

A Deep Catalogue of Marsquakes

Nikolaj L. Dahmen^{1*}, John F. Clinton², Men-Andrin Meier¹, Simon C. Stähler¹, Savas Ceylan¹, Doyeon Kim¹, Alexander E. Stott³, and Domenico Giardini¹

¹Institute of Geophysics, ETH Zurich, Zurich, Switzerland

²Swiss Seismological Service, ETH Zurich, Zurich, Switzerland

³Institut Supérieur de l'Aéronautique et de l'Espace SUPAERO, Toulouse, France

Key Points:

- Marsquakes recorded by InSight's seismometer are challenging to detect and analyse due to typically low signal-to-noise-ratio
- We present MarsQuakeNet - a convolutional neural network for marsquake detection and denoising - trained on synthetic data
- Our catalogue is consistent with the existing manual catalogue, extends it by 60% and confirms significant changes in event rate across Martian years

*Institut für Geophysik, NO E 17, Sonneggstrasse 5, 8092 Zürich, Switzerland

Corresponding author: Nikolaj Dahmen, nikolaj.dahmen@erdw.ethz.ch

Abstract

NASA’s InSight seismometer has been recording Martian seismicity since early 2019, and to date, over 1300 marsquakes have been catalogued by the Marsquake Service (MQS). Due to typically low signal-to-noise ratios (SNR) of marsquakes, their detection and analysis remain challenging: while event amplitudes are relatively low, the background noise has large diurnal and seasonal variations and contains various signals originating from the interactions of the local atmosphere with the lander and seismometer system. Since noise can resemble marsquakes in a number of ways, the use of conventional detection methods for catalogue curation is limited. Instead, MQS finds events through manual data inspection.

Here, we present MarsQuakeNet (MQNet), a deep convolutional neural network for the detection of marsquakes and the removal of noise contamination. Based on three-component seismic data, MQNet predicts segmentation masks that identify and separate event and noise energy in time-frequency domain. As the number of catalogued MQS events is small, we combine synthetic event waveforms with recorded noise to generate a training data set.

We apply MQNet to the entire continuous 20 samples-per-second waveform data set available to date (>1000 Martian days), for automatic event detection and for retrieving denoised amplitudes. The algorithm reproduces all high quality-, as well as majority of low quality events in the manual, carefully curated MQS catalogue. Furthermore, MQNet detects ~60 % additional events that were previously unknown with mostly low SNR, that are verified in manual review. Our analysis on the event rate confirms seasonal trends and shows a substantial increase in the second Martian year.

Plain Language Summary

InSight’s seismometer on Mars has recorded over 1300 marsquakes since its full deployment in early 2019. Marsquakes are often weak compared to the seismic background noise, which makes their detection and analysis challenging. For this reason, the current event catalogue relies on identifying events in manual data review, which can result in an inconsistent event catalogue with weak events being missed and quality standards changing over time.

In this study, we use a type of artificial neural network for the automatic detection of marsquakes and to separate seismic signal and background noise. Since these artificial neural networks usually require many examples to learn from but the number of known marsquakes is low, we generate synthetic marsquake examples to train the network. We run the detection algorithm across the mission and compare its performance to the manually compiled event catalogue: the algorithm can also detect the majority of identified marsquakes and additionally finds many weaker, previously missing events, thereby extending the number of known marsquakes by ~60 %, from 1297 to 2079. Further, our results show substantial variations in event numbers throughout the mission.

1 Introduction

Seismology allows us to explore the interior of the Earth by studying the numerous earthquakes that occur each year and are recorded by a worldwide network of thousands of seismometers. In contrast, extraterrestrial seismology has been relying on the limited data sets collected during the Apollo (Latham et al., 1970), Viking (Anderson et al., 1977) or Venera (Ksanfomaliti et al., 1982) missions to the Moon, Mars and Venus. While numerous moonquakes were found in the Apollo data set, no quake was unambiguously identified on Mars or Venus due to the high seismic noise level or the short mission duration. On Mars, this has changed with the arrival of NASA’s InSight (National Aeronautics and Space Administration; Interior Exploration using Seismic Investigations, Geodesy and Heat Transport) mission in 2018 (Banerdt et al., 2020). With the lessons learned during the Viking mission, the InSight lander installed a seismometer package directly on the ground and covered it with a wind and thermal shield to reduce atmospheric noise (Lognonné et al., 2019).

Since the complete deployment in beginning of 2019, more than 1300 marsquake have been detected (InSight Marsquake Service, 2022), giving us a much better understanding of the Martian seismicity (Giardini et al., 2020) and providing constraints on the interior of the planet (Knapmeyer-Endrun et al., 2021; Khan et al., 2021; Stähler et al., 2021; Durán et al., 2022; Khan et al., 2022; Kim, Lekić, et al., 2021; Li et al., 2022).

The Marsquake Service (MQS) carefully monitors the seismic data and identifies, locates and catalogues marsquakes (Clinton et al., 2018, 2021). InSight’s noise levels are low, and at times much lower than on Earth at periods of 1-20 s, due to the absence of ocean microseismic noise (Lognonné et al., 2020, Fig. 2). This allows for the detection of a large number of small marsquakes with moment magnitudes between 1.5 and 2.5 from approximately 1500 km distance during the quiet evenings of the Martian day. During other times of the day, local winds, pressure fluctuations and extreme temperature changes elevate the seismic noise level by 30 dB and more and consequently contaminate and almost entirely obscure marsquake waveforms.

The MQS identifies seismic events by manual inspection of the data, as conventional event detectors such as the STA/LTA (Allen, 1982, e.g.) perform poorly on the single-station records due to high noise fluctuations. Onboard wind and pressure sensors aid in differentiating seismic events from atmospheric signals and can further quantify the atmospheric noise (Charalambous et al., 2021). As wind and pressure recordings are not available throughout the mission, the atmospheric noise can also be directly estimated from the seismic imprint of wind-excited eigenmodes of the InSight lander (Dahmen, Zenhäusern, et al., 2021). Lander mode excitation is an important event discriminator and its absence is often used to identify marsquakes in the manual MQS review, as well to find another class of short-duration, high-frequency signals that are interpreted as local, thermal events (Dahmen, Clinton, et al., 2021). Since these thermal events fall into families with near-identical waveforms, template matching techniques are effective in completing their catalogue (Dahmen, Clinton, et al., 2021; Compaire et al., 2022). With a similar method and using 9 distant marsquakes as templates, Sun and Tkalčić (2022) proposed the detection of 47 new events that are not currently included in the MQS catalogue.

Recently, various studies (Perol et al., 2018; Ross et al., 2018; Mousavi et al., 2019; Zhu & Beroza, 2019; Mousavi et al., 2020) have demonstrated the power of deep learning methods, and in particular, supervised convolutional neural networks (CNN), for seismic event detection and phase picking. While these studies could resort to massive amounts of labelled earthquake and noise data for network training, the volume of extraterrestrial recordings is much smaller. Possible approaches to tackle this issue were shown by Civilini et al. (2021), who detected high quality moonquakes with a CNN trained on earthquakes, or by Knapmeyer-Endrun and Hammer (2015), who used for the same task a Hidden Markov Model-based detector relying on only few template waveforms. Unsupervised methods that operate on unlabelled data and are used for clustering or dimensionality reduction have been applied to identify dominant lander or atmospheric signals in the InSight data set (Barkaoui et al., 2021).

In this study, we use a CNN, named MarsQuakeNet (MQNet), for the detection of marsquake energy in the time-frequency domain, and highlight its value for waveform denoising. MQNet is based on the UNet architecture developed by Ronneberger et al. (2015) for medical image segmentation. The UNet consist of a set of convolutional layers that are arranged in an encoder-decoder setup. In medical image analysis (Ronneberger et al., 2015), this technique is used to build segmentation masks that pixel-wise classify cells or organs in the human body. Jansson et al. (2017); Chandna et al. (2017) and others showed how this method can decompose the time-frequency representation of audio tracks into different instruments or vocals. Similar to the mix of different instruments in an audio track, seismic data can be understood as a superposition of signals from various sources, which can potentially be separated by this method. Zhu et al. (2019) introduced this approach to seismic data analysis and demonstrated its value for earthquake signal denoising. Later, Tibi et al.

(2021) and Yang et al. (2022) used this method to denoise regional earthquake records, and earthquake records made in urban settings, respectively.

As we do not have a large labelled dataset readily available to train the network - the MQS catalogue contains only about 1300, mostly low signal-to-noise ratio (SNR) events - we train MQNet on synthetic data. We generate event waveforms with characteristics following the different marsquake types and combine them with recorded noise, which allows the network to learn the idiosyncrasies of the InSight data. The efficacy of this approach is demonstrated on existing catalogued MQS events and then used for event detection on the complete 20 samples-per-second (sps) InSight dataset, spanning 1043 sols. We compare the performance to the MQS event catalogue and review and report our new detections. The product of this study is the first automatically-generated marsquake catalogue, which extends the manually curated MQS catalogue from 1297 to 2078 events in the investigated time period (+60 %). Our extended catalogue is consistently produced across the mission and allows us to study the temporal distribution of marsquakes.

2 InSight Dataset

2.1 Instruments

The InSight lander carried the Seismic Experiment for Interior Structure (SEIS), as part of its payload, which consists of three-component Very Broadband (VBB) and a three-component Short-Period (SP) instruments (Lognonné et al., 2019). SEIS was placed on the Martian surface at a 1.8 m distance from the lander and subsequently covered by the wind and thermal shield to protect it from the harsh atmospheric conditions. SEIS, and in particular the VBB sensor, were fully-operational since about Sol 72.

The available VBB data are mainly sampled at 10 sps until Sol 181 and at 20 sps from thereon; here, a sol corresponds to a Martian day (~ 24 h 40 min) counted from landing. Higher sampling rates are only transmitted on request for short periods due to bandwidth limitations. The recordings are nearly continuous and only occasionally contain gaps lasting minutes to hours caused by competing science activities, limited power or transmission failures. Two longer gaps were caused by the loss of communication to the lander during the solar conjunction between Sol 267-288, and power constraints during Sol 1107-1129. Data could be stored on the lander during the solar conjunction in the second year, and were transmitted afterwards. We focus on the continuous 20 sps data set from Sol 182 (June 2020) to 1224 (May 2022) (InSight Marsquake Service, 2022), after which recordings are only sporadically available due to increasingly severe power constraints.

Next to the seismic data streams, atmospheric conditions are recorded by a pressure sensor, and two sets of temperature and wind sensors, all mounted on the lander deck. The atmospheric sensors had to be turned off during large portions of the later mission (starting from around Sol 700 for wind, and Sol 850 for pressure) due to power restrictions.

2.2 Seismic Data

Throughout the manuscript, we show the VBB seismic data in velocity, following the instrument deconvolution (and a frequency taper) and the rotation to vertical (Z), north (N) and east (E) components using the ObsPy package (See Open Research). The seismic data contain various noise signals that might resemble seismic events in duration, frequency content and amplitude and are potentially challenging to differentiate from seismic events with automated detection algorithms. We explain them by taking the example of Sol 923, shown in Figure 1.

The VBB instrument records extremely small ground motions (Lognonné et al., 2020) and is consequently also sensitive to the local environment. The main source of background noise in the bandwidths of interest for seismic events (0-10 Hz) is the wind that predomi-

nantly couples to the ground through lander induced vibrations (Murdoch et al., 2018; Stott et al., 2021), as modelled pre-mission (Mimoun et al., 2017) and observed on mission data (Charalambous et al., 2021). This, wind driven background noise varies throughout the day by several orders of magnitudes. Sol 923 shows a division of the day into periods with moderate noise levels during the night, high noise levels from sunrise to just before sunset, and the lowest noise conditions in the evening. The high noise level is relatively constant from sunrise to noon, whereas in the afternoon, fluctuations are much stronger. We observe the broadband signature of individual minute-long wind gusts, with short quiet periods in between. The marsquakes found on this sol are representative of other events, having amplitudes that are only marginally above low noise periods. As a consequence, seismic events are primarily seen in the quiet evening period.

An outstanding feature of the dataset is the excitation of resonances which correspond to i) wind-excited eigenmodes of the lander (Dahmen, Zenhäusern, et al., 2021; Murdoch et al., 2018) or seismometer system (Hurst et al., 2021), some of which are associated with different positions of InSight’s robotic arm and change during the mission (Dahmen, Zenhäusern, et al., 2021; Stott et al., 2021); ii) a measurement artefact that leads to thin spectral peaks at 1 Hz with overtones (Zweifel et al., 2021), and iii) a broad resonance at 2.4 Hz that has been associated with a subsurface structure (Hobiger et al., 2021; Dahmen, Zenhäusern, et al., 2021; Compaire et al., 2021). In contrast to the lander modes, the 2.4 Hz resonance is only seen in quiet periods and is excited by seismic events.

Additionally, the recordings are corrupted by transient high-amplitude signals - glitches - which are caused by SEIS-internal relaxation of thermal stresses and instrument tilts (Scholz et al., 2020; Kim, Davis, et al., 2021; Barkaoui et al., 2021). Glitches usually have duration of 10-30 s and are found below 1 Hz, but sometimes also extend to the Nyquist frequency (10 Hz) and beyond. Other second-long signals resulting from stress releases in the lander and seismometer system occur at high frequencies (>10 Hz) and occasionally leak below 10 Hz (Ceylan et al., 2021; Compaire et al., 2021).

A number of other common noise signals are produced by pressure drops and long-period pressure fluctuations, the motion of InSight’s robotic arm, various forms of sensor cross-talks, sensor calibrations and others, all summarised in Ceylan et al. (2021). Among these are several signal types that can resemble seismic events or individual phase arrivals.

2.3 Seasonal Changes of Background Noise

The atmospheric conditions apparent in the fluctuations of the seismic noise level shown for Sol 923 (Fig. 1) are representative for the Martian spring and summer in Elysium Planitia. During fall and winter, winds are strong throughout the sol and noise levels are persistently high except for occasional calmer periods (see Supporting Information, Fig. S1). This follows the pattern of turbulence outlined in Chatain et al. (2021), which shows convective turbulence to dominate during the day time throughout the mission and shear driven turbulence becoming almost as strong during the night in the fall and winter. As pointed out in Charalambous et al. (2021) these different turbulence mechanisms exhibit a different sensitivity in the seismic data. One Martian year (668 sols) into the mission, we find that noise patterns starts to closely repeat the previous seasons. A noticeable difference in the second Mars year is the increased occurrence of high-amplitude glitches that corrupt the recordings, likely caused by the reduction of the SEIS-internal temperatures after switching off the heater.

2.4 Marsquake Event Types and Qualities

MQS catalogues seismic events by their frequency content and distinguishes between low- and high frequency event families (Clinton et al., 2021), and five event types, examples of which are displayed in Figure 2. Events with high frequency energy are interpreted as

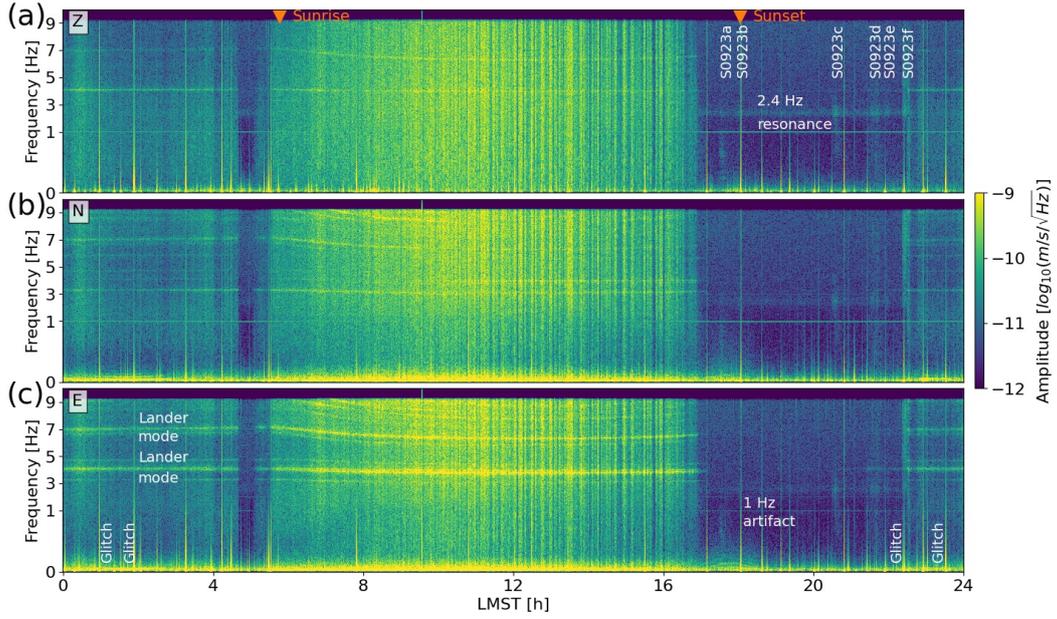


Figure 1: Seismic data from Sol 923. Day-long spectrograms showing (a) vertical, (b) north and (c) east components. Key features are indicated - six seismic events from MQS catalogue (between 17:30-22:30 LMST), typical features of the noise, and sunrise and sunset times. MQS: Marsquake Service, LMST: local mean solar time.

shallow events propagating in the crust (van Driel et al., 2021), while events with low-frequency energy are thought to be deeper events that generate body wave energy travelling through the interior of the planet (Giardini et al., 2020). Event signals of all families and types typically last several minutes to tens of minutes.

High frequency event family (Clinton et al., 2021; van Driel et al., 2021; InSight Marsquake Service, 2022):

- (a) Very high frequency (VF) events usually exhibit a broad excitation of the 2.4 Hz resonance and above, often reaching the 10 Hz Nyquist frequency and beyond (sometimes to 30 Hz); characteristic is the excess of horizontal energy at high frequencies; located at distances of $\sim 1\text{-}45^\circ$ ($\sim 50\text{-}2600$ km).
- (b) High frequency (HF) events show a broad excitation of the 2.4 Hz resonance, and are typically above the noise level from 1 Hz to at least 4 Hz; occasionally also with energy at lower frequencies; located in $\sim 15\text{-}32^\circ$ distance ($\sim 940\text{-}1870$ km).
- (c) 2.4 Hz (2.4) events are interpreted as low amplitude HF events and only show a narrow excitation of the 2.4 Hz resonance; located in $\sim 13\text{-}41^\circ$ ($\sim 780\text{-}2430$ km).

Low frequency event family (Clinton et al., 2021; Giardini et al., 2020; InSight Marsquake Service, 2022):

- (d) Broadband (BB) events have most of their energy below 1 Hz but additionally include energy at the 2.4 Hz resonance and above; 2.4 Hz excitation is sometimes only observed around the event start; located in distances of 27° (1600 km) and beyond.
- (e) Low frequency (LF) events have most energy below 1 Hz; located in distances of 27° and beyond.

All examples in Fig. 2 are from high quality events, while the majority of the events in the MQS catalogue have a much lower SNR. MQS assigns location qualities to the events, ranging from QA - clear phases and polarisation (full location), QB - clear P- and S-phases or clear polarisation (distance or back azimuth), QC - clear signal but phase picks are uncertain or unclear, to QD - low SNR signals or signals possibly of non-seismic origin (Clinton et al., 2021). The highest quality for events from the high frequency family is QB (MQS catalogue V10) because scattering at high frequencies prevents meaningful estimation of back azimuths.

About 93% of the catalogued events belong to the high frequency event family, most of which are low-amplitude 2.4 Hz events. 14% of the events are of high quality (QA/B).

MQS further catalogues a separate class of short-duration signals which are interpreted as local thermal cracking. We do not consider this signal type, as they are efficiently detected with simple template matching techniques (Dahmen, Clinton, et al., 2021).

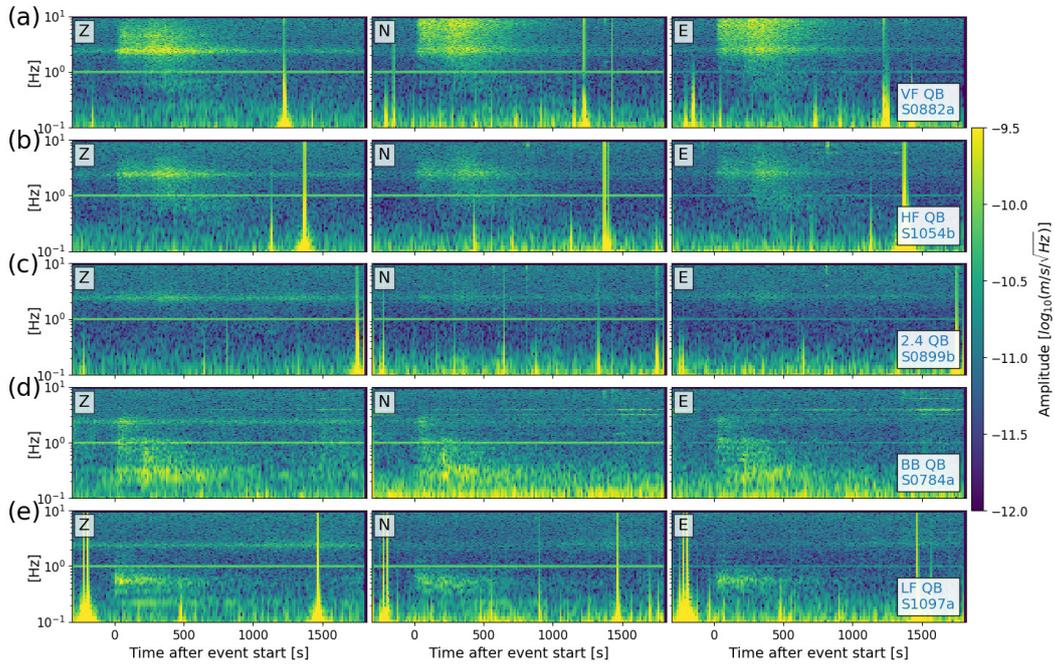


Figure 2: Marsquake types. Vertical, north and east spectrograms of (a) very high frequency (S0882a), (b) high frequency (S1054b), (c) 2.4 Hz (S0899b), (d) broadband (S0764a), and (e) low frequency events (S1097a); all catalogued as Quality B by the Marsquake Service. All windows are 35 minutes long and include 5 min of pre-event noise.

3 Method

3.1 UNet

Deep learning methods, and in particular convolution neural networks (CNN), are nowadays widely used for various complex tasks such as natural language processing or visual object recognition (LeCun et al., 2015). Unlike conventional machine learning techniques that require a certain level of domain knowledge and careful feature engineering, deep learning operates on raw data and automatically finds the optimal data representation for each task. CNNs are composed of simple non-linear operations that transform the data into more abstract representations, and by combining many operations, can learn complex functions (LeCun et al., 2015).

Fully-convolutional networks typically consist of a sequence of layers along a contracting path where successively, convolution, non-linear activation and down-sampling operations are performed. Advantageous attributes of CNNs are their sparse connectivity, shared parameters and the equivariance to translations (Goodfellow et al., 2016): convolutional layers contain a set of weights, called kernels or filters, which are element-wise multiplied (dot product) with only small subsets of the input data - the receptive fields ('sparse connectivity'). This operation is systematically repeated over the complete input data to produce the feature maps. Thereby, the same kernel or set of kernels is used across the input ('shared parameters'), which allows the CNN to find features independent of their location ('equivariance to translation') and simultaneously acts as a regularizer. This structure makes CNNs substantially more efficient than fully-connected neural networks and less prone to overfitting (Goodfellow et al., 2016). In addition, activations impart non-linear expression functionality to the model, and down-sampling operations gradually reduce the dimensions of the feature maps and allow the model to find meaningful features across different scales.

The UNet (Ronneberger et al., 2015) builds on the fully-convolutional architecture, but is extended by an expansive path that leads to high-resolution localisation of features. The network typically consumes 2D data in form of images or time-frequency representations of time series. The output is not a single label but a segmentation mask that pixel-wise classifies 2D data into regions with common class labels.

In this study, we use the UNet architecture to detect marsquakes by casting the problem as a segmentation task. When working with the InSight seismic data, we are primarily dealing with low SNR events that are hidden in highly fluctuating noise conditions. As will be shown later, the predicted masks are more informative than a single event/noise label and allow us to detect, analyse and confirm weak events. Besides the event detection, we can also use the masks to remove the noise and estimate the frequency content, duration and amplitude of the marsquake signals.

Following Zhu et al. (2019), we define the recorded seismic signal Y as the background noise N that is superimposed with a marsquake signal S . The time-frequency representation of the recorded signal $Y(t, f)$ can be decomposed by finding the masks $M_{event}(t, f)$ and $M_{noise}(t, f)$ that separate the recorded signal into $S(t, f)$ and $N(t, f)$. The masks for event and noise have the same dimension as the recorded signal with values between 0 and 1 and are defined as:

$$M_{event}(t, f) = \frac{|S(t, f)|}{|S(t, f)| + |N(t, f)|} \quad (1)$$

$$M_{noise}(t, f) = \mathbb{1} - M_{event} \quad (2)$$

To separate the recorded signal into event and noise, the time-frequency representation is element-wise multiplied with event and noise masks, respectively. In typical binary image segmentation tasks, objects are not overlapping and segmentation masks have a rather

binary value distributions. In our case, seismic events and noise overlap in time-frequency domain. Consequently, the mask values indicate the contribution of the event and noise to the total spectral amplitude observed.

We train the UNet to take $Y(t, f)$ as input to predict the segmentation masks $M_{event}(t, f)$ and $M_{noise}(t, f)$. We can then use the masks for *i*) signal detection: we check the event mask if it contains a substantial amount of non-zero elements; *ii*) signal denoising: we can multiply the masks with $Y(t, f)$ to separate marsquake and noise signals.

The processing pipeline for mask prediction and application is schematically shown in Figure 3. Starting from the raw waveforms, we deconvolve the instrument response and rotate to ZNE. We fix the length of each sample’s waveform to approximately 27 min (1628 s) - a duration that is longer than the average event signal. We compute the short-time Fourier transform (STFT, window length of 12.8 s, 50% overlap) of the waveforms and collect the real and imaginary parts of the complex-valued time-frequency coefficients. This leads to dimensions of 256x256x6 for each sample (1628 s x 0-10 Hz x real and imaginary for ZNE components).

The complex-valued coefficients have a large dynamic range (the interquartile range of absolute values of a full sol is typically between $1e-12$ to $1e-10$ $m/s/\sqrt{Hz}$) and include many outliers associated with glitches (partly $>1e-6m/s/\sqrt{Hz}$). We sample-wise standardise the data by removing the median and scaling to the interquartile range. Extreme values that typically correspond to glitches or other signals of non-seismic origin are clipped.

The UNet ingests these standardised STFT coefficients and predicts the event and noise masks for each component. The event masks detect, locate and scale marsquake energy in time-frequency domain. By element-wise multiplication of the STFT coefficients with the predicted masks, we can separate the event and noise. Finally, we can transform the extracted event and noise back to time domain by computing the inverse STFT.

3.2 UNet Implementation

We implement our UNet architecture with the open-source machine learning library TensorFlow (Abadi et al., 2015). As shown in Figure 3 (‘Mask Prediction’), the contracting path consist of five blocks of convolutional layers with kernel size 3x3 and ReLU activations (two of each), followed by 2x2 Maxpooling for down-sampling. At the end of the contracting path, the feature maps are condensed to a size of 8x8x512. Five blocks of upsampling layers, concatenations, convolutional layers and ReLU activation constitute the expansive path. The upsampling is performed by transposed convolution operations with a kernel size of 3x3 and a stride of 2, which effectively doubles the height and width of the feature map, and halves the channel number. The concatenations combine feature maps from the contracting and expansive path for the precise localisation of features. The final convolutional layer has a kernel size of 1x1 and is followed by a sigmoid activation that produces mask values in the range of 0 to 1.

3.3 Training Dataset

Deep learning models usually perform best when trained on large datasets. For example, Zhu et al. (2019) built a data set by combining more than 56k earthquake and 179k noise waveforms. When labelled training samples are sparse, data sets are commonly augmented by generating copies with slightly altered properties (Ronneberger et al., 2015, e.g.). Similarly, the MQS catalogue currently contains only about 1300 events that are very unevenly distributed among the different event types, and most are highly noise-contaminated. In order to train a robust model, we build a data set based on synthetic event waveforms replicating features of real marsquake waveforms that are combined with noise recorded by the VBB sensor.

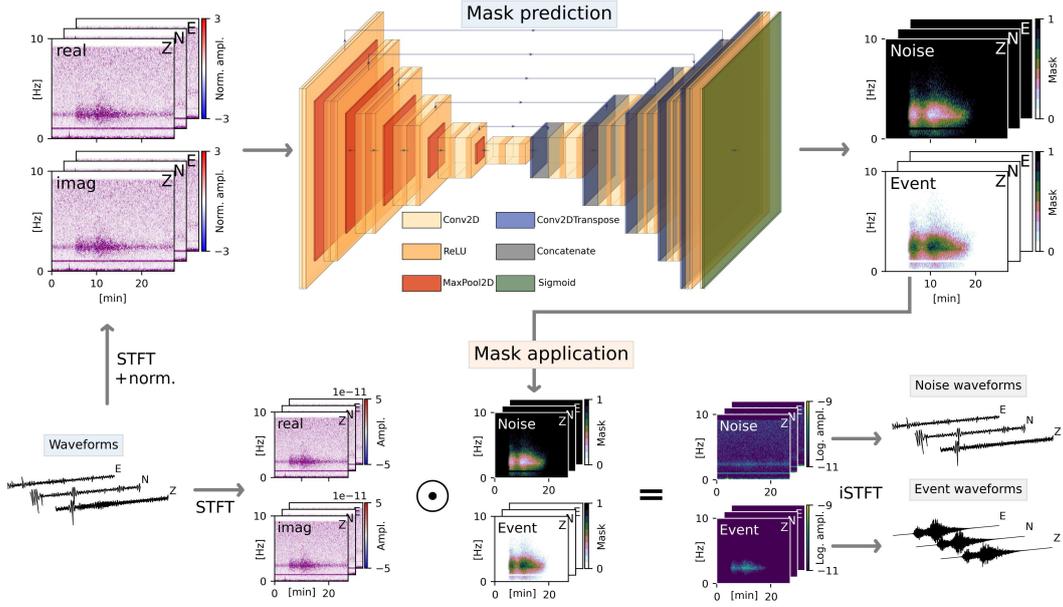


Figure 3: Workflow for event detection and denoising. The normalised time-frequency coefficients of the recorded waveforms serve as the input for the MQNet to predict the event and noise masks. The predicted masks are then element-wise multiplied (denoted by \odot) with the time-frequency coefficients to separate the event and noise. Note that the shown network architecture indicates varying dimensions of the feature maps, but is not to scale. Conv2D(Transpose): 2D (transposed) convolutional layer; ReLU: rectified linear unit (activation); Sigmoid: sigmoid activation; Concatenate: concatenation of encoder and decoder feature maps; MaxPool2D: 2D max. pooling operation; iSTFT: inverse short time Fourier transformation.

We create synthetic event waveforms using a modified version of the stochastic waveform simulation presented by Boore (2003). The method is particularly suited for the simulation of high-frequency ground motions and allows us to create event signals that mimic the time- and frequency-domain characteristics of real recorded signals for all five marsquake types.

Figure 4 shows the procedure for a 2.4 Hz event, the most common event type in the MQS catalogue. The inputs are the normalised time-domain envelope from a real, recorded event signal, an amplitude spectrum resembling the required event type, and a white noise time series. At first, the white noise time series is scaled in time-domain by the envelope. The modulated time series is then transformed to time-frequency domain where we scale the normalised real and imaginary parts by the spectrum. By applying the inverse STFT we obtain the event waveform. Different from the method proposed by Boore (2003), we work with the time-frequency representation instead of the amplitude and phase spectra since this is simpler for further processing. To increase the variability of the event waveforms, we randomise the time-domain envelopes and spectra used as input.

Following this procedure, we generate waveforms that match the different marsquake types (compare Fig. 2):

1. High frequency event family: to simulate VF events, we take the smoothed envelopes and spectra from three high quality VF events located in representative distances - S0128a (8°), S0334a (20°) and S0264e (35°). For 2.4 Hz and HF events, we use the envelopes from six high quality HF events - S0423c, S0371b, S0340a, S0323a, S0314b,

S0228c (21-27°) - and use a Lorentz peak centred at 2.4 Hz with variable width as the spectrum. We then randomly vary the envelope in duration by stretching or compressing, and in relative amplitude by adding an exponent in the range of 0.5-2 to each amplitude value. The amplitudes on the three components are scaled to match the relative amplitude ratio from the template events.

2. Low frequency event family: we compute smoothed envelopes from Instaseis (See Open Research) by generating events with random source parameters at different distances and back azimuths. The event spectra are approximated by narrow- to wide peaks at low frequencies, and partly include a Lorentz peak at 2.4 Hz to represent BB events. The three component amplitudes are controlled by the relative amplitude of the Instaseis envelopes.

We build a synthetic event waveform catalogue consisting of 40% LF/BB (MQS catalogue: 7%), 30% VF (5%) and 30% HF/2.4 Hz events (88%). The different distribution compared to the MQS catalogue, accounts for the much higher variability of LF, BB, and VF events in terms of frequency content, duration and S-P time.

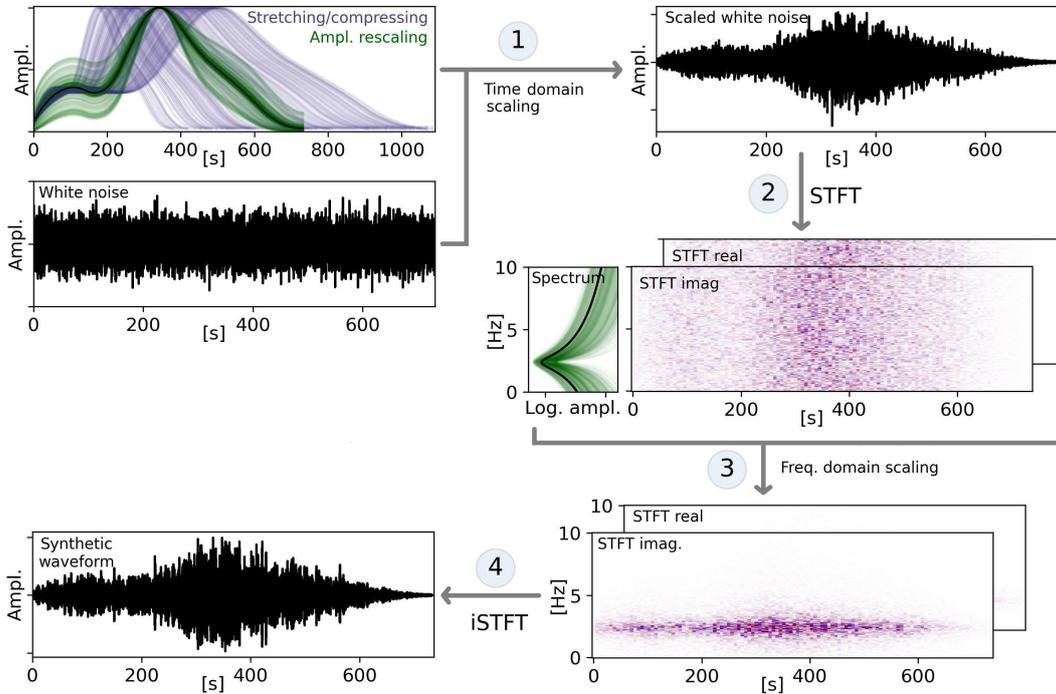


Figure 4: Generating synthetic event waveforms. The inputs are an envelope and a spectrum matching the required event type, and a white noise time series. We select a randomised envelope (stretched/compressed, re-scaling of relative P and S amplitude, separately shown) and multiply it with the white noise time series to obtain the scaled white noise time series (1). After transformation to time-frequency domain (2), the real and imaginary parts of the STFT are multiplied with the spectrum to scale the frequency content (3). Back-transformation to time domain results in the synthetic event waveform (4). iSTFT: inverse short time Fourier transformation.

Since InSight’s background noise is complex and highly fluctuating, we use real VBB noise to capture all types of environmental and instrumental signals in the training dataset. We collect 30 min periods from a full Martian year (Sol 290-960), as some features have a

seasonal character (Ceylan et al., 2021), and exclude periods that contain catalogued MQS events.

Each noise sample is randomly cut to a 27 min window and superimposed with a randomly shifted event sample in time-frequency domain. Thereby, the event amplitude is rescaled to produce samples with a SNR in the range of 0.67-5. In this context, we define the SNR as the ratio of the event and noise standard deviation and compute it only over time-frequency bins that contain the event energy. In the last step, we compute the event and noise masks following equations 1 and 2.

Following this procedure, we generate two separated data sets with noise from even- and odd-numbered sols, each consisting of about 50k samples. We train a model on each set and can thus ensure that data used for network training and subsequent event detection do not overlap.

3.4 Training and Validation

We perform the training separately for two models, named *even-sol* and *odd-sol*, which are trained on two different data sets containing noise recordings from even and odd-numbered sols. Each of these data sets is further split into approx. 40k training and 10k validation samples. The training is done in mini-batches of size 32 using the Adam optimiser (Kingma & Ba, 2014) with a learning rate of 0.001, while minimising the binary cross entropy loss that is typically optimised in segmentation task (Ronneberger et al., 2015). We test several modifications of the model on a smaller subset and, in particular, explored the effect of different down-sampling layers, network sizes, optimiser and learning rates, as well as the addition of batch normalisation layers (Supporting Information Fig. S2). The previously described architecture (see 3.2 UNet Implementation) led to the best results, though some modification (e.g. average pooling instead of max pooling) produced to comparable results.

The training on the full data set takes approx. 73 mins per epoch on a GPU (Tesla T4). We terminate the process after the validation loss stopped decreasing for 3 epochs and select the model with the best performance (Supporting Information, Fig. S3). The final models have a test loss of 0.087 for *even-sol* (computed on odd-sol validation set) and 0.086 *odd-sol* (computed on even-sol validation set) which is similar to the training (*even-sol*: 0.086; *odd-sol*: 0.088) and validation loss (*even-sol*: 0.086; *odd-sol*: 0.087) and indicates that the models performs equally well on unseen data.

Several examples from the training set with their target and predicted event masks are given in the Supporting Information, Figures S4 and S5.

4 Application to InSight Data

4.1 Test on MQS Events

At first, we test MQNet on known events from the MQS catalogue to verify that the synthetics-trained model can indeed detect the real marsquake recordings. We show the example of the catalogued HF event S1038b (QC) in Figure 5 with the (a)-(b) normalised model input, (c)-(d) the predicted event and noise masks, and (e)-(h) the original signal, the extracted event and noise signals both in time-frequency and time domain. S1038b is strongly corrupted by wind noise, visible in the form of broadband noise and lander mode excitation, the 1 Hz artefact, and large spikes at low- and high frequencies. By element-wise multiplying the STFT coefficients with the predicted event and noise masks, we extract the event and noise signals, respectively. Both in time-frequency domain and time domain, the extracted event indicates the commonly observed properties of HF events (compare Fig. 2). We note that the broadband noise and in particular the ambient 2.4 Hz excitation that overlap with the event are removed and are correctly attributed to the noise.

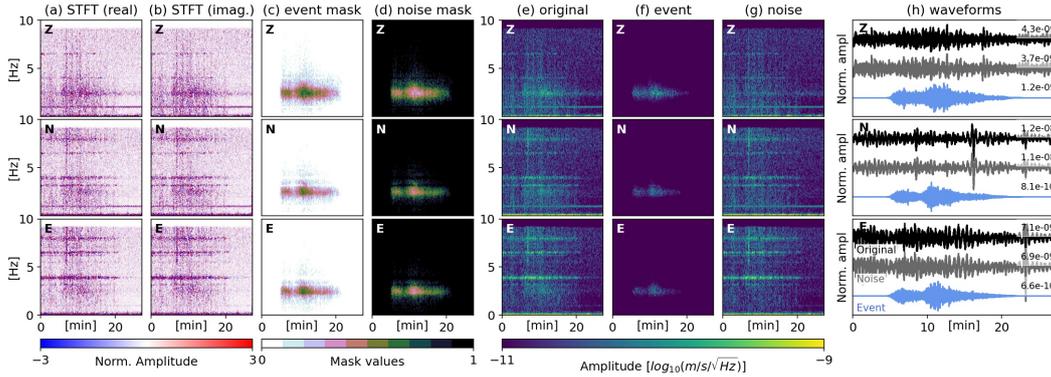


Figure 5: Example of HF event S1038b, (QC). The model inputs are the (a) real and (b) imaginary coefficients of the STFT; the model outputs are the masks for (c) event and (d) noise. The spectrogram of (e) original signal, (f) extracted event and (g) extracted noise are shown, along with the (h) three component waveforms of original signal (black), and extracted event (blue) and noise (grey). The waveforms are normalised by their maximum value (given on right side in m/s) and offset for clarity. STFT: short time Fourier transformation.

We present further successful detections of catalogued MQS events in Figure 6 for all event types and quality classes. The examples are all clearly visible in the event masks with an energy distribution that matches the typical frequency range of each event type (compare with Fig. 2). The single exception is the QD BB event S1033a, where the prediction does not indicate long period energy, but instead shows an energy distribution closer to a 2.4 Hz event. Events commonly have the highest SNR on the vertical component and should be detectable on one-component alone - with the exception of the VF event type that is often more clearly observed on the horizontal components. Using all three components allows MQNet to identify the typical high-frequency (>6 Hz) energy on the horizontal components, and consequently, the model can better discriminate VF events from wind gusts (S0712a) or find weak VF events that are barely visible on the vertical component (S1048f). This also aids in distinguishing HF and VF events, which can have a very similar spectra on the vertical component. The model also successfully detects and extracts all relevant energy from outlier events like S1094b, a BB event with remarkably broadband excitation and unusual long duration; S1073d, a HF event with three energy packages that might be associated to the superposition of two individual events; and S01048f, a VF event without 2.4 Hz energy. Since we did not generate synthetic event waveforms with these specific characteristics, their successful detection indicates the model’s ability to generalise and find unusual events.

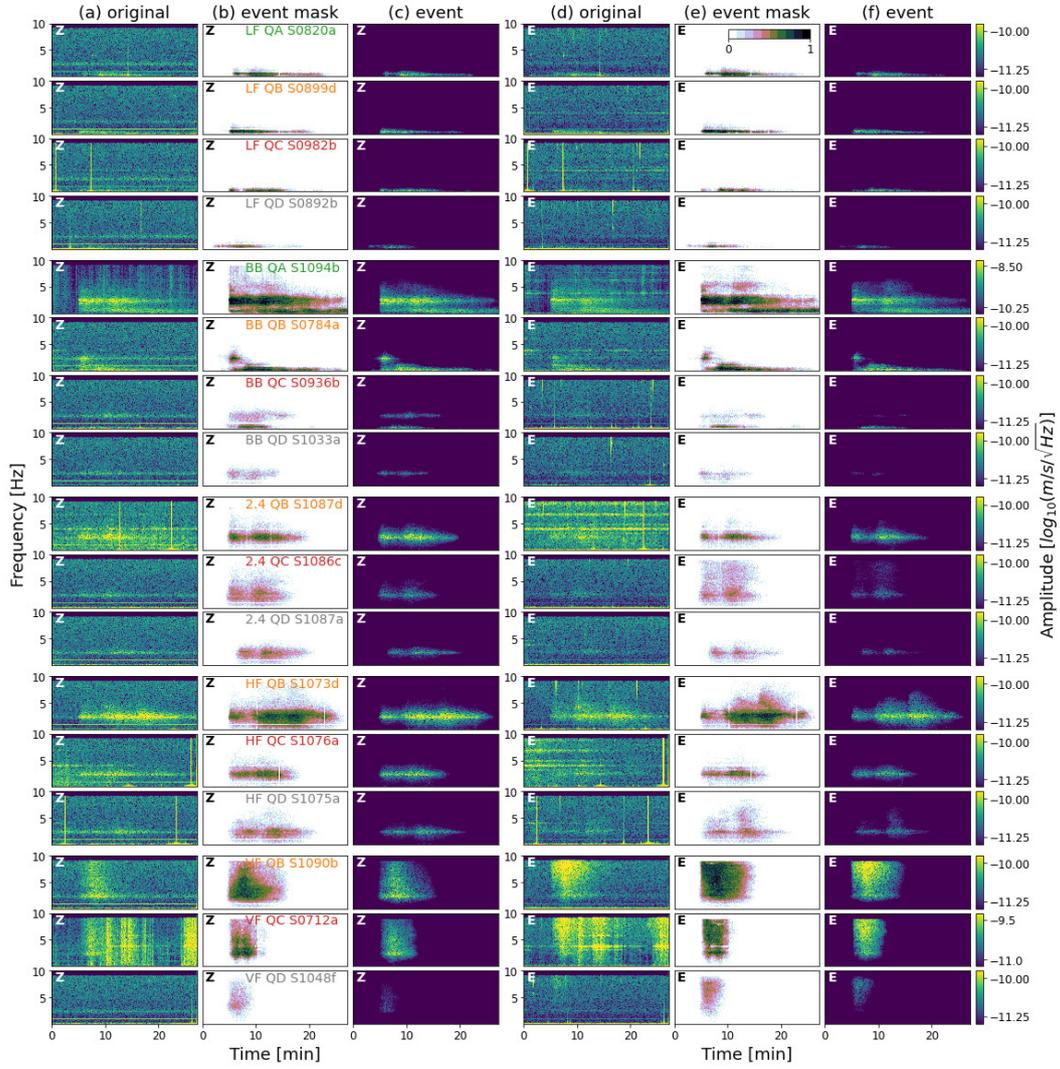


Figure 6: Examples of events for each MQS event type and quality class. The recorded data are given in (a) and (d), the predicted event masks in (b) and (e), and the extracted events in (c) and (f) for Z and E components, respectively. Labels in (b) indicate event type, quality and name in MQS catalogue. Note that the spectrogram range is fixed within each row but is different for S1094b and S0712a. MQS: Marsquake Service.

4.2 Event Detection on Continuous Data

We now apply MQNEt to the full 20 sps waveform data set, spanning 1043 sols. For each sol, we split the data into 110 time windows of 27 minutes length each, with 50% overlap. The event masks are then predicted for each of these time windows.

Figure 7 shows the procedure for Sol 923 (also shown in Fig. 1). The MQS catalogue includes six events on this sol (S0923a, QC, BB; S0923b, QD, 2.4; S0923c, QB, HF; S0923d, QC, VF; S0923e, QD, 2.4; S0923f, QC, VF), which are all clearly visible in the predicted event masks. The high-amplitude HF and the two VF events are detected in both overlapping rows, whereas the weaker events are not (fully) found when they are close to the window edge. The start and end of the event masks approximately match the MQS start and end times, which are defined as the earliest and latest possible event energy (Clinton et al., 2021). As the MQS start time includes phase pick uncertainty, it is often earlier than our detections. Most parts of this sol are (assumably) correctly identified as noise. However, there are a few weak detections visible in the event masks (and outside of the MQS event windows), which may correspond to real, low SNR events.

We can now use the predicted masks for event detection: for each time step, we sum the mask values at all frequencies and for both overlapping time windows. Thereby, we ignore small mask values (as noise does not exactly have a value of zero in the event masks), as well as mask values close to the window edge, which in rare cases can include spurious detections. This leads to a time series of added mask values (Fig. 7 (c)), from which we collect the peaks and the associated start and end times. We then compute an event-wise detection score D , by summing up the timeseries from start to end time. D is proportional to the event SNR and increases with longer duration and broader frequency content of an event signal. Catalogued MQS events on this sol have D values between 580-13809, whereas the remaining detections have much lower values ($D < 382$).

We use this procedure on the complete 20 sps data set (Sol 182-1224) and collect all detections which are above an initially very low detection threshold $D_{thresh} = 100$. This value is somewhat arbitrarily selected and we will optimise it in the next section as we review the detections. All detections are then associated either to the low- or high frequency event family, depending on whether they have more energy below 1 Hz or above 2 Hz, respectively. In addition to the event family, we provide the event type (LF, BB, 2.4/HF, VF) in manual review (see Open Research).

Since we trained the two models with noise from Sol 290-960, we alternately use the *even-sol* (trained on even-numbered sols) and *odd-sol* model (trained on odd-numbered sols) to avoid predicting on data seen during training. We remove time windows with gaps and lander or SEIS activity. Finally, events around LMST midnight that are detected twice on consecutive sols are merged into one detection.

4.3 Review of Detections and Selection of Detection Threshold

To optimise the detection threshold, we need to understand whether the new detections outside of catalogued MQS events correspond to true or false positive detections. Therefore, we select a representative part of the mission (Sol 182-600) and manually review all 1071 detections with a detection value above the initial threshold score of 100 (414 of these match MQS catalogue events). During event review, the 2.4 Hz resonance serves as an indicator for marsquakes, as it is excited by most events (Clinton et al., 2021), while the excitation of wind-driven lander modes suggest noise contamination (Dahmen, Zenhäusern, et al., 2021; Charalambous et al., 2021). However, we expect that newly identified events are generally in noisier periods where they were either overlooked or fell below the MQS STA/LTA threshold (Clinton et al., 2021), and consequently have higher noise contamination. Besides these two features, we check for other similarities with the known event types (Fig. 2) and specifically, the frequency content, duration, envelope, possible presence of primary and

secondary energy packages from P- and S-phases, and the signal energy distribution on all three components. We show examples of this review process in the Supplemental Material, Figures S6-S14 and describe the features that we consider in confirming these detections. We classify each detection as correct, unclear or false. Correct detections exhibit typical event characteristics in time-frequency domain, and false detections are clearly associated with signals of non-seismic origin (Ceylan et al., 2021), and in particular often includes an energy onset that is coincident with an increase in lander mode excitation. Unclear detections correspond to signals which we cannot confidently associate with correct detections, but which are also not obviously false, and are often associated with detections around the 2.4 Hz resonance.

Then we label the detections as true positive (TP), false positive (FP) or false negative (FN) and compute the precision ($TP/(TP+FP)$) and recall ($TP/(TP+FN)$) as a function of the detection threshold D_{thresh} . All detections matching catalogued MQS events and new correct detections are classified as true positive when their detection score $D > D_{thresh}$, and otherwise as false negatives. All new unclear or false detections are labelled as false positives when above the threshold. In Figure 8, we show the precision-recall curve and the F_1 score, which is defined as the harmonic mean of precision and recall ($2 \cdot (precision \cdot recall)/(precision + recall)$), separately for LF and HF family detections. The precision-recall curve shows the trade-off of false positives (precision) and false negatives (recall) with varying detection thresholds. The optimal threshold is indicated by the maximum of the F_1 score curve. We notice that the method detects LF family events more reliably and the optimal threshold is better defined compared to the HF family events. For the latter group, the area under precision-recall curve is smaller and the F_1 score has a lower and broader peak. Based on the F_1 score, we decide to select a detections threshold of $D_{thresh,HF} = 450$ for HF family events, which is in between the maximum and other peaks at higher detection thresholds but with similar F_1 scores. The threshold for LF family events is set to a value of $D_{thresh,LF} = 319$, which corresponds to the maximum F_1 score.

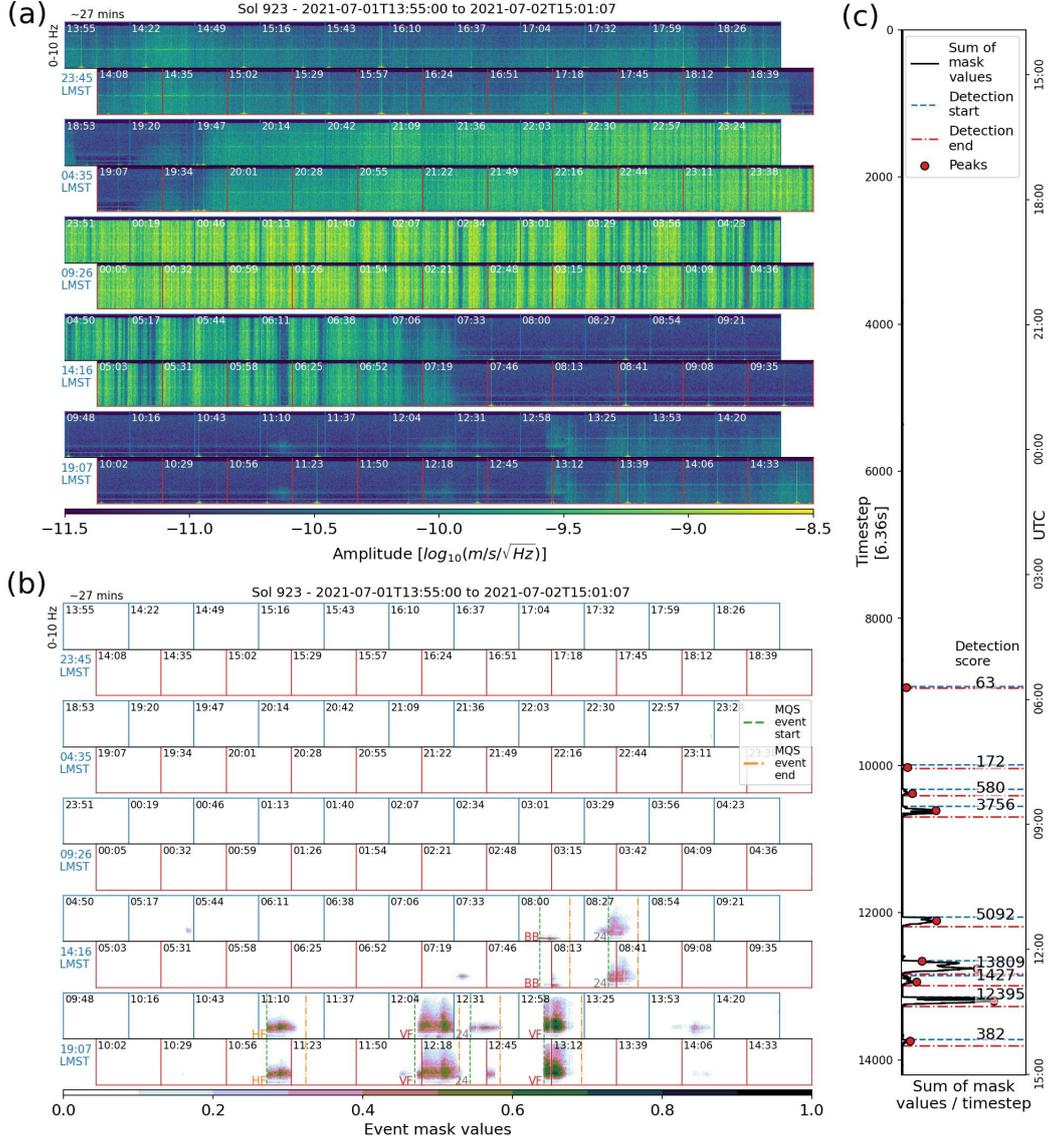


Figure 7: Event detection on Sol 923. The input data are shown in (a), represented by the Z-component, spectrograms split into 110 ~ 27 min time windows with 50% overlap (0-10 Hz); compare with Fig. 1 (a). The start time (UTC) of each window is given on top left. Indented rows (red boxes) show same data as rows directly above (blue boxes), but with a 50% time shift. Start time (LMST) of non-indented rows is given on left side (blue). Data include short overlap with previous/following sols. The corresponding predicted event masks are shown in (b) together with the catalogued MQS events (vertical dashed lines). The integration over all masks values per time step is shown in (c), annotated with the detection score D , and the start and end times for each peak (horizontal dashed lines). MQS: Marsquake Service; LMST: local mean solar time.

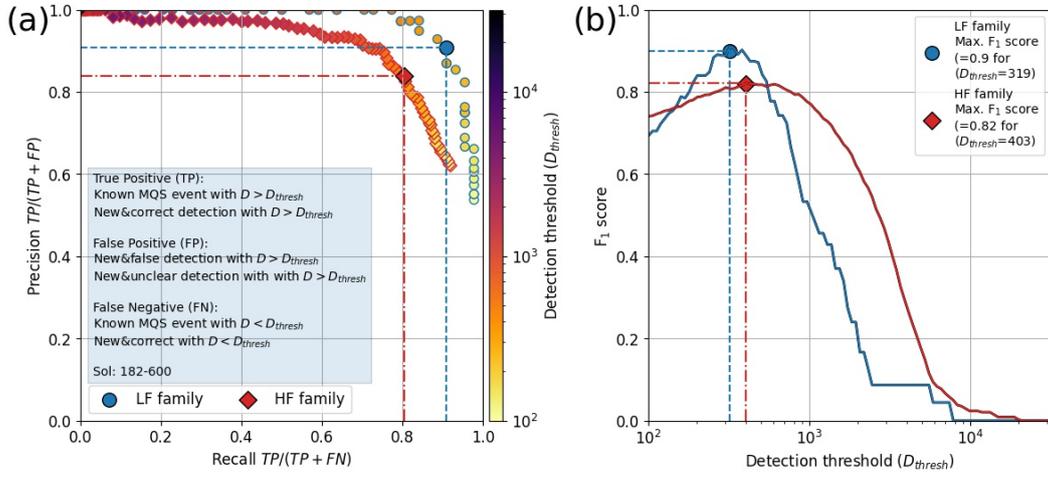


Figure 8: Detection performance as function of detection threshold D_{thresh} . Precision-recall curve in (a), and F_1 score in (b) for Sol 182-600, separately for LF and HF family events. Note that new&correct detections are included as the false negatives when below D_{thresh} . MQS: Marsquake Service; LF: low frequency; HF: high frequency.

5 Results

5.1 Detection Statistics

On the 1002 sols with data between Sol 182 - 1224 we collect 2187 detections with D scores above the selected thresholds. Of these, 2075 belong to the HF family and 112 to the LF family. First, we compare the detections to the MQS catalogue (InSight Marsquake Service, 2022) in Figure 9 (a), while distinguishing event types and qualities. Across all quality classes, the recall is $\sim 88\%$ for HF family events, and $\sim 92\%$ for LF family events. We detect all high quality (QA&B) and the majority of low quality events (QC&D). Further, we find all but five BB (1), HF (1) and VF events (3, where S1143a, QC, VF was originally detected but removed as it occurred during lander activity). The recall of MQS events is lower for LF (2 QC, 4 QD undetected) and 2.4 Hz events. Low quality events from both of these event types are usually only faintly visible and do not include clear phases (or might even be of non-seismic origin (Clinton et al., 2021)). The detected MQS events include six cases where two events overlap (e.g. nearly simultaneous LF and 2.4 Hz events). The detection scheme combines them to a single detection, although each individual event is visible in the event masks.

We then review all other detections not matching catalogued MQS events and classify them as correct, unclear and false, considering the criteria described earlier. We detect 632 additional, correct events that were missing in the MQS catalogue (precision of 86% for LF, and 81% for HF events). An additional 147 events were confirmed in review, though, they fell below our detection threshold and are not listed here.

Figure 9 (b) summarises the performance considering all detections separately for LF and HF family events: for LF detections, 68% correspond to MQS events, 18% to new correct events, and the remainder to mainly unclear (11%) and few clearly false detections (3%). When it comes to HF family events, about half of the detections match MQS events (52%), while the shares of new correct (29%) and unclear (18%) detections are higher (1% false detections).

We provide our extended catalogue, as well as sol-wise solutions (Fig. 7 (b)) for the complete processed data set in the Supporting Information (see Open Research).

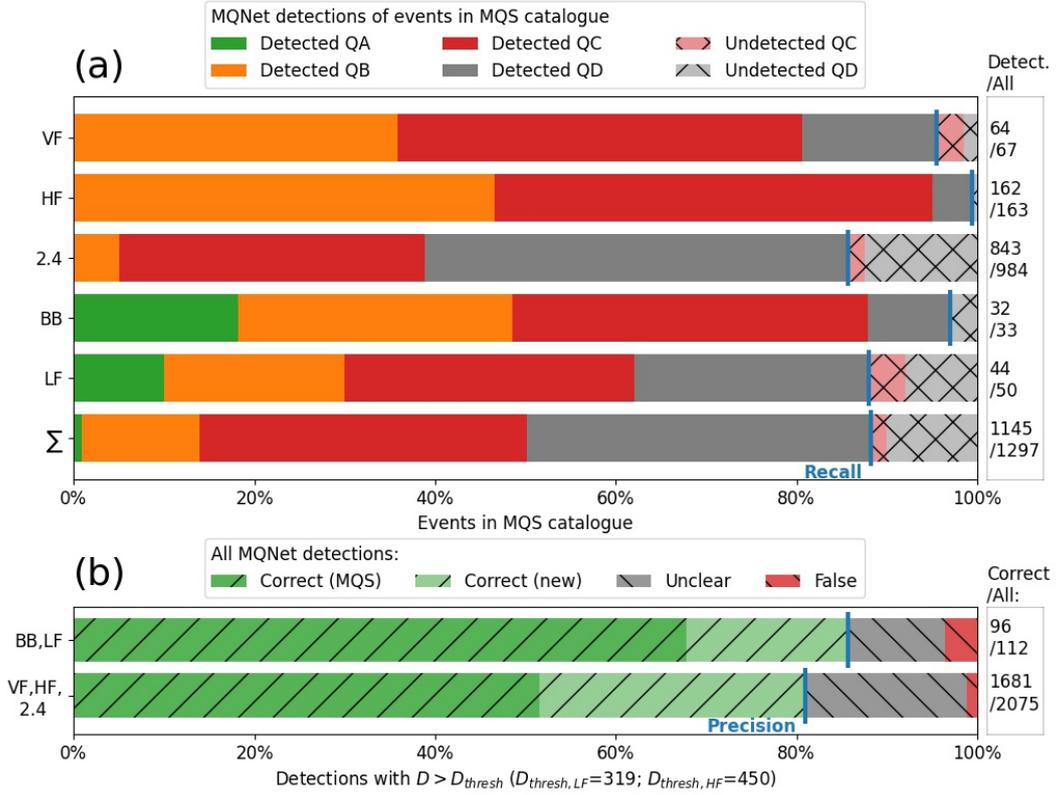


Figure 9: Detection performance from Sol 182-1224 based on chosen detection thresholds. (a) MQNet detections compared to MQS catalogue: detected and undetected events by MQS event type (very high-, high frequency, 2.4 Hz, broadband and low frequency events) and quality (Quality A-D). Here, the recall (blue marker) considers only MQS events; right side indicates the numbers of detected events by MQNet compared to the MQS catalogue. (b) Statistics of our review of all MQNet detections, separately for LF and HF family with precision (blue marker). Right side indicates number of correct detections (matching MQS event or confident new detection) compared to number of all detections.

5.2 New Detections

Similar to the event distribution in the manual MQS catalogue, most newly identified events correspond to weak 2.4 Hz events. In addition, we find 23 new LF family and 14 new VF events. Following the MQS convention, we assign all confirmed new detections an alpha-numerical name D0000x, which consists of the 4-digit sol number on which the event took place, and an alphabetical character x to uniquely identify each events.

Figure 10 presents examples from our new detections: we find several VF events which are more easy to miss in manual review, as they are only visible on the horizontal components (D0774a) or in between heavy wind gusts (D0497a). New 2.4/HF events are usually weak (D1027c), in windier periods (D1028a), or contaminated by glitches (D1068b). Added LF events are weak (D0190a, D0345a) or in periods with increased noise levels, but still distinctly different from wind noise (D0943a, see Supporting Information, Figures S6-S14).

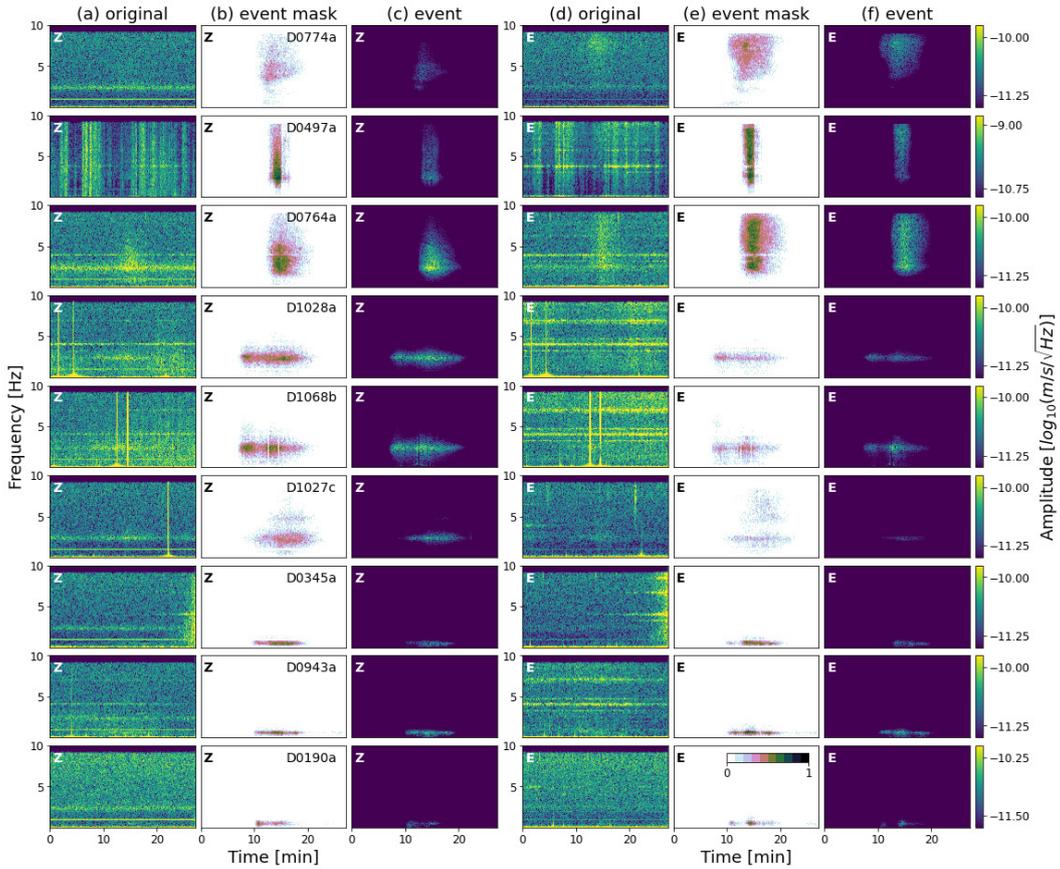


Figure 10: Examples of new events found by MQNet but missing from MQS catalogue. The recorded data are given in (a) and (d), the predicted event masks in (b) and (e), and the extracted events in (c) and (f) for Z and E components, respectively. Potential new VF events D0774a, D0497a, D0764a, 2.4 Hz/HF events D1028a, D1068b, D1027c, and LF events D0345a, D0943a, D0190a. Note that the spectrogram range is fixed within each row but adapted for each event; D0774a and D0190a were added to the MQS catalogue following this analysis, and are labelled as S0774a and S0190a.

5.3 Temporal Distribution of Events

Figure 11 summarises all of our new and confirmed detections that are above and below our detection threshold (D_{thresh}), together with the catalogued MQS events from Sol 182-1224 (see Open Research for extended catalogue). Similar to the distribution of MQS events, the majority of our new detections are in the quiet evenings. We do though observe new events in periods with increased noise levels, in particular, the nights during spring and summer of the first (Sol 182-850) and second Martian year (from Sol 850). However, very few new events are found during the high noise periods during the day, or throughout fall and winter, when the noise levels are much higher than commonly observed event amplitudes.

The cumulative event count shows substantive changes throughout the mission, which are clearly influenced by the seasonal noise pattern. However, it appears that event rates tail off already before the high noise season starts around Sol 487 and Sol 1156 (start of fall), and increase only many sols after its end around Sol 783 (start of spring). Further, we see a considerable increase in the event rate in the second Martian spring/summer period. This is despite a highly repetitive noise pattern for both years - apart from glitches, which are actually much stronger in the second year.

Figure 12 shows the cumulative HF family event count for MQS catalogue events and our new detections for each Martian year. There are two periods with significant data gaps in each year (Sol 267-288 and Sol 1107-1129), and the rates are given both including and correcting for the gap by linear interpolation of the event rate assuming a similar rate as in the preceding and following ten sols. The HF events in the MQS catalogue shows a seasonality in rate that is distinct from the noise pattern, as well as a substantial rate increase in the second year (+66% during Sol 850-1224 compared to Sol 182-556). Both of these observations are corroborated by the new detections, in fact the rate increase in the second year is even more pronounced when our new detections are included (+88%).

There are also few more LF family events in the second Martian year, but the difference does not appear significant (Supporting Information, Fig. S15).

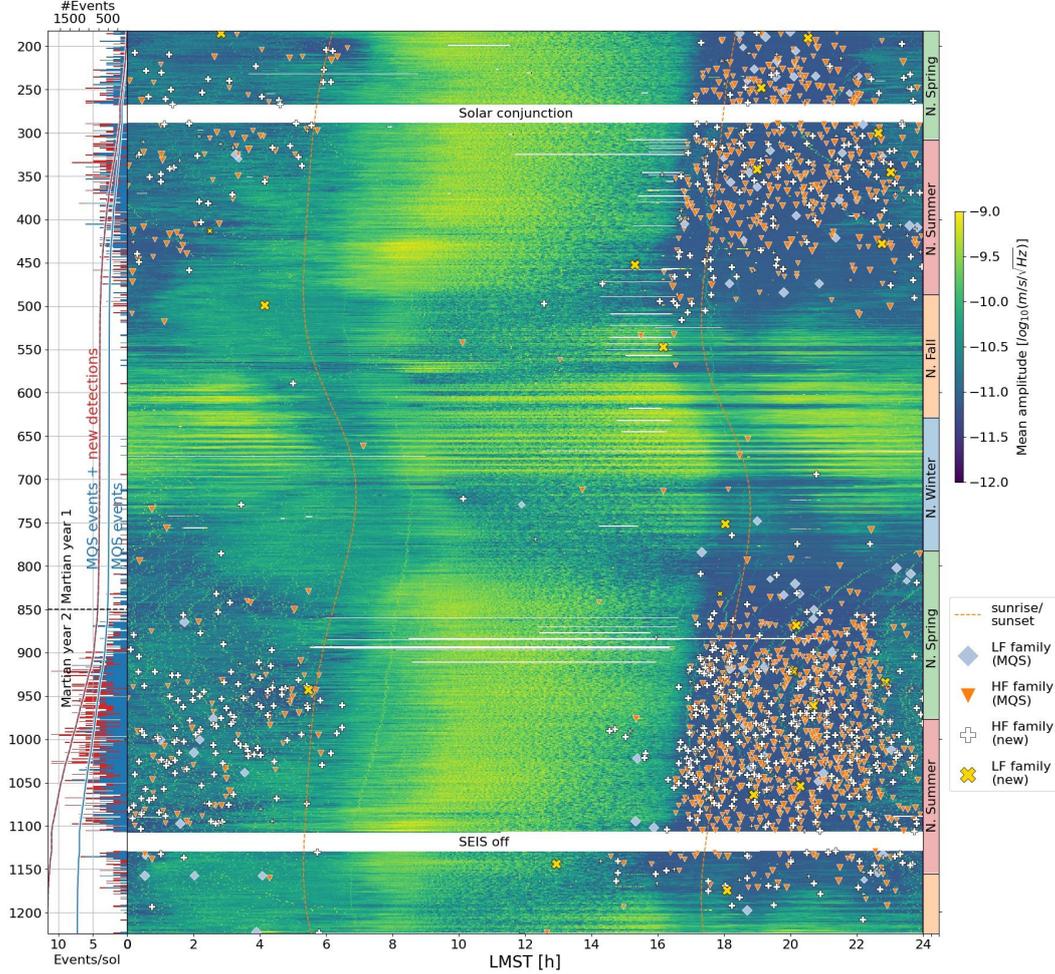


Figure 11: Overview of our new detections and MQS events. Shown are the new, correct LF and HF family detections (small/large markers indicate events below/above threshold), together with MQS catalogue LF and HF family events (small/large markers for low/high quality events). The background represents the mean noise level on the vertical component between 0.1-4.5 Hz. The left side indicates the event count per sol and the cumulative event numbers (MQS: blue, new: red). The start of the second Martian year is indicated, corresponding to 668 sols after start of plot / beginning of continuous 20 sps data. Martian seasons on Northern hemisphere are given on right side. MQS: Marsquake Service; LF: low frequency; HF: high frequency, SEIS: Seismic Experiment for Interior Structure; LMST: local mean solar time.

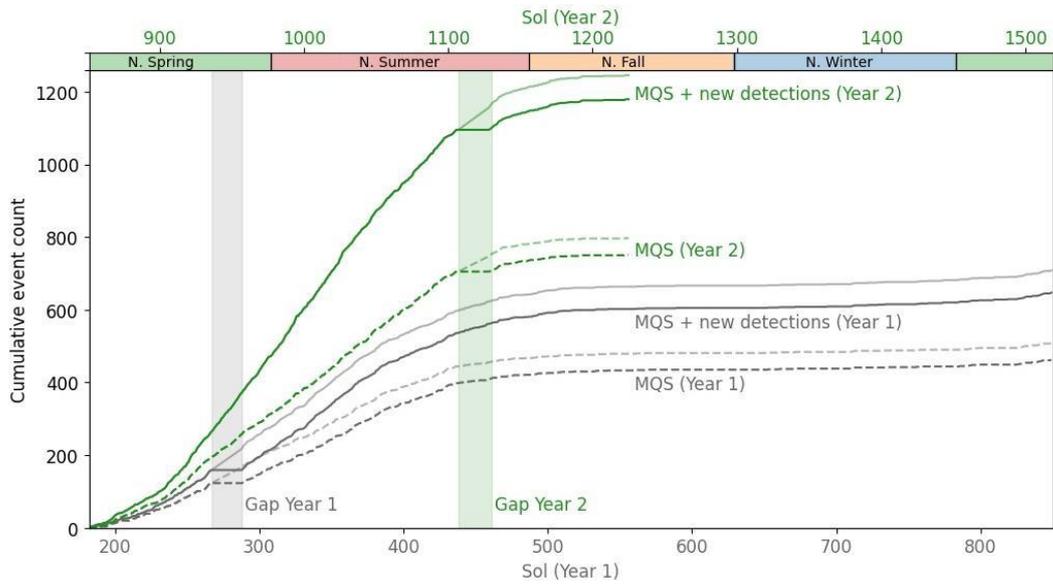


Figure 12: Cumulative event count for HF family events. Shown are the cumulative HF event numbers of all new, correct detections (above detection threshold) and MQS events for first (solid-grey line, Sol 182-850) and second Martian year (solid-green line, from Sol 850); dashed lines show numbers in MQS catalogue alone. Faint-coloured lines include a correction of the event numbers during the gaps in the first and second year (by linear interpolation of event rate). HF: high frequency. MQS: Marsquake Service.

An alternative view on the events and the changing background noise is given in Figure 13, where we show the noise variations represented by the daily 10th, 33rd, 67th and 90th spectral amplitude percentiles, separately for the LF (0.2-0.8 Hz) and HF (2.2-2.6 Hz) bands. Detected MQS events and our new events are given with their denoised spectral amplitudes. For their calculation, we compute the denoised event spectrograms and concatenate consecutive time windows for detections spanning multiple windows. The spectral amplitudes are then selected as the 90th percentile in the LF and HF frequency bands, which is less affected by potential remnants of extreme amplitude glitches providing more stable values. The figure follows Fig. 3 of Clinton et al. (2021), but since we report denoised spectral amplitude at the 90th percentile and that figure uses the time domain peak value of the bandpass filtered data, manually selected to avoid spurious peaks, our reported values are lower.

No new very significant amplitude event is discovered for either the LF or HF families. The new LF detections have relatively low amplitude, though the amplitudes are similar to other low-amplitude MQS events. Similarly, our new HF detections are overall lower in amplitude than MQS events, but share a wide overlapping range. A subset of them is found in the quietest times and are weaker than MQS events. However, many new detections with moderately high amplitudes are made in the night periods during increased noise levels (compare Fig. 11). Extremely low amplitudes might correspond to either very weak detections or events without detected energy in the LF or HF bands, e.g. VF events with little 2.4 Hz energy (S1048f, Fig. 6).

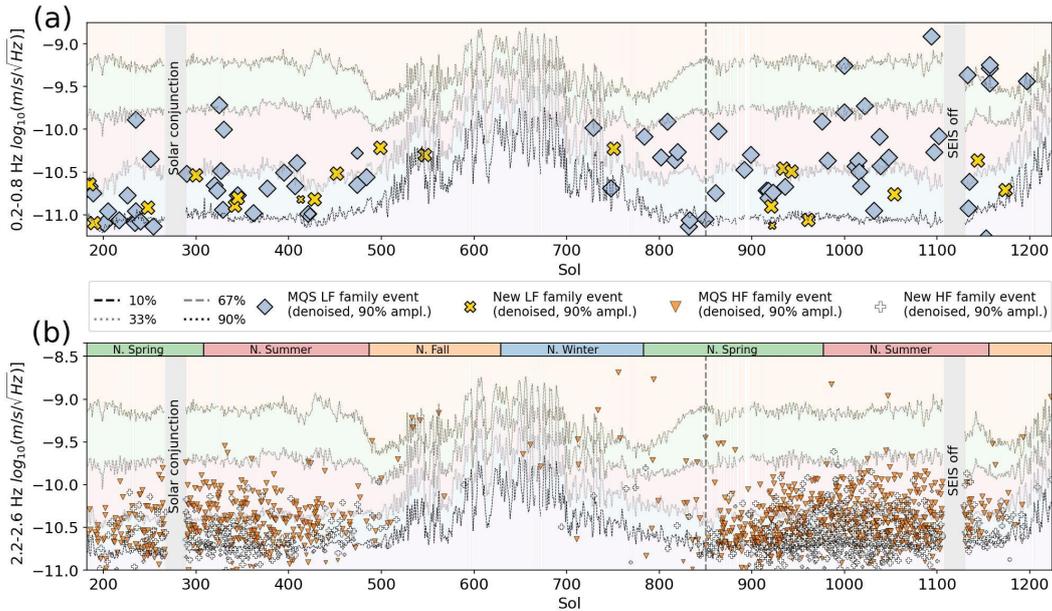


Figure 13: Event amplitudes and background noise variations. (a) Newly-identified and detected MQS LF family events with denoised 90th percentile amplitude (Z), together with the background noise represented by the daily 10th, 33rd, 67th, and 90th percentile of the spectral amplitude (Z) between 0.2-0.8 Hz (small/large markers for events below/above threshold). (b) same plot for HF family events showing the background noise around the 2.4 Hz resonance. Note that some events have amplitudes below or above the shown range, among them are the MQS VF events S0976b (-7.85) and S1034a (-7.69), as well as BB S1222a (-7.32, all in $\log_{10}(m/s/\sqrt{Hz})$). HF: high frequency; LF: low frequency.

6 Discussion

Detecting and analysing marsquakes is challenging because of their low amplitudes and relatively high and fluctuating background noise. We choose a deep learning network for this tasks, which we train on a dataset of synthetic event waveforms, because the number of available recorded waveforms is small. We combine synthetic event waveforms with recorded InSight noise, which allows the network to learn the full range of signal characteristics.

We design synthetic event waveforms to have similar envelopes and spectral contents to those known for observed event types (Clinton et al., 2021) using a modified version of the stochastic waveform modelling (Boore, 2003). This allows us to replicate various features of marsquakes, in particular, variable coda shapes, weak and emergent onsets, high frequency event energy, and narrow to wide spectral peaks. The synthetic event waveforms have a constant frequency spectrum throughout the event duration (apart from our synthetic BB events) and do not include dispersive signals. These are not commonly observed for marsquakes, but could potentially be added in the future for a targeted search.

MQNet shows a high detection performance, even though it is exclusively trained with relatively simplistic synthetic event waveforms. We notice also that the network detects unusual events that differ considerably from our synthetic samples (Fig. 6), demonstrating the networks ability to recognise event energy on various scales in time-frequency domain and to generalise to unseen event recordings. We optimise the UNet architecture for the detection of event and noise energy, and by working in time-frequency domain, we can detect events that are often far below the noise level in time domain (Fig. 5).

In the training data set, periods outside of catalogued MQS events are labelled as noise, even though they may still include undetected events. In practise, this appears to be less problematic, as the network does not seem to overfit the data. Further, we train two models that include noise only from either even or odd-numbered sols, to avoid predicting on data seen during training.

To find marsquakes on the continuous data, we collect detections from the event masks, which is faster and more stable compared to the detection on denoised time series. This is due to the wide dynamic range of the background noise, which means, predicted mask values close to zero can still produce high amplitudes during extremely noisy periods and might partly remain in the denoised time series. Alternatively, the UNet architecture could also be modified to produce a time series with labels for event and noise. However, we prefer the current setup as it allows us to use the same model for event detection and denoising.

A concern with a manually curated catalogue is, over time, methods and quality can change leading to inconsistencies. The product of this study is the first automatically generated, but manually reviewed marsquake catalogue over the complete 20 sps waveform dataset (see Open Research). This new catalogue extends the manually curated MQS catalogue (1297 events between Sol 182-1224) by 779 previously undetected events (632 above-, 147 below threshold D_{thresh}), most of which have low SNR. We have reviewed all detections manually, first, to optimise the detections threshold on parts of the mission, and, second, to confirm or reject all other detections. During review, we follow a conservative approach and only confirm detections showing characteristics matching the different event types (Figs. S6-S14). This way, a considerable number of detections (383) are labelled as unclear, and are treated in the detection statistics as false positive detections. It is plausible that a substantial fraction of these events are in fact correct detections, although this is difficult to verify.

The share of unclear detections is larger for high frequency events (Figs. 8, 9) and likely arises from the presence of the ambient 2.4 Hz resonance, which itself shows some diurnal and seasonal variations (Dahmen, Zenhäusern, et al., 2021), and may complicate the detection of weak 2.4 Hz events and their discrimination from background noise. Unclear detections might include additional events with even lower SNR, which is also supported by a similar

temporal distribution across the mission as correct detections (Supporting Information, Fig. S16). Excepted from this is the period between Sol 700-800, in which the number of unclear detections seems to be somewhat increased. This coincides with a part of the mission with stronger ambient excitation of the 2.4 Hz resonance compared to other periods (e.g., see ambient 2.4 Hz amplitude in Fig. S8 as against Figs. S13 and S14).

We compare our detections against the carefully manually curated MQS catalogue (Clinton et al., 2021) and find that our MQNet detection method reproduces all high quality events and the majority of the low quality events (Fig. 9), with the best performance for HF, BB and VF events. The recall is lower for low quality LF and 2.4 Hz events, but is still high considering that these events are often extremely weak and just above the background noise. Our analysis suggests that the MQS catalogue is nearly complete for high SNR events (apart from a few clear events, e.g. D0764a, Fig. 10). Despite this careful data review by MQS, we can still add a substantial amount of low SNR events to the catalogue, many of which are similar to MQS events in amplitude (Fig. 13).

Comparing the noise level to the expected atmospherically-induced noise (Charalambous et al., 2021) has proven to be a valuable tool for event identification and confirmation, although this approach cannot easily be automated for event detection as it does not consider other noise sources. In general, any detection method applicable for the whole mission would have to be based on seismic data alone, since atