and Transfer Learning to Bridge Scales for Seismic Phase Picking.
- 883 *Geophysical Research Letters*, 47(16). <https://doi.org/10.1029/2020GL088651>
- 884 Chollet, F., & and others. (2015). Keras (software available from <https://keras.io>).
- 885 D'souza, R. N., Huang, P.-Y., & Yeh, F.-C. (2020). Structural Analysis and Optimization of Convolutional Neural
- 886 Networks with a Small Sample Size. *Scientific Reports*, 10(834). <https://doi.org/10.1038/s41598-020-57866-2>
- 887 Daumé, H. (2007). Frustratingly easy domain adaptation. In *45th Annual Meeting of the Association for*
- 888 *Computational Linguistics (ACL)* (pp. 256–263). <https://arxiv.org/abs/0907.1815>
- 889 Dokht, R. M. H., Kao, H., Visser, R., & Smith, B. (2019). Seismic Event and Phase Detection Using Time–
- 890 Frequency Representation and Convolutional Neural Networks. *Seismological Research Letters*, 90(2A), 481–
- 891 490. <https://doi.org/10.1785/0220180308>
- 892 Donovan, A., Blundy, J., Oppenheimer, C., & Buisman, I. (2018). The 2011 eruption of Nabro volcano, Eritrea:
- 893 perspectives on magmatic processes from melt inclusions. *Contributions to Mineralogy and Petrology*, 173(1),
- 894 1–23. <https://doi.org/10.1007/s00410-017-1425-2>
- 895 Dramsch, J. S., & Lüthje, M. (2018). Deep-learning seismic facies on state-of-the-art CNN architectures. In *SEG*
- 896 *Technical Program Expanded Abstracts* (pp. 2036–2040). Anaheim, CA, USA.
- 897 Efremova, Di. B., Sankupellay, M., & Konovalov, D. A. (2019). Data-Efficient Classification of Birdcall through
- 898 Convolutional Neural Networks Transfer Learning. In *2019 Digital Image Computing: Techniques and*
- 899 *Applications (DICTA)* (pp. 1–8). Perth, Australia. <https://doi.org/10.1109/DICTA47822.2019.8946016>
- 900 El Zini, J., Rizk, Y., & Awad, M. (2020). A Deep Transfer Learning Framework for Seismic Data Analysis: A Case
- 901 Study on Bright Spot Detection. *IEEE Transactions on Geoscience and Remote Sensing*, 58(5), 3202–3212.
- 902 <https://doi.org/10.1109/TGRS.2019.2950888>
- 903 Fromm, M., Kablick III, G., Nedoluha, G., Carboni, E., Grainger, R., Campbell, J., & Lewis, J. (2014). Correcting
- 904 the record of volcanic stratospheric aerosol impact: Nabro and Sarychev Peak. *Journal of Geophysical*
- 905 *Research: Atmospheres*, 119, 10343–10364. <https://doi.org/10.1002/2014JD021507>
- 906 Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., ... Lempitsky, V. (2016). Domain-
- 907 Adversarial Training of Neural Networks. *Journal of Machine Learning Research*, 17(59), 1–35.
- 908 <https://arxiv.org/abs/1505.07818>
- 909 Gibbons, S. J., & Ringdal, F. (2006). The detection of low magnitude seismic events using array-based waveform
- 910 correlation. *Geophysical Journal International*, 165(1), 149–166. [https://doi.org/10.1111/j.1365-](https://doi.org/10.1111/j.1365-246X.2006.02865.x)
- 911 [246X.2006.02865.x](https://doi.org/10.1111/j.1365-246X.2006.02865.x)
- 912 Global Volcanism Program. (2013). *Volcanoes of the World*, v.4.9.3 (01 Feb 2021). Venzke, E (ed.). Smithsonian

- 913 Institution. Downloaded 10 Feb 2021. <https://doi.org/10.5479/si.GVP.VOTW4-2013>
- 914 Glorot, X., Bordes, A., & Bengio, Y. (2011). Domain adaptation for large-scale sentiment classification: A deep  
915 learning approach. In *Proceedings of the 28th International Conference on Machine Learning, ICML 2011*  
916 (pp. 513–520). Bellevue, WA, USA.
