



# A Dynamic Extreme Value Model with Application to Volcanic Eruption Forecasting

Michele Nguyen<sup>1</sup> · Almut E. D. Veraart<sup>2</sup> · Benoit Taisne<sup>1,3</sup> ·  
Chiou Ting Tan<sup>3</sup> · David Lallemand<sup>1,3</sup>

Received: 27 September 2022 / Accepted: 17 September 2023 / Published online: 30 October 2023  
© The Author(s) 2023

## Abstract

Extreme events such as natural and economic disasters leave lasting impacts on society and motivate the analysis of extremes from data. While classical statistical tools based on Gaussian distributions focus on average behaviour and can lead to persistent biases when estimating extremes, extreme value theory (EVT) provides the mathematical foundations to accurately characterise extremes. This motivates the development of extreme value models for extreme event forecasting. In this paper, a dynamic extreme value model is proposed for forecasting volcanic eruptions. This is inspired by one recently introduced for financial risk forecasting with high-frequency data. Using a case study of the Piton de la Fournaise volcano, it is shown that the modelling framework is widely applicable, flexible and holds strong promise for natural hazard forecasting. The value of using EVT-informed thresholds to identify and model extreme events is shown through forecast performance, and considerations to account for the range of observed events are discussed.

**Keywords** Extreme value theory · High-frequency data · Forecasting · Seismic data

## 1 Introduction

Natural hazards and extreme financial loss can be seen as extreme events (i.e., events that have low probabilities of occurring under normal circumstances). Because of their imbalanced number of occurrences, one usually has more data on common, non-extreme events than on extreme events. However, it can be shown that fitting a

---

✉ Michele Nguyen  
michele.nguyen@ntu.edu.sg

<sup>1</sup> Asian School of the Environment, Nanyang Technological University, Singapore 637459, Singapore

<sup>2</sup> Department of Mathematics, Imperial College London, London SW7 2AZ, UK

<sup>3</sup> Earth Observatory of Singapore, Nanyang Technological University, Singapore 639798, Singapore

distribution to all the data via classical statistical methods leads to a reasonable fit to the bulk of the data at the expense of a poor fit to the tails in which the extremes lie (Ribatet 2016). Extreme value theory (EVT) and its corresponding models seek to remedy this disparity by offering guidance on when an extreme regime kicks in and how it can be modelled. For example, under known conditions, the distribution of excesses above a fixed threshold can be shown to converge to a generalised Pareto distribution (GPD) as the threshold increases (Embrechts et al. 2013; Beirlant et al. 2004). By knowing this asymptotic distribution, one can conduct model checks and generate sensible estimates of extremal behaviour.

EVT is increasingly used in financial applications and has been shown to give more accurate tail-risk predictions (Danielsson and De Vries 1997; Longin 2000). To address the fact that there is time dependence in financial returns which goes against the traditional assumption of independence (Diebold et al. 1998), Bee et al. (2019) propose a dynamic extreme value model which uses high-frequency realised measures of the daily asset price variation as covariates to model the probability of exceeding a high threshold and the size of the excesses. Since the realised variation is time-varying, the estimates of threshold exceedance and excesses are also time-varying.

The extreme value model developed by Bee et al. (2019) can be adapted to wider settings. In this paper, this is demonstrated by adapting the model for natural hazard forecasting, specifically for volcanic eruptions. Just as how extreme loss can be defined as an exceedance over some financial threshold, extreme volcanic activity can be defined as the threshold exceedance of some energy index.

The contributions of this model to the eruption forecasting literature are manifold. In addition to the short overview of existing methods provided in the Supplementary Information, the key differences and contributions of the proposed model are highlighted here. The dynamic extreme value model draws on techniques from time series analysis, EVT and machine learning, of which full and varied potential for short-term eruption forecasting has yet to be realised (Malfante et al. 2018; Carniel and Guzmán 2020; Whitehead and Bebbington 2021). To adapt the model from its original financial context to wider settings, choices need to be made on the following:

- (i) a suitable index to compute threshold exceedances;
- (ii) the look-ahead window or forecast horizon;
- (iii) the auxiliary information or covariates used to inform future behaviour; and
- (iv) the time periods to compute these from (i.e., the covariate window).

In the seismic context, threshold exceedance for event detection is synonymous with first-arrival picking algorithms such as the classical short-time average over long-time average (STA/LTA) method and a recently introduced method using trace envelopes (Withers et al. 1998; Trnkoczy 1999; Al-Mashhor et al. 2019). The trace envelope can be seen as the instantaneous amplitude of a seismic trace and related to the amount of energy in the signal (Ktonas and Papp 1980; Taner 2001). It is computed by taking the square root of the sum of the squared real and imaginary parts of the seismic trace. In this paper, trace envelopes will be used as eruption indices from which to take threshold exceedances. These exceedances would hopefully relate to extreme regimes leading up to volcanic eruptions and can be forecasted using covariates. Note that the exceedance of an eruption index is modelled rather than a physical monitoring

signal because the latter does not necessarily have a monotonic relationship with the hazard. In addition, since different frequency bands within the seismic signal represent different physical phenomena (Bormann et al. 2013; Salvage et al. 2019), envelopes of frequency-filtered data will be considered.

While existing techniques like the failure forecast method estimate the eruption onset time directly, others including event trees define a look-forward window within which the probability of an eruption occurring is estimated. The latter approach is taken here. In particular, for illustration, 1-h-ahead forecasts are produced to complement other longer-term forecasts. Although a longer lead time allows for more time to make emergency management decisions, notify the public and implement evacuations (Wild et al. 2021), forecasts are typically more accurate when made closer to the actual eruption time due to temporal divergence at larger lags (see for example, Sugihara and May 1990).

Eruption forecasting methods such as event trees, belief networks and process/source models presuppose precursors or associations between source mechanisms and time series signals (Brennguier et al. 2008). In contrast, the proposed methodology selects combinations of covariates that could represent precursors for any volcano and type of eruption if the model is trained on corresponding data. Specifically, covariates inspired by machine learning classification algorithms for seismic signals (Malfante et al. 2018) are tested. By combining these covariates, different aspects of the seismic data and relationships across different frequency bands are represented. An objective stepwise selection procedure is then used to determine which covariates are more informative for forecasting eruptions.

The rest of this paper is organised as follows. The key components of the dynamic extreme value model introduced by Bee et al. (2019) are outlined in Sect. 2. In Sect. 3, the case study, the Piton de la Fournaise volcano, is introduced to illustrate how the model can be adapted for eruption forecasting. By comparing the effect of different choices of the threshold on the model fit and training performance, the value of using EVT to guide threshold choice is highlighted in Sect. 4. In Sect. 5, the broader applicability of the method is evaluated by refitting the model using three training event sets and testing the calibrated model on both event and non-event sets. In Sect. 6, the results and future areas for research are discussed. The code used for the analysis is publicly available at <https://github.com/ntu-dasl-sg/dynamic-EV-forecasting>.

## 2 The Dynamic Extreme Value Model

Following Bee et al. (2019), let  $\{Y_t\}_{t=1,\dots,T}$  denote a time series of an index where higher values are associated with extreme events. Based on the selected threshold  $u \in \mathbb{R}$  of the index, exceedances are defined as the binary indicators of whether the index is higher than the threshold and excesses are defined as the numerical values of the exceedances. The conditional probability that the index at time  $t$ ,  $Y_t$ , exceeds  $u$  by some excess  $z > 0$  given prior information available at time  $t$ ,  $\mathcal{F}_{t-1}$ , can be written as

$$P(Y_t > u + z | \mathcal{F}_{t-1}) = P(Y_t > u | \mathcal{F}_{t-1})P(Y_t - u > z | Y_t > u, \mathcal{F}_{t-1})$$

$$= \phi_t | \mathcal{F}_{t-1} \times GPD(\xi_t, \nu_t) | \mathcal{F}_{t-1}. \quad (1)$$

Here,  $\phi_t | \mathcal{F}_{t-1} = P(Y_t > u | \mathcal{F}_{t-1})$  represents a time-varying binomial exceedance probability which can be modelled with a logit function

$$\phi_t | \mathcal{F}_{t-1} = \frac{\exp(\psi_0 + \sum_{i=1}^p \psi_i x_{t-1}^{(i)})}{1 + \exp(\psi_0 + \sum_{i=1}^p \psi_i x_{t-1}^{(i)})}, \quad (2)$$

where  $\mathbf{x}_{t-1} = (x_{t-1}^{(1)}, \dots, x_{t-1}^{(p)})$  denotes a vector of  $p$  covariates from the previous time step and is used to project the future probability. The parameters  $\{\psi_i\}_{i=0, \dots, p}$  can be estimated by maximising the likelihood function

$$\mathcal{L}(\boldsymbol{\psi}; I_t, \mathbf{x}_t) = \prod_{t=l+1}^T \left( \exp(\psi_0 + \sum_{i=1}^p \psi_i x_{t-1}^{(i)}) \right)^{I_t} \times \frac{1}{1 + \exp(\psi_0 + \sum_{i=1}^p \psi_i x_{t-1}^{(i)})}, \quad (3)$$

where  $l$  is the lag at which the covariates  $\mathbf{x}_t$  become available and  $I_t$  is the indicator of an exceedance at time  $t$  (it is equal to 1 if there is an exceedance and 0 otherwise). This is equivalent to using logistic regression to model the probability of threshold exceedance.

In (1), the model for excesses of the threshold is given by a generalised Pareto distribution (GPD)

$$P(Y_t - u > z | Y_t > u, \mathcal{F}_{t-1}) = GPD(\xi_t, \nu_t) | \mathcal{F}_{t-1}. \quad (4)$$

The shape parameter  $\xi_t = \xi$  is estimated assuming a non-time-varying GPD distribution for the excesses and kept constant for stability. To account for the time-varying nature of the excess distribution, the scale parameter is modelled as

$$\nu_t = \exp \left( \kappa_0 + \sum_{i=1}^q \kappa_i x_{t-1}^{(i)} \right). \quad (5)$$

This involves  $q$  covariates,  $(x_{t-1}^{(1)}, \dots, x_{t-1}^{(q)})$  in a log-linear function.

When  $\kappa_i = 0$  for  $i > 0$ , this means that it is sufficient to model the distribution of the excesses statically. From the definition of a GPD with a non-zero shape parameter

$$P(Y_t - u > z | Y_t > u, \mathcal{F}_{t-1}) = \left( 1 + \frac{\xi z}{\exp(\kappa_0 + \sum_{i=1}^q \kappa_i x_{t-1}^{(i)})} \right)^{-1/\xi}, \quad (6)$$

where  $z = y_t - u$  denotes the excess. If  $\xi > 0$ ,  $z \geq 0$  and if  $\xi < 0$ ,  $0 \leq z \leq -\exp(\kappa_0 + \sum_{i=1}^q \kappa_i x_{t-1}^{(i)})/\xi$  (i.e., the excesses have an upper bound). The parameters

$\{\kappa_i\}_{i=1,\dots,q}$  can be estimated by maximising the likelihood function

$$\begin{aligned} &\mathcal{L}(\kappa, \xi; \mathbf{z}_t, \mathbf{x}_t) \\ &= \prod_{t=l+1}^T \left( \frac{1}{\exp(\kappa_0 + \sum_{i=1}^q \kappa_i x_{t-1}^{(i)})} \times \left[ \left( 1 + \frac{\xi z_t}{\exp(\kappa_0 + \sum_{i=1}^q \kappa_i x_{t-1}^{(i)})} \right)^{-1/\xi - 1} \right]_+ \right)^{I_t}, \end{aligned} \tag{7}$$

where  $[x]_+ = \max(0, x)$ , and the time subscript is added to the excesses,  $\mathbf{z}$ , to denote their temporal indices. Henceforth, the maximisation of (7) to estimate  $\{\kappa_i\}_{i=1,\dots,q}$  will be referred to as GPD regression. When the estimated shape parameter  $\hat{\xi}$  is not significantly different from zero, it is set to zero and an exponential regression is used instead.

### 3 Case Study: Piton de la Fournaise Volcano

#### 3.1 Data

The dynamic extreme value model is used to forecast eruptions at the Piton de la Fournaise volcano. Situated on La Réunion Island, Piton de la Fournaise is one of the most active basaltic volcanoes, with an average of one eruption every 10 months (Roult et al. 2012). In addition to the existing seismic monitoring stations, 15 broadband stations were installed on the volcano as part of the Understanding Volcanic Processes (UnderVolc) project in 2009–2010 (Taisne et al. 2011). Figure 1 shows their locations together with other stations of the UnderVolc project and Station Sismologique de Riviere de l’Est (RER) of the Geoscope seismic network. The data collected are available at [https://www.fdsn.org/networks/detail/YA\\_2009/](https://www.fdsn.org/networks/detail/YA_2009/) and cover the years 2009 to 2011. Between 2009 and 2011, five eruptions were recorded at Piton de la Fournaise.

To avoid issues arising from varying energy index magnitudes at different recording stations due to their respective distances from the eruption centres, data from one station were used to train the model. This is the UV05 station, which was the closest available station to the January 2010 eruptive centre (Journeau et al. 2020) and which is approximately 571 m from the recorded starting point of the first fissure (Roult et al. 2012).


Data from four out of the five recorded eruptions were used in the analysis. Data from the December 9, 2010 eruption were not used for training or testing because the location of the eruption (Flank N) is relatively far away from the rest of the event locations and hence could skew the results.

Table 1 shows the key characteristics of the four eruption events as outlined by Roult et al. (2012). To evaluate the forecast performance in Sect. 5, three of the events will be used for training the model, and the other event will be used for testing. Three non-events are used for further validation. The non-event dates were chosen to be roughly halfway between the selected four events and represent quiet periods for which one should not expect threshold exceedances of the eruption indices.

**Table 1** Events and non-events used to train and test the forecasting model

	Data period	Eruption			Duration	
		Date	Location	Type	SC	SS
Training event 1	04–06/11/2009	05/11/2009	Flank S SE	E-L	1 h 30 min	30 min
Training event 2	13–14/12/2009	14/12/2009	Flank S/SW	E-L	1 h 10 min	32 min
Training event 3	01–03/01/2010	02/01/2010	Dolomieu W	E-C	2 h 30 min	42 min
Test event 1	13–14/10/2010	14/10/2010	Flank S	E-L	5 h 30 min	40 min
Test non-event 1	30/11/2009	–	–	–	–	–
Test non-event 2	22–23/12/2009	–	–	–	–	–
Test non-event 3	08–09/05/2010	–	–	–	–	–

The dates are given in dd/mm/yyyy. E-C refers to an eruption that remains inside the summit craters (Dolomieu or Bory) and E-L refers to an eruption that starts anywhere else outside of the summit craters. SS refers to seismic swarm, while SC refers to seismic crisis

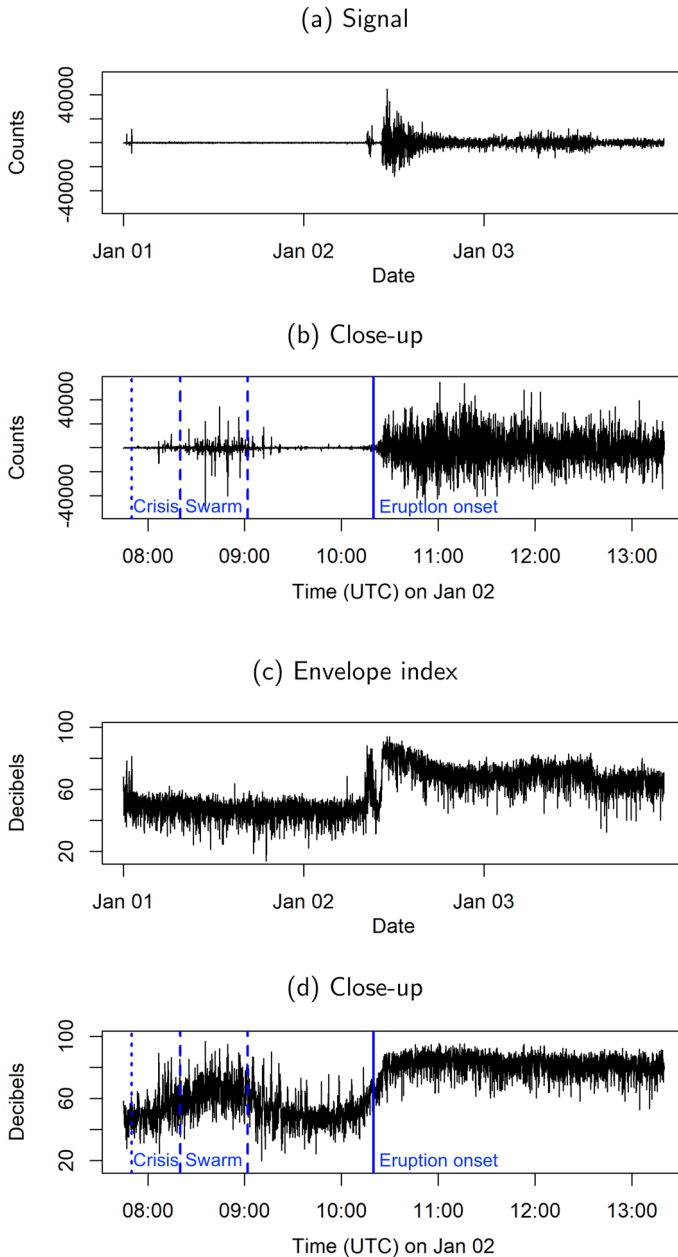


**Fig. 1** Map of the Piton de la Fournaise volcano, corresponding to Fig. 2 of Seydoux et al. (2016). The locations of the Understanding Volcanic Processes stations are represented by red triangles, and that of Station Sismologique de Riviere de l'Est (RER) of the Geoscope seismic network is represented by the blue triangle

### 3.2 Illustration of Method

To illustrate the use of the dynamic extreme value model for eruption forecasting, the method is first applied to the January 2010 event data (training event 3) since it is best documented out of the chosen events. Figure 2a shows the raw seismic signal for the 1–5 Hz frequency-filtered data. This is given in counts which corresponds to the number read off the seismometer (i.e., the voltage read from a sensor). To scale it to physical units (e.g., m/s), a number of factors including the frequency of motion being measured and the calibrated zero-point of the instrument would need to be considered. Since only a scaling factor is involved, the same modelling results would be obtained when the raw seismic signal is used in counts or in physical units.

Figure 2b focuses on January 2 when the eruption occurred, and the dotted, dashed and bold blue vertical lines denote the start and end times of the seismic crisis, swarm and eruption onset, respectively. As documented by Roult et al. (2012), a seismic crisis took place from 07:50 local time. This means that there is an increase in seismic events such that this exceeds background seismicity. Between 08:10 and 09:02, a seismic swarm was recorded. This corresponds to when seismic events overlap and become hard to differentiate (McNutt and Roman 2015) and is reflected in the seismic signal as larger fluctuations in the readings between the dashed blue vertical lines. The swarm was followed by a relatively quiet phase that directly preceded the onset of the eruption at about 10:20 (the eruption is indicated by the continuous seismic tremor in the figure). The whole eruption was estimated to have lasted 9.6 days, ending at 00:05 on January 12, 2010.



**Fig. 2** **a** Seismic signal for the 1–5 Hz frequency band for the January 2010 eruption event with **b** a close-up look at the period near the time of eruption on January 2. **c** The corresponding envelope index in decibels (dB), with **d** a zoom into the same time period on January 2. The dotted, dashed and bold blue vertical lines denote the start of the seismic crisis, the start and end of the seismic swarm and the eruption onset respectively. Note that the time series are downsampled such that every 5,183th reading and every 402nd reading are shown for the top and bottom plots, respectively



### 3.2.1 Trace Envelope as an Eruption Index

Before fitting the dynamic extreme value model for eruption forecasting, the eruption index from which to consider threshold exceedances needs to be chosen. Here, the trace envelope  $e_t$  for  $t = 1, \dots, T$  is used. This can be seen as the instantaneous amplitude of the seismic trace  $\mathbf{s} = (s_1, \dots, s_T)$ , and can be computed as follows:

1. First, the discrete Fourier transform (DFT) of  $\mathbf{s}$  is computed for  $t \in \{1, \dots, T\}$  (this is implemented by the R function ‘fft’)

$$f_t = \text{DFT}(\mathbf{s})_T = \sum_{k=1}^T s_k \exp(-2\pi i(k-1)(t-1)/T), \tag{8}$$

where  $T$  is the length of the seismic trace. Set  $\mathbf{f} = (f_1, \dots, f_T)$ .

2. Next, the complex Hilbert fast Fourier transform (FFT) of  $\mathbf{s}$  is computed as

$$H_t = \text{IFT}(\mathbf{fh})/T, \tag{9}$$

where the length- $T$  series  $\mathbf{h} = (1, 2, 2, \dots, 2, 2, 1)$  if  $T$  is even and  $(1, 2, 2, \dots, 2, 2, 2)$  if  $T$  is odd and the inverse Fourier transform (IFT) is defined as

$$\text{IFT}(\mathbf{fh})_t = \sum_{k=1}^T f_k h_k \exp(2\pi i(k-1)(t-1)/T), \tag{10}$$

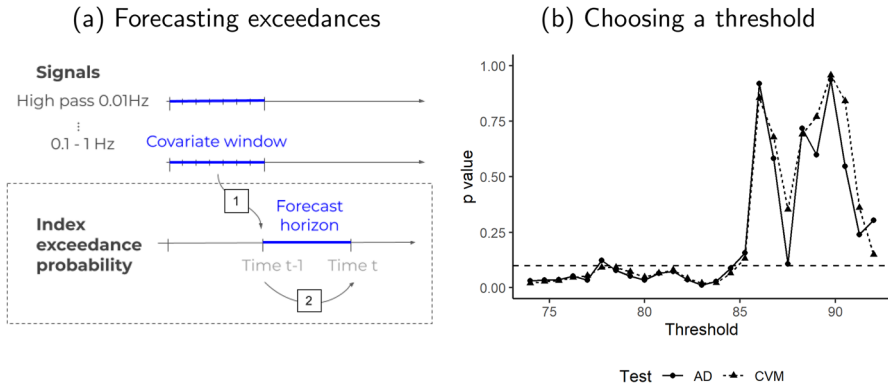
for  $t \in \{1, \dots, T\}$ .

3. For  $t \in \{1, \dots, T\}$ , the trace envelope of the seismic trace  $\mathbf{s}$  is defined as

$$e_t = \sqrt{\text{Re}^2(H_t) + \text{Im}^2(H_t)} = \text{Mod}(H_t). \tag{11}$$

The above steps are in line with those used within the R function ‘envelope’ in the IRISseismic package.

In Al-Mashhor et al. (2019), the first arrival travel time picking algorithm was based on the envelope in decibels. Similarly, the trace envelope  $e_t$  is converted to decibels via  $Y_t = 20 \log_{10}(e_t)$  (Moore 1995, p. 11) and  $Y_t$  is used as the eruption index. Figure 2c shows the trace envelope time series computed from the 1–5 Hz frequency-filtered data for the January 2010 event. The 1–5 Hz frequency band is used because it is strongly associated with volcanic/magmatic activity and fluid resonances (Salvage et al. 2019). Results for the 5–15 Hz frequency band are given in the Supplementary Information. In Fig. 2c, it is seen that although there are some spikes at the start of January 1, the envelope index remains relatively low around 50 dB until the recorded seismic crisis on January 2 (represented by the dotted blue vertical line in the bottom plot). In Fig. 2d, the index increases from the time of the seismic crisis to a first peak during the seismic swarm before waning slightly. About 10:00, the index increases steadily to plateau slightly above 80 dB, the timing of which coincides with the recorded eruption onset (represented by the bold blue vertical line in the bottom plot). The relationship between



**Fig. 3** **a** Model conceptual diagram for forecasting exceedances of the 1–5 Hz envelope index. At time  $t - 1$ , the covariates are computed from signals across multiple frequency bands within a past time window (step 1). These are then used to forecast threshold exceedances at time  $t$  (step 2). **b**  $p$  values from the Anderson–Darling (AD) and Cramér–von Mises (CVM) tests for goodness of fit of the excesses to the GPD distribution. The lowest threshold for which the  $p$  value exceeds the 10% significance level is chosen

the increases in the index and the seismic events enables its threshold exceedances to be used to forecast eruptions.

### 3.2.2 Forecast Horizon and Covariate Window

In addition to computing the index from which to take exceedances, choices need to be made on the following:

- (i) The forecast horizon  $\delta_t$ : the time period between the time window where the covariates are computed and the forecast time.
- (ii) The time window within which past observations contribute to the covariates,  $W$ .

For illustration, 1-h-ahead forecasts are made with 1 h of past data to inform the covariates in the model ( $\delta_t = W = 1$  h). This mimics the settings used by Brenguier et al. (2008), Malfante et al. (2018) and Ren et al. (2020). Their analyses involved the classification of 1-h-long signals or the generation of the covariates by scanning a moving window of length 1 h across the seismic signals. This framework is pictured in Fig. 3a, where the time period between  $t - 1$  and  $t$  is  $\delta_t = 1$  h, and the covariate window  $W = 1$  h.

Although the covariates were computed from the original high-frequency (100 Hz) data, forecasts were only produced every 10 s to reduce any unnecessary computational burden. For a practically useful workflow, this should be longer than the time required to compute the covariates from the past hour and use the fitted model to generate the forecast.

### 3.2.3 Threshold Selection

Next, the threshold to define the extreme regime is selected. This will be associated with the covariates to forecast extreme behaviour (i.e., exceedance). Based on the

Anderson–Darling (AD) and Cramér–von Mises (CVM) tests for goodness of fit of the excesses to the GPD distribution (Barder et al. 2018), Fig. 3b suggests that under a significance level of 10%, a threshold of 85 dB is reasonable, since this is the lowest value from the right for which the p-value exceeds the significance level. At this threshold, the null hypothesis of the excesses coming from a GPD distribution is not rejected.

### 3.2.4 Covariates

To forecast threshold exceedances of the eruption index, the covariates suggested by Malfante et al. (2018) are computed for the frequency-filtered data (0.1–1 Hz, 1–5 Hz, 5–15 Hz, 0.1–20 Hz and high pass 0.01 Hz) and their associated trace envelopes. Malfante et al. (2018) introduce three domains of representation of the seismic time series which could be useful: the original temporal domain, the frequency domain where spectral content is obtained via a Fourier transform, and the cepstral domain where the Fourier transform is computed twice to highlight the harmonic properties of a given signal.

For each of these three representations of the seismic traces and their trace envelopes, statistical, entropy and shape descriptor features are computed. These are listed together with their definitions in Table 2. Each covariate or feature tries to capture a particular aspect of the signal within the covariate window. For example, while kurtosis captures the transition between two signals, Shannon entropy describes the distribution of the amplitude levels of a given signal.

A covariate can also take on different meanings depending on the domain it is computed on. For example, the feature ‘i of Central Energy’ which is the time around which the signal energy is centred or the time centroid in the temporal domain, can be interpreted as the fundamental frequency in the frequency domain and the harmonic frequency in the cepstral domain. In addition, while the ratio of the maximum value over the mean value can describe the contrast and relate to the cause of the event in the original temporal domain, in the frequency domain, it describes the spectral richness of the signature and in the cepstral domain, the harmonic content of an observation.

To ensure that the model is not too sensitive towards extreme covariate values, the covariates were transformed before use. Box–Cox analyses were used to select which power or log-transformation was required to make their distributions more similar to Gaussian distributions. After transformation, the covariates were standardised using their mean and standard deviations to be on similar scales. To account for multicollinearity, the covariates were ordered according to the increasing Akaike information criterion (AIC) of their univariate models (for threshold exceedance and excesses separately). Then, the pool of covariates was reduced by removing covariates that had more than 0.6 in absolute correlation to covariates that were deemed more informative than themselves.

### 3.2.5 Regression Models

As outlined in Sect. 2, a logistic regression and a GPD regression were fitted for threshold exceedances and threshold excesses, respectively. The shape parameter of the GPD

**Table 2** Covariates considered for the dynamic extreme value models

Covariate	Definition
Mean	$\mu_s = \frac{1}{n} \sum_{i=1}^n s_i$
Standard deviation	$\sigma_s = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (s_i - \mu_s)^2}$
Skewness	$\frac{1}{n} \sum_{i=1}^n \left( \frac{s_i - \mu_s}{\sigma_s} \right)^3$
Kurtosis	$\frac{1}{n} \sum_{i=1}^n \left( \frac{s_i - \mu_s}{\sigma_s} \right)^4$
Energy	$E = \sum_{i=1}^n s_i^2$
<i>i</i> of central energy	$\bar{i} = \frac{1}{E} \sum_{i=1}^n i \cdot E_i$
Root mean square bandwidth	$B = \sqrt{\frac{1}{E} \sum_{i=1}^n i^2 \cdot E_i - \bar{i}^2}$
Mean skewness	$\sqrt{\frac{\sum_{i=1}^n (i - \bar{i})^3 \cdot E_i}{E \cdot B^3}}$
Mean kurtosis	$\sqrt{\frac{\sum_{i=1}^n (i - \bar{i})^4 \cdot E_i}{E \cdot B^4}}$
Shannon entropy	$-\sum_{i=1}^n p(s_i) \log_2(p(s_i))$
Rate of attack (ROA)	$\max_i \left( \frac{s_i - s_{i-1}}{n} \right)$
Maximum	$\max_i s_i$
Minimum	$\min_i s_i$
Ratio of maximum over mean	$\frac{\max_i s_i}{\mu_s}$

Each is computed for the frequency-filtered time series and their associated trace envelopes in the temporal, frequency and cepstral domains, so that a total of  $14 \times 5 \times 2 \times 3 = 420$  covariates are considered for covariate selection. In the definitions,  $\{s_i : i = 1, \dots, n\}$  refers to the signal within the covariate window,  $E_i = s_i^2$  denotes the energy at index *i*, and  $p(s_i)$  is the empirical probability of amplitude level  $s_i$  within the covariate window

is fixed to the value of the estimate obtained using maximum likelihood estimation for a constant GPD. For the January 2010 training event, this has an asymptotically normal 95% confidence interval of  $(-0.300, -0.160)$  which does not include 0. The negative shape parameter implies that the distribution of the excesses is Pareto type II and lies within the Weibull domain of attraction which contains distributions with short tails (i.e., finite endpoints).

The covariates in the models were chosen by stepwise selection based on AIC. There are three options for stepwise variable selection: forward, backward and bidirectional. The default configuration of the ‘stepAIC’ function in the R package ‘MASS’ is bidirectional (Venables and Ripley 2002). For the logistic regression, a backward search is conducted before considering forward selection. The backward direction is often preferred over the forward direction because the full model and the effect of all candidate variables are considered (Steyerberg 2009; Harrell 2015; Chowdhury and Turin 2020). For the GPD regression, a forward search is used before backward selection to avoid singularity issues that arise from having large numbers of covariates and excess uninformative covariates in the model.

The five steps outlined in Sect. 3.2.1–3.2.5 can be repeated to model threshold exceedances and excesses for the other frequency-filtered envelopes.

### 3.3 Results

The black line in Fig. 4a shows the 1-h ahead probabilistic forecasts for threshold exceedances of the 1–5 Hz envelope. The forecasts are highest for January 2, the day the eruption started. Focusing on January 2 in Fig. 4b, the exceedance probability jumps about 1 h before the recorded eruption onset at 10:20. This means that the fitted forecast model is able to give a warning about 1 h ahead of the eruption.

Similar to Bee et al. (2019), the goodness of fit of the logistic regression is checked using a deviance chi-squared test. The p-value was  $e^{-1,113.53}$ , indicating that the fitted model is significantly different from a null model (a logistic regression with an intercept term but no covariates). The usefulness of the covariates for explaining the temporal dependence in the occurrence of the threshold exceedances can also be seen through the reduction in the autocorrelation of the Pearson residuals in Fig. S1a of the Supplementary Information. In contrast, little temporal dependence was observed for the excess residuals in Fig. S1b. Hence, there was no real benefit of using covariates to inform a dynamic GPD and a constant GPD would have sufficed. As will be seen later, this is threshold-specific: when multiple events are used to train the model in Sect. 5 and the lowest EVT-informed threshold is selected among the training events, there will be autocorrelation in the excess residuals and hence benefits of modelling with a dynamic GPD.

## 4 Value of Extreme Value Theory

EVT was used to select the threshold which defined the exceedances and excesses being forecasted by the dynamic extreme value model. Figure 5a shows the goodness-of-fit plots comparing the modelled and empirical probabilities and return levels when 50% and 100% of the threshold informed by EVT were used. The latter provided a better fit to the model assumptions.


The threshold choice is also important for determining what kind of phenomena is being modelled and what covariates are chosen to best explain it. Figure 5b shows that for the exceedance forecasting of the 1–5 Hz, 0.1–20 Hz and high pass 0.01 Hz envelope indices, the forecast performance, as measured via the area under the curve (AUC), generally improves as the threshold is increased towards that informed by EVT. Hence, EVT has benefits for modelling in terms of both goodness of fit to the data and forecast performance.

## 5 Evaluating Forecast Performance

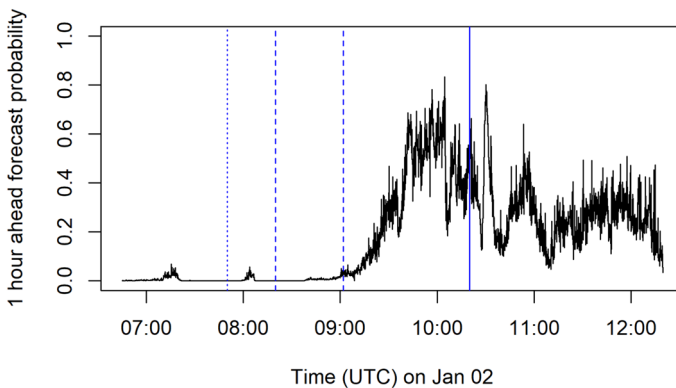
### 5.1 Using Multiple Training Events

To assess the forecast capability of the dynamic extreme value model more formally, the model will be fitted to all three training events, and tested on the remaining test event and the three non-events.

## (a) Training forecasts



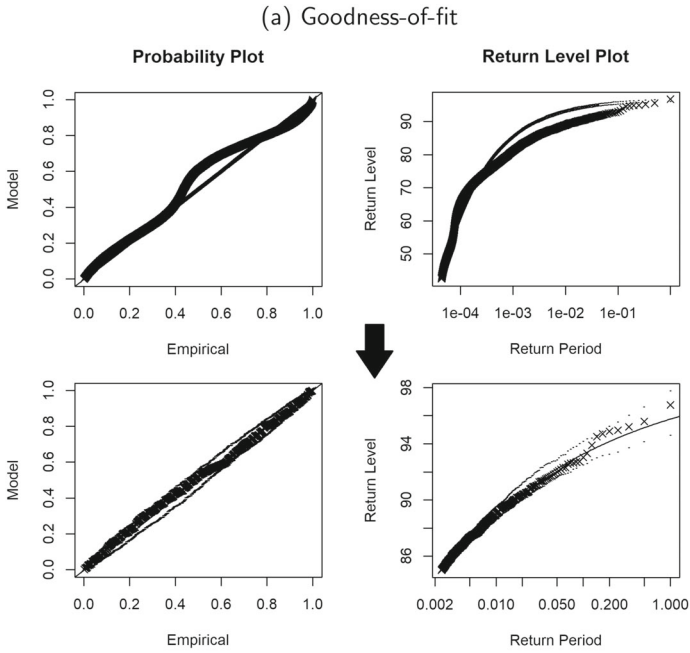
## (b) Zoomed in



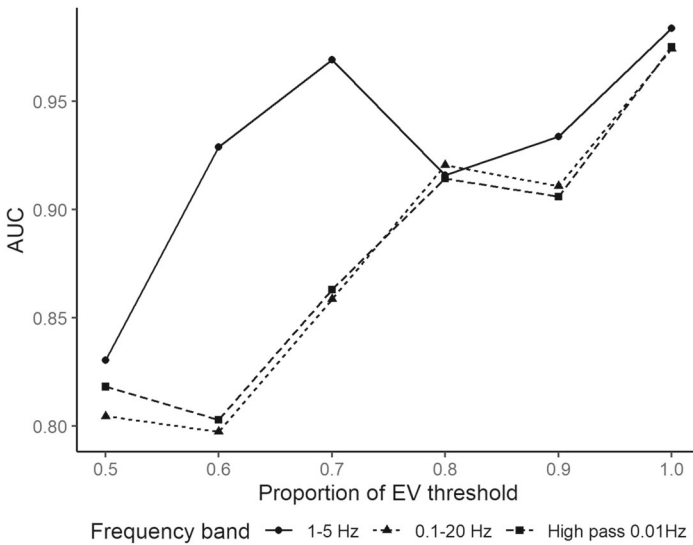
**Fig. 4** **a** Training 1-h-ahead exceedance forecasts based on the logistic regression for the 1–5 Hz envelope index and the January 2010 eruption; **b** zoomed into the period of significant volcanic activity (the dotted, dashed and bold blue vertical lines denote the start and end of the seismic crisis, swarm and eruption onset respectively). Here, the value of the black line at 09:00, for example, indicates the forecasted probability of exceedance for 10:00

There are a few ways to determine a suitable threshold based on data from multiple events. An initial approach might be to simply combine the data across events and use all exceedances to inform the threshold. However, this led to a relatively high threshold estimate of 95 with training event 2 dominating the model fit because there were comparatively more exceedances from event 2 than events 1 and 3 (see Sect. 4 of the Supplementary Information).

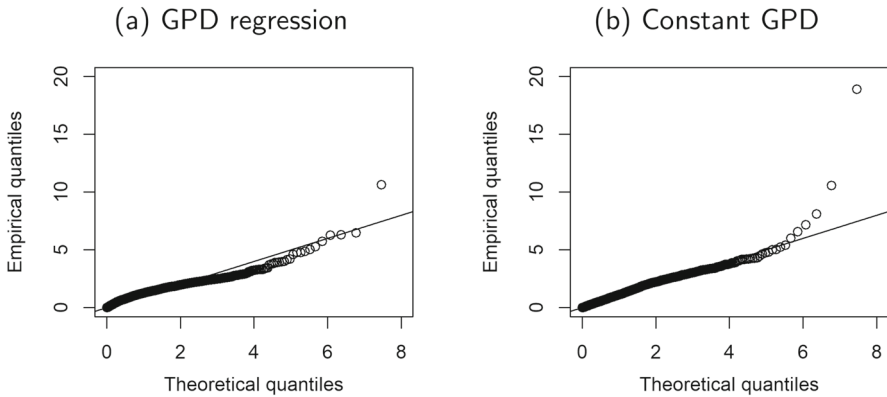
An alternative approach, which will be used, would be to estimate thresholds for the training events separately and use the lowest estimate across events. This will ensure sufficient exceedances to represent each event. For the 1–5 Hz envelope index, the lowest threshold among the three training events was 85 dB, and the estimated GPD



(b) Training performance



**Fig. 5** **a** Improvement in goodness of fit when the threshold used increases from 50% (top) to 100% (bottom) of the value chosen by extreme value theory; **b** improvement in the training performance in terms of area under the curve (AUC) for the 1–5 Hz, 0.1–20 Hz and high pass 0.01 Hz index exceedance models when the threshold increases



**Fig. 6** With multiple training events: comparison of goodness of fit in terms of the empirical and theoretical quantiles of the standardised excesses when **a** the GPD regression is used for the excess distributions instead of **b** treating the excess distribution as static through a constant GPD. Since the extreme quantiles lie closer to the one-to-one diagonal line for the GPD regression, it fits the data better

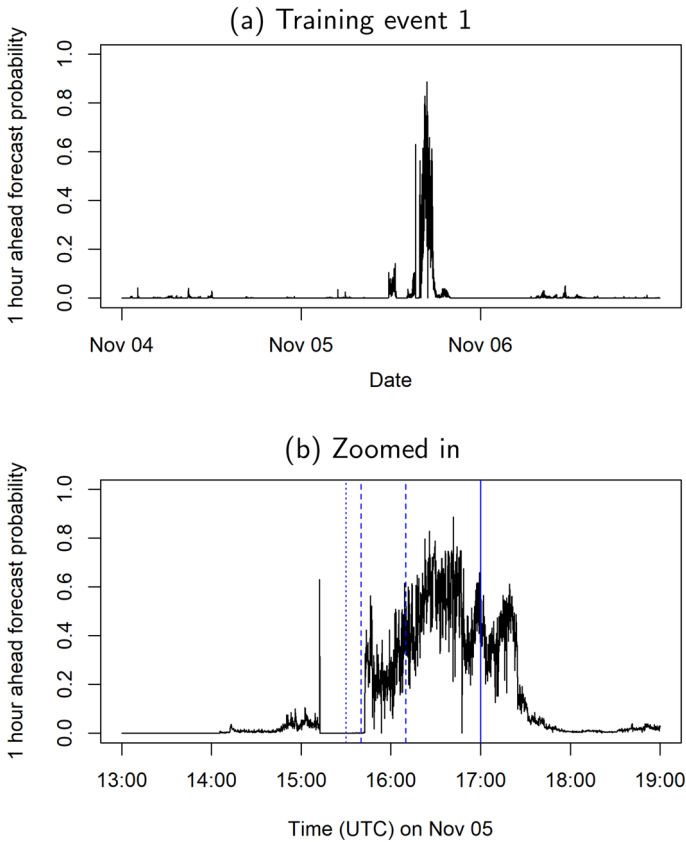
shape parameter is  $\hat{\xi} = -0.125$  with 95% confidence interval  $(-0.146, -0.104)$ . Tables 1 and 2 of the Supplementary Information show the chosen covariates with their transformations and parameter estimates for the fitted logistic and GPD regressions, respectively.

As will be seen in the next section, this threshold choice leads to 1-h-ahead forecast probabilities which increase before the time of eruption for all three training events and the test event while remaining low for the three test non-events. By lowering the threshold to identify extremes across all three training events, there is more benefit of modelling the threshold excesses dynamically because there is autocorrelation in the excess residuals, particularly for training event 2 (see Sect. 2 of the Supplementary Information). In contrast to the initial illustration with just training event 3 in Sect. 3, modelling the excess distribution dynamically leads to better estimation of extreme quantiles, though there is still room for improvement. This is illustrated in Fig. 6.

## 5.2 Training and Test Performance

After fixing the threshold for which to model exceedances, the dynamic extreme value model is fitted to data from the three training events. Figures 7, 8 and 9 shows that apart from an outlier on the first day of training event 3, the forecast probabilities remain low (e.g., below 0.3) for all three events until the time of their recorded seismic events. For training event 1 (referring to Fig. 7b), sustained high forecast probabilities start during the seismic swarm (between the dashed vertical lines), slightly more than 1 h before the recorded eruption at 17:00. For event 2, 1-h-ahead eruption warnings can also be made as the forecast probabilities begin to take high values during the seismic swarm, about 1 h before the recorded eruption at 14:40 (see Fig. 8b). For event 3, the forecast probabilities gradually increase from the time of the seismic swarm around 08:30 before jumping up to a higher plateau about an hour before the recorded eruption onset at 10:20 as shown in Fig. 9b.

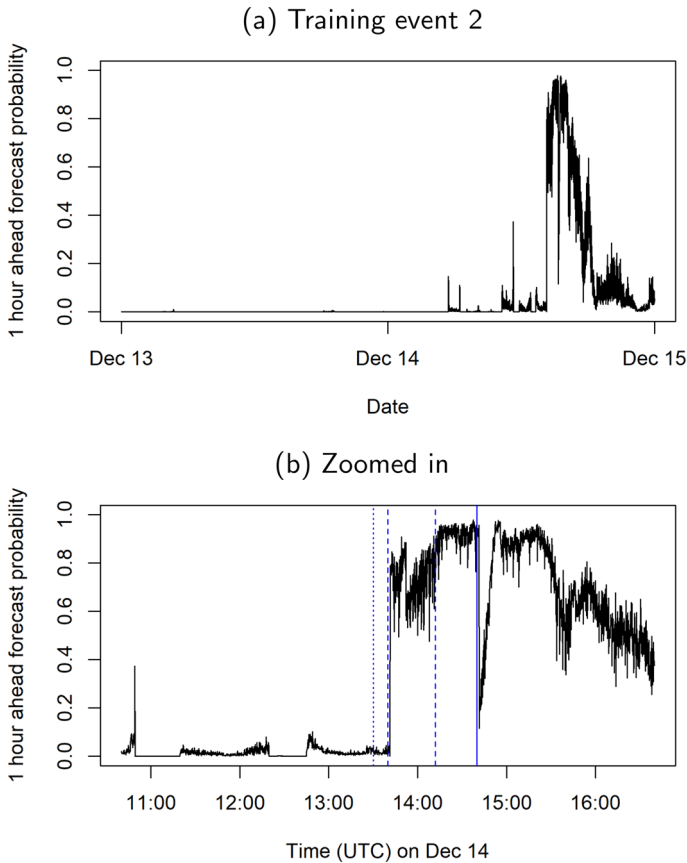




**Fig. 7** Training event 1: 1-h-ahead threshold exceedance forecasts using the lowest threshold estimated across events. The dotted, dashed and bold blue vertical lines denote the times of the seismic crises, start and end of the seismic swarms and the eruption onset respectively

The features of the forecast probabilities, namely the sharp jumps from near-zero for training events 1 and 2, the gradual increase for training event 3, the presence of outliers and the tendency to increase 1 h before the recorded eruption onsets, stem from the chosen covariates of the logistic regression. As can be inferred from their high coefficient estimates in Table 1 of the Supplementary Information, the logistic regression for threshold exceedance has three covariates that contribute to forecast probabilities more than the others: 0.1–20 Hz cepstral kurtosis, 0.1–20 Hz cepstral skewness and high pass 0.01 Hz energy.

Figures S12 to S15 in the Supplementary Information show the time series of these covariates for the training and test events. Unlike the 0.1–20 Hz cepstral kurtosis and skewness, the high pass 0.01 Hz energy has more block-like features in its time series. This influences the sharp jumps from near-zero for training events 1 and 2 during their seismic swarms. In contrast, the change in high pass 0.01 Hz energy during the period of the seismic events on January 2 of training event 3 was more smooth, resulting in a smoother increase in forecast probabilities. The block-like features result from one

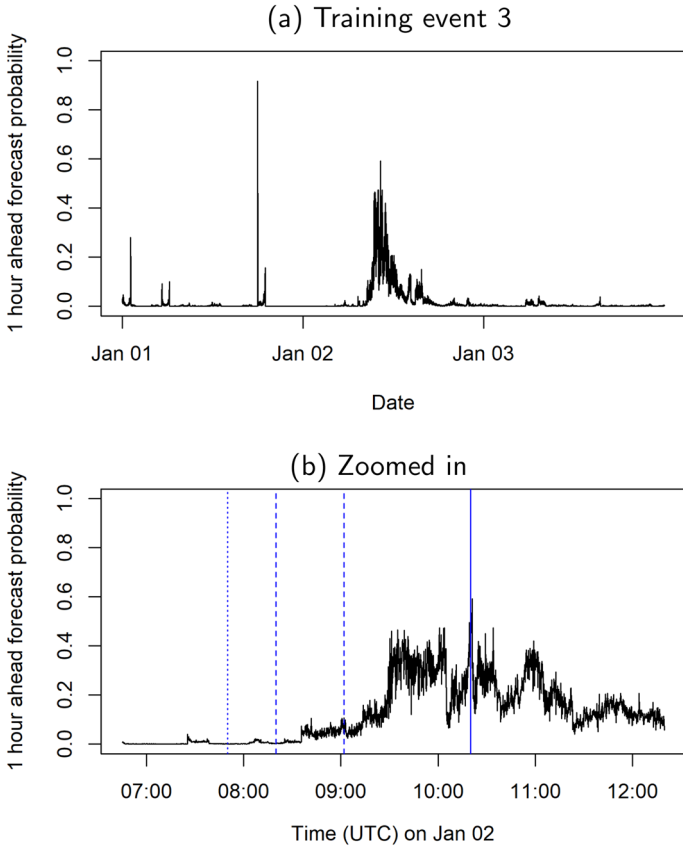


**Fig. 8** Training event 2: 1-h-ahead threshold exceedance forecasts using the lowest threshold estimated across events. The dotted, dashed and bold blue vertical lines denote the times of the seismic crises, start and end of the seismic swarms and the eruption onset respectively


extreme value in the high pass 0.01 Hz signal which causes high energy values for the length of the moving covariate window (1 h).

Since energy was defined as the sum of the squared signal, this covariate is also very sensitive to outliers. Future exploration can be done for making the covariates more robust to outliers. For example, instead of using the mean or average operation  $\frac{1}{n} \sum_{i=1}^n f_i$  on any sequence  $\{f_i : i = 1, \dots, n\}$ , the median operation which is less sensitive to outliers can be used. For covariates involving minimum or maximum values, the 0.1 or 0.9 quantiles can be considered.

Still focusing on high pass 0.01 Hz energy, Figs. S12 to S15 show that the covariate values tend to increase during the seismic swarms. Since the seismic swarms for training events 1, 2 and 3 occur slightly more than 1 h prior to their recorded eruptions, monitoring the energy values seems to provide good 1-h-ahead forecasts for the eruptions. It seems likely that the length of the covariate window and the forecast horizons can be optimised depending on the expected seismic crisis and swarm durations at



**Fig. 9** Training event 3: 1-h-ahead threshold exceedance forecasts using the lowest threshold estimated across events. The dotted, dashed and bold blue vertical lines denote the times of the seismic crises, start and end of the seismic swarms and the eruption onset respectively



**Fig. 10** Test event: 1-h-ahead threshold exceedance forecasts using the lowest threshold estimated across events. The dotted, dashed and bold blue vertical lines denote the times of the seismic crisis, seismic swarm and the eruption onset

a volcano. If the seismic swarms precede the eruption by a longer time period, as in the test event, the one-hour ahead forecast probabilities may not be so temporally accurate. This is observed in Fig. 10 for the 1-h-ahead forecast probabilities for the test event. Here, the seismic crisis lasts 5 h 30 min instead of 1–2.5 h for the training events. Referring to Fig. S15 of the Supplementary Information, the sole forecast probability outlier on October 13 and the spikes in forecast probabilities in the earlier part of October 14 are in line with spikes in the high pass 0.01 Hz energy covariate. However, while the largest energy value was recorded during the seismic swarm (see Fig. S15f), the forecast probability was higher nearer to the actual eruption onset. This shows the effect of the other covariates, including the 0.1–20 Hz cepstral kurtosis and skewness, which work together to moderate the forecast probabilities. By nature of the covariates involved, the forecast probabilities are sensitive to different aspects of seismicity.

In addition to the training and test events, the potential for the dynamic extreme value model for eruption forecasting is examined by looking at its performance for the three non-events. In line with expectations, the corresponding threshold exceedance forecasts remain very low, compared to the magnitudes during the training and test events (see Fig. S16 of the Supplementary Information). Similar training and test results were obtained for the 5–15 Hz frequency-filtered data. The corresponding plots are given in Sect. 8 of the Supplementary Information.

## 6 Discussion and Outlook

In Sect. 3, the dynamic extreme value model was fitted to training event 3, the seismic time series for the January 2010 eruption at Piton de la Fournaise. Promising results were obtained with spikes in the probabilistic forecasts about an hour prior to the eruption onset. This means that in this case, 1 h ahead warnings can be made with the chosen set-up: using the 1–5 Hz trace envelope as an eruption index with a 1-h covariate window and 1 h forecast horizon. Similar performance was also seen when the model was fitted to more data in Sect. 5.

In addition to the good forecast results, it was seen that EVT is useful for choosing the threshold from which to model exceedances. An appropriate threshold is important because it determines the balance between the number of exceedances used to inform covariate selection and the adherence of the threshold excesses to the asymptotic theory. In general, with a higher threshold, there are fewer exceedances to inform the model which could mean higher estimation uncertainty but the excess distribution becomes closer to a GPD.

By examining the training performance of the logistic regression for different threshold values, the analysis in Sect. 3 showed that using EVT to choose the threshold can improve forecast performance. When the threshold was increased towards its EVT-informed value, the AUC, a measure for how well the model can distinguish between exceedances and non-exceedances, generally increased for the 1–5 Hz, 0.1–20 Hz and high pass 0.01 Hz envelope indices. Intuitively, the threshold determines what it means to be extreme; hence it would affect the covariates selected and thus the forecast performance.

A related note to defining what it means to be extremes is that the chosen threshold and forecast probabilities are volcano-specific. In this paper, the dynamic extreme value model was fitted for one volcano, Piton de la Fournaise. This was based on its seismic data and hence the chosen thresholds and the forecast probabilities are specific to this volcano. If one were to fit similar models for other volcanoes separately, their forecast probabilities would not be directly comparable since they are specific to each volcano's natural characteristics such as eruption frequencies.

In practice, one would aim to train the forecast model with as much relevant data as possible. Section 5 presents some considerations to make when incorporating data from different events. For example, the p-values of GPD goodness-of-fit tests identify the threshold beyond which the excesses can be seen to follow a GPD. This procedure assumes that the same constant GPD applies to all the training events. However, as observed, different events can suggest different thresholds. Specifically, for training event 2, a higher threshold was inferred for the energy-related envelope index. This difference in envelope index values between events could be because some eruptions occur closer to the measurement station or because the eruption itself has a higher flux.

The proposed strategy to deal with the different threshold choices is to take the lowest identified threshold. The GPD regression component of the dynamic extreme value model would then model the non-stationarity of the GPD with appropriate covariates. In fact, modelling the excess distribution dynamically was seen to be more useful when the model was trained using multiple events and the lower threshold as compared to previously with just a single training event in Sect. 3. If instead a higher threshold was used, as suggested when all the training data were combined, only large events would be forecasted because training event 2 would dominate the model fit, resulting in the inability to forecast training events 1 and 3 well.

The proposed modelling framework is still far from operational use and can be extended in various ways. In the analysis, 1 h covariate and forecast windows were used. More work can be done to optimise these durations which are likely to be related to the seismic crisis and swarm durations. While the training events had seismic crises which lasted about 1–2 h, the test event had a much longer seismic crisis duration of 5 h 30 min. This could explain why the training forecasts were more temporally precise than the test forecasts. One could also experiment with computing the trace envelope in different ways. For example, different sliding windows to compute the Hilbert transform can be tested. Non-Hilbert transform methods of estimating amplitudes can also be used (Rosenblum et al. 2021). In addition, one might investigate the effect of different significance levels in the goodness-of-fit tests.

The analysis focused on seismic signals in the 1–5 Hz frequency range. As noted in the literature (Bormann et al. 2013) and observed from the AUC comparison in Fig. 5b, some frequency bands can be more useful for eruption forecasting than others. By combining information across useful frequency ranges through a joint, multivariate model, more comprehensive eruption forecast can be provided. Similarly, one can incorporate different monitoring signals such as gas emissions and ground deformation in the model. So far, only indices and covariates based on seismic signals have been used due to the prevalence of seismometers as volcano monitoring tools. Future extensions

can include different monitoring signals and account for their interdependence and shared covariates via joint models.

Data from one seismic station (UV05) were used in the analysis. The promising results indicate that such methods can be useful for volcanoes where there is only one measuring station. To check that the performance of the modelling approach is not dependent on the choice of UV05, which is the closest available station to the January 2010 eruptive centre, and a product of hindcasting, additional analyses were conducted for two other stations (UV08 and UV11). The results are provided in Section 9 of the Supplementary Information.

The training and test event results for the UV08 and UV11-fitted model indicate that a model fitted to another station can be used to forecast eruptions that occur closer to other stations and this can be useful for mitigating the issue of a damaged seismometer. Nevertheless, it is preferable to use the model fitted to station data which is closest to the eruptions in the training data and to use the model to forecast at the same station. This is because when the station is further away from the eruptive centre, the threshold exceedances of the envelope may be linked with other mechanisms which precede or occur after the eruption.

With more data, each eruption event can be matched to the nearest station and this station-specific training data can be used to build station-specific models. Future work can look into linking these station-specific models in a multivariate framework, taking into account their spatiotemporal relations. Since the distance from the eruption location is related to higher detected energies, an additional benefit of extending the framework to model signals from the monitoring network as a whole may be the ability to inform not just the timing, but also the location of the eruption through the varying forecast probabilities at different stations. In line with the extreme value approach, a multivariate generalised Pareto framework, similar to that by Rootzén and Tajvidi (2006), could be useful for extensions to multivariate and spatial models.

Given that the dynamic extreme value model has worked well for financial forecasting (Bee et al. 2019) and can be adapted for volcanic eruption forecasting, it is postulated to have high potential for wider applications. In particular, high sampling-rate data were used in both the financial and volcanic contexts. For the former, it was used to compute realised variations over relatively short horizons while in the latter, it helped to separate different frequency bands of interest. High sampling rate data relevant to other natural hazards and crises are also being made available; however, what is deemed as a high sampling rate is highly context-specific. For example, for sea levels, data were previously publicly available only at the monthly or annual scales. Hence, 1–15-min resolution is deemed a high sampling rate. These are increasingly sought after to study extreme sea levels and coastal flooding (Woodworth et al. 2016; Ozsoy et al. 2016; Zemunik et al. 2021). Using such data, the dynamic extreme value model can be adapted to forecast extreme sea levels and their impact on coastal communities. While high sampling rates are good to have, the framework itself does not depend on this, and in the absence of such data, one can still forecast, albeit on a coarser scale.

To adapt the dynamic extreme value model to wider settings, there are several general considerations, namely, what is a suitable index to compute threshold exceedances? What are reasonable forecast horizons and covariate windows? What covariates can be used to inform future behaviour? One might also consider using

algorithms for selecting the threshold automatically (see for example, Barder et al. 2018).

A general practical consideration that is shared across all contexts, be it finance, volcanoes or other hazards and crises, is the translation of the forecast probabilities into decisive action. What forecast probability warrants a warning or more drastic measures such as evacuation? The optimal strategy may not be straightforward but may involve many competing priorities and constraints, and should be determined on a case-by-case basis with multiple stakeholders.

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1007/s11004-023-10109-2>.

**Acknowledgements** The data used for the analysis were collected by the Institut de Physique du Globe de Paris, Observatoire Volcanologique du Piton de la Fournaise (IPGP/OVVPF) and the Laboratoire de Géophysique Interne et Tectonophysique (LGIT) within the framework of ANR\_08\_RISK\_011/UnderVolc project. The sensors are properties of the réseau sismologique mobile Français, Sismob (INSU-CNRS). This work comprises Earth Observatory of Singapore contribution no. 530. The research is supported by the NTU-Imperial collaboration grant INCF-2020-010 and the National Research Foundation, Prime Minister's Office, Singapore under the NRF-NRFF2018-06 award. AV would like to acknowledge funding by the Alan Turing Institute-Lloyd's Register Foundation Programme on Data-Centric Engineering.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Al-Mashhor AA, Al-Shuhail AA, Hanafy SM, Mousa WA (2019) First arrival picking of seismic data based on trace envelope. *IEEE Access* 7:128806–128815
- Barder B, Yan J, Xuebin Z (2018) Automated threshold selection for extreme value analysis via ordered goodness-of-fit tests with adjustment for false discovery rate. *Ann Appl Stat* 12(1):310–329
- Bee M, Dupuis DJ, Trapin L (2019) Realized peaks over threshold: a time-varying extreme value approach with high-frequency-based measures. *J Financ Econom* 17(2):254–283
- Beirlant J, Goegebeur Y, Segers J, Teugels JL (2004) *Statistics of extremes: theory and applications*, vol 558. Wiley, New York
- Bormann P, Wendt S, Klinge K (2013) Data analysis and seismogram interpretation. In: *New Manual of Seismological Observatory Practice 2 (NMSOP-2)*, Deutsches GeoForschungsZentrum GFZ, pp 1–151
- Brenguier F, Shapiro NM, Campillo M, Ferrazzini V, Duputel Z, Coutant O, Nercessian A (2008) Towards forecasting volcanic eruptions using seismic noise. *Nat Geosci* 1(2):126–130
- Carniel R, Guzmán SR (2020) Machine learning in volcanology: a review. In: Németh K (ed) *Updates in volcanology*, chapter 5. IntechOpen, Rijeka

- Chowdhury M, Turin T (2020) Variable selection strategies and its importance in clinical prediction modelling. *Fam Med Community Health* 8(1):66
- Danielsson J, De Vries CG (1997) Tail index and quantile estimation with very high frequency data. *J Empir Finance* 4(2–3):241–257
- Diebold FX, Schuermann T, Stroughair JD (1998) Pitfalls and opportunities in the use of extreme value theory in risk management. *Decision technologies for computational finance*. Springer, Berlin, pp 3–12
- Embrechts P, Klüppelberg C, Mikosch T (2013) *Modelling extremal events: for insurance and finance*, vol 33. Springer, Berlin
- Harrell FE (2015) *Multivariable modeling strategies*. Springer, Cham, pp 63–102
- Journeau C, Shapiro NM, Seydoux L, Soubestre J, Ferrazzini V, Peltier A (2020) Detection, classification, and location of seismovolcanic signals with multicomponent seismic data: example from the Piton de La Fournaise Volcano (La Réunion, France). *J Geophys Res Solid Earth* 125(8):1–19
- Ktonas PY, Papp N (1980) Instantaneous envelope and phase extraction from real signals: theory, implementation, and an application to EEG analysis. *Signal Process* 2(4):373–385
- Longin FM (2000) From value at risk to stress testing: the extreme value approach. *J Bank Finance* 24(7):1097–1130
- Malfante M, Dalla Mura M, Metaxian JP, Mars JI, Macedo O, Inza A (2018) Machine learning for volcano-seismic signals: challenges and perspectives. *IEEE Signal Process Mag* 35(2):20–30
- McNutt S, Roman D (2015) Chapter 59-volcanic seismicity. In: Sigurdsson H (ed) *The encyclopedia of volcanoes*, 2nd edn. Academic Press, Amsterdam, pp 1011–1034
- Moore B (1995) Hearing. Academic Press, San Diego, *Handbook of perception and cognition*
- Ozsoy O, Haigh ID, Wadey MP, Nicholls RJ, Wells NC (2016) High-frequency sea level variations and implications for coastal flooding: a case study of the Solent, UK. *Contin Shelf Res* 122:1–13
- Ren CX, Peltier A, Ferrazzini V, Rouet-Leduc B, Johnson PA, Brenguier F (2020) Machine learning reveals the seismic signature of eruptive behavior at piton de la fournaise volcano. *Geophys Res Lett* 47(3):e2019GL085523
- Ribatet M (2016) Crash course on univariate extreme value theory. Lecture notes from the June 2016 Extreme Value Modeling and Water Resources Summer School held in Lyon, France
- Rootzén H, Tajvidi N (2006) Multivariate generalized Pareto distributions. *Bernoulli* 12(5):917–930
- Rosenblum M, Pikovsky A, Kühn A, Busch J (2021) Real-time estimation of phase and amplitude with application to neural data. *Sci Rep* 11(1):18037
- Roult G, Peltier A, Taisne B, Staudacher T, Ferrazzini V, Di Muro A (2012) A new comprehensive classification of the Piton de la Fournaise activity spanning the 1985–2010 period. Search and analysis of short-term precursors from a broad-band seismological station. *J Volcanol Geotherm Res* 241–242:78–104
- Salvage RO, Karl S, Neuberg JW (2019) *Volcano seismology: detecting unrest in wiggly lines*. Springer, Cham, pp 185–201
- Seydoux L, Shapiro N, de Rosny J, Brenguier F, Landès M (2016) Detecting seismic activity with a covariance matrix analysis of data recorded on seismic arrays. *Geophys J Int* 204(3):1430–1442
- Steyerberg E (2009) *Selection of main effects*. Springer, New York, pp 191–211
- Sugihara G, May RM (1990) Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series. *Nature* 344(6268):734–741
- Taisne B, Brenguier F, Shapiro NM, Ferrazzini V (2011) Imaging the dynamics of magma propagation using radiated seismic intensity. *Geophys Res Lett* 38(4):2–6
- Taner MT (2001) Seismic attributes. *CSEG Rec* 26(7):48–56
- Trnkoczy A (1999) Understanding and parameter setting of sta/Ita trigger algorithm. [https://gfzpublic.gfz-potsdam.de/rest/items/item\\_4097/component/file\\_4098/content](https://gfzpublic.gfz-potsdam.de/rest/items/item_4097/component/file_4098/content). Accessed 27 Apr 2021
- Venables WN, Ripley BD (2002) *Modern applied statistics with S*, 4th edn. Springer, New York
- Whitehead MG, Bebbington MS (2021) Method selection in short-term eruption forecasting. *J Volcanol Geotherm Res* 66:107386
- Wild AJ, Bebbington MS, Lindsay JM, Charlton DH (2021) Modelling spatial population exposure and evacuation clearance time for the Auckland Volcanic Field. *New Zealand. J Volcanol Geotherm Res* 416:107282
- Withers M, Aster R, Young C, Beiriger J, Harris M, Moore S, Trujillo J (1998) A comparison of select trigger algorithms for automated global seismic phase and event detection. *Bull Seismol Soc Am* 88(1):95–106



- Woodworth PL, Hunter JR, Marcos M, Caldwell P, Menéndez M, Haigh I (2016) Towards a global higher-frequency sea level dataset. *Geosci Data J* 3(2):50–59
- Zemunik P, Šepić J, Pellikka H, Čatipović L, Vilibić I (2021) Minute Sea-Level Analysis (MISELA): a high-frequency sea-level analysis global dataset. *Earth Syst Sci Data* 13(8):4121–4132