- 917 Goitom, B. (2017). The Nabro volcano, tectonic framework and seismic hazard assessment of Eritrea (doctoral  
918 thesis). University of Bristol.
- 919 Goitom, B., Oppenheimer, C., Hammond, J. O. S., Grandin, R., Barnie, T., Donovan, A., ... Berhe, S. (2015). First  
920 recorded eruption of Nabro volcano , Eritrea , 2011. *Bulletin of Volcanology*, 77(85).  
921 <https://doi.org/10.1007/s00445-015-0966-3>
- 922 Hamlyn, J. E., Keir, D., Wright, T. J., Neuberg, J. W., Goitom, B., Hammond, J. O. S., ... Grandin, R. (2014).  
923 Seismicity and subsidence following the 2011 Nabro eruption, Eritrea: Insights into the plumbing system of an  
924 off-rift volcano. *Journal of Geophysical Research: Solid Earth*, 119, 8267–8282.  
925 <https://doi.org/10.1002/2014JB011395>
- 926 Hansen, S. M., & Schmandt, B. (2015). Automated detection and location of microseismicity at Mount St. Helens  
927 with a large-N geophone array. *Geophysical Research Letters*, 42(18), 7390–7397.  
928 <https://doi.org/10.1002/2015GL064848>
- 929 He, T., Zhang, Z., Zhang, H., Zhang, Z., Xie, J., & Li, M. (2019). Bag of Tricks for Image Classification with  
930 Convolutional Neural Networks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*  
931 *(CVPR)* (pp. 558–567). Long Beach, CA, USA. <https://doi.org/10.1109/CVPR.2019.00065>
- 932 Hibert, C., Provost, F., Malet, J. P., Maggi, A., Stumpf, A., & Ferrazzini, V. (2017). Automatic identification of  
933 rockfalls and volcano-tectonic earthquakes at the Piton de la Fournaise volcano using a Random Forest  
934 algorithm. *Journal of Volcanology and Geothermal Research*, 340, 130–142.  
935 <https://doi.org/10.1016/j.jvolgeores.2017.04.015>
- 936 Huot, F., Biondi, B., & Beroza, G. C. (2018). Jump-starting neural network training for seismic problems. In *SEG*  
937 *Technical Program Expanded Abstracts* (pp. 2191–2195).
- 938 Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal  
939 covariate shift. In *32nd International Conference on Machine Learning (ICML)* (pp. 448–456). Lille, France.  
940 <https://arxiv.org/abs/1502.03167>
- 941 Kingma, D. P., & Ba, J. (2014). Adam: A Method for Stochastic Optimization. In *3rd International Conference on*  
942 *Learning Representations (ICLR)* (pp. 1–15). <http://arxiv.org/abs/1412.6980>
- 943 Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., ... Hadsell, R. (2017).  
944 Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*,  
945 114(13), 3521–3526. <https://doi.org/10.1073/pnas.1611835114>
- 946 Krischer, L., Megies, T., Barsch, R., Beyreuther, M., Lecocq, T., Caudron, C., & Wassermann, J. (2015). ObsPy: A  
947 bridge for seismology into the scientific Python ecosystem. *Computational Science and Discovery*, 8, 014003.  
948 <https://doi.org/10.1088/1749-4699/8/1/014003>
- 949 Lahr, J. C., Chouet, B. A., Stephens, C. D., Power, J. A., & Page, R. A. (1994). Earthquake classification, location,

- 950 and error analysis in a volcanic environment: implications for the magmatic system of the 1989–1990  
 951 eruptions at redoubt volcano, Alaska. *Journal of Volcanology and Geothermal Research*, 62(1–4), 137–151.  
 952 [https://doi.org/10.1016/0377-0273\(94\)90031-0](https://doi.org/10.1016/0377-0273(94)90031-0)
- 953 Lapins, S. (2021). Python notebooks to accompany paper ‘A Little Data goes a Long Way: Automating Seismic  
 954 Phase Arrival Picking at Nabro Volcano with Transfer Learning’ (Version v1.0.0) [Archived GitHub  
 955 repository]. Zenodo. <https://doi.org/10.5281/zenodo.4558121>
- 956 Lapins, S., Goitom, B., Kendall, J.-M., Werner, M. J., Cashman, K. V. & Hammond, J. O. S. (2021). Training,  
 957 Validation and Test Sets for paper ‘A Little Data goes a Long Way: Automating Seismic Phase Arrival  
 958 Picking at Nabro Volcano with Transfer Learning’ [Dataset]. Zenodo. <https://doi.org/10.5281/zenodo.4498549>
- 959 Lapins, S., Roman, D. C., Rougier, J., De Angelis, S., Cashman, K. V, & Kendall, J.-M. (2020). An examination of  
 960 the continuous wavelet transform for volcano-seismic spectral analysis. *Journal of Volcanology and  
 961 Geothermal Research*, 389, 106728. <https://doi.org/10.1016/j.jvolgeores.2019.106728>
- 962 Lara, F., Lara-Cueva, R., Larco, J. C., Carrera, E. V, & León, R. (2020). A deep learning approach for automatic  
 963 recognition of seismo-volcanic events at the Cotopaxi volcano. *Journal of Volcanology and Geothermal  
 964 Research*, (xxxx), 107142. <https://doi.org/10.1016/j.jvolgeores.2020.107142>
- 965 Lengliné, O., Duputel, Z., & Ferrazzini, V. (2016). Uncovering the hidden signature of a magmatic recharge at Piton  
 966 de la Fournaise volcano using small earthquakes. *Geophysical Research Letters*, 43(9), 4255–4262.  
 967 <https://doi.org/10.1002/2016GL068383>
- 968 Li, H., Xu, Z., Taylor, G., Studer, C., & Goldstein, T. (2017). Visualizing the Loss Landscape of Neural Nets.  
 969 *Advances in Neural Information Processing Systems (NIPS)*, 6389–6399. <https://arxiv.org/abs/1712.09913>
- 970 Li, W., Duan, L., Xu, D., & Tsang, I. W. (2014). Learning with augmented features for supervised and semi-  
 971 supervised heterogeneous domain adaptation. *IEEE Transactions on Pattern Analysis and Machine  
 972 Intelligence*, 36(6), 1134–1148. <https://doi.org/10.1109/TPAMI.2013.167>
- 973 Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollar, P. (2017). Focal Loss for Dense Object Detection. In 2017  
 974 *IEEE International Conference on Computer Vision (ICCV)* (pp. 2999–3007). Venice, Italy: IEEE.  
 975 <https://doi.org/10.1109/ICCV.2017.324>
- 976 Lomax, A., Virieux, J., Volant, P., & Berge, C. (2000). Probabilistic earthquake location in 3D and layered models:  
 977 Introduction of a Metropolis-Gibbs method and comparison with linear locations. In C. H. Thurber & N.  
 978 Rabinowitz (Eds.), *Advances in Seismic Event Location* (pp. 101–134). Amsterdam: Kluwer.
- 979 Maqsood, M., Nazir, F., Khan, U., Aadil, F., Jamal, H., Mehmood, I., & Song, O. (2019). Transfer Learning  
 980 Assisted Classification and Detection of Alzheimer’s Disease Stages Using 3D MRI Scans. *Sensors*, 19(11),  
 981 2645. <https://doi.org/10.3390/s19112645>
- 982 McNutt, S. R., & Roman, D. C. (2015). Volcanic Seismicity. In *The Encyclopedia of Volcanoes* (2nd Edition, pp.  
 983 1011–1034). Elsevier. <https://doi.org/10.1016/B978-0-12-385938-9.00059-6>
- 984 Megies, T., Beyreuther, M., Barsch, R., Krischer, L., & Wassermann, J. (2011). ObsPy - what can it do for data  
 985 centers and observatories? *Annals of Geophysics*, 54(1), 47–58. <https://doi.org/10.4401/ag-4838>
- 986 Mousavi, S. M., Zhu, W., Sheng, Y., & Beroza, G. C. (2019). CRED: A Deep Residual Network of Convolutional

- 987 and Recurrent Units for Earthquake Signal Detection. *Scientific Reports*, 9, 10267.  
 988 <https://doi.org/10.1038/s41598-019-45748-1>
- 989 Nair, V., & Hinton, G. E. (2010). Rectified Linear Units Improve Restricted Boltzmann Machines. In *Proceedings*  
 990 *of the 27th International Conference on Machine Learning* (pp. 807–814). Haifa, Israel.  
 991 <https://doi.org/10.5555/3104322.3104425>
- 992 Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data*  
 993 *Engineering*, 22(10), 1345–1359. <https://doi.org/10.1109/TKDE.2009.191>
- 994 Razavian, A. S., Azizpour, H., Sullivan, J., & Carlsson, S. (2014). CNN features off-the-shelf: An astounding  
 995 baseline for recognition. In *2014 IEEE Computer Society Conference on Computer Vision and Pattern*  
 996 *Recognition Workshops* (pp. 512–519). Columbus, OH, USA: IEEE.  
 997 <https://doi.org/10.1109/CVPRW.2014.131>
- 998 Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional Networks for Biomedical Image  
 999 Segmentation. In N. Navab, J. Hornegger, W. M. Wells, & A. F. Frangi (Eds.), *Medical Image Computing and*  
 1000 *Computer-Assisted Intervention (MICCAI) 2015, Part III, Lecture Notes in Computer Science* (Vol. 9351, pp.  
 1001 234–241). Springer International Publishing. [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
- 1002 Ross, Z. E., Meier, M., Hauksson, E., & Heaton, T. H. (2018). Generalized Seismic Phase Detection with Deep  
 1003 Learning. *Bulletin of the Seismological Society of America*, 108(5A), 2894–2901.  
 1004 <https://doi.org/10.1785/0120180080>
- 1005 Shelly, D. R., Beroza, G. C., & Ide, S. (2007). Non-volcanic tremor and low-frequency earthquake swarms. *Nature*,  
 1006 446(7133), 305–307. <https://doi.org/10.1038/nature05666>
- 1007 Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on Image Data Augmentation for Deep Learning. *Journal of*  
 1008 *Big Data*, 6(1). <https://doi.org/10.1186/s40537-019-0197-0>
- 1009 Sun, B., Feng, J., & Saenko, K. (2016). Return of frustratingly easy domain adaptation. In *30th AAAI Conference on*  
 1010 *Artificial Intelligence (AAAI)* (pp. 2058–2065). <https://arxiv.org/abs/1511.05547>
- 1011 Sun, C., Shrivastava, A., Singh, S., & Gupta, A. (2017). Revisiting Unreasonable Effectiveness of Data in Deep  
 1012 Learning Era. In *2017 IEEE International Conference on Computer Vision (ICCV)* (pp. 843–852). Venice,  
 1013 2017. <https://doi.org/10.1109/ICCV.2017.97>
- 1014 Theys, N., Campion, R., Clarisse, L., Brenot, H., van Gent, J., Dils, B., ... Ferrucci, F. (2013). Volcanic SO<sub>2</sub> fluxes  
 1015 derived from satellite data: a survey using OMI, GOME-2, IASI and MODIS. *Atmospheric Chemistry and*  
 1016 *Physics*, 13, 5945–5968. <https://doi.org/10.5194/acp-13-5945-2013>
- 1017 Titos, M., Bueno, A., García, L., Benítez, C., & Segura, J. C. (2020). Classification of Isolated Volcano-Seismic  
 1018 Events Based on Inductive Transfer Learning. *IEEE Geoscience and Remote Sensing Letters*, 17(5), 869–873.  
 1019 <https://doi.org/10.1109/LGRS.2019.2931063>
- 1020 Tompson, J., Goroshin, R., Jain, A., LeCun, Y., & Bregler, C. (2015). Efficient Object Localization Using  
 1021 Convolutional Networks. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp.  
 1022 648–656). Boston, MA, USA: IEEE. <https://doi.org/10.1109/CVPR.2015.7298664>
- 1023 Tran, K. T., Griffin, L. D., Chetty, K., & Vishwakarma, S. (2020). Transfer learning from audio deep learning

- 1024 models for micro-Doppler activity recognition. In *2020 IEEE International Radar Conference, (RADAR)* (pp.  
1025 584–589). Washington, DC, USA. <https://doi.org/10.1109/RADAR42522.2020.9114643>
- 1026 Tzeng, E., Hoffman, J., Darrell, T., & Saenko, K. (2015). Simultaneous Deep Transfer Across Domains and Tasks.  
1027 In *2015 IEEE International Conference on Computer Vision (ICCV)* (pp. 4068–4076). Santiago, Chile: IEEE.  
1028 <https://doi.org/10.1109/ICCV.2015.463>
- 1029 van den Ende, M. P. A., & Ampuero, J. P. (2020). Automated Seismic Source Characterization Using Deep Graph  
1030 Neural Networks. *Geophysical Research Letters*, *47*(17), 1–11. <https://doi.org/10.1029/2020GL088690>
- 1031 van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., ... Kavukcuoglu, K. (2016).  
1032 WaveNet: A Generative Model for Raw Audio. <https://arxiv.org/abs/1609.03499>
- 1033 Withers, M., Aster, R., Young, C., Beiriger, J., Harris, M., Moore, S., & Trujillo, J. (1998). A comparison of select  
1034 trigger algorithms for automated global seismic phase and event detection. *Bulletin of the Seismological*  
1035 *Society of America*, *88*(1), 95–106.
- 1036 Woollam, J., Rietbrock, A., Bueno, A., & De Angelis, S. (2019). Convolutional Neural Network for Seismic Phase  
1037 Classification, Performance Demonstration over a Local Seismic Network. *Seismological Research Letters*, 1–  
1038 12. <https://doi.org/10.1785/0220180312>
- 1039 Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). How transferable are features in deep neural networks?  
1040 *Advances in Neural Information Processing Systems*, *27*. <https://arxiv.org/abs/1411.1792>
- 1041 Yu, F., & Koltun, V. (2016). Multi-scale context aggregation by dilated convolutions. In *4th International*  
1042 *Conference on Learning Representations (ICLR)*. <https://arxiv.org/abs/1511.07122>
- 1043 Zamir, A. R., Sax, A., Shen, W., Guibas, L., Malik, J., & Savarese, S. (2018). Taskonomy: Disentangling Task  
1044 Transfer Learning. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 3712–  
1045 3722). Salt Lake City, UT, USA: IEEE. <https://doi.org/10.1109/CVPR.2018.00391>
- 1046 Zech, J. R., Badgeley, M. A., Liu, M., Costa, A. B., Titano, J. J., & Oermann, E. K. (2018). Variable generalization  
1047 performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study.  
1048 *PLOS Medicine*, *15*(11), 1–17. <https://doi.org/10.1371/journal.pmed.1002683>
- 1049 Zhu, W., & Beroza, G. C. (2019). PhaseNet: a deep-neural-network-based seismic arrival-time picking method.  
1050 *Geophysical Journal International*, *216*(1), 261–273. <https://doi.org/10.1093/gji/ggy423>
- 1051 Zhuang, F., Cheng, X., Luo, P., Pan, S. J., & He, Q. (2015). Supervised representation learning: Transfer learning  
1052 with deep autoencoders. In *IJCAI International Joint Conference on Artificial Intelligence* (pp. 4119–4125).  
1053 Buenos Aires, Argentina.
- 1054 Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., ... He, Q. (2020). A Comprehensive Survey on Transfer  
1055 Learning. *Proceedings of the IEEE*, 1–34. <https://doi.org/10.1109/JPROC.2020.3004555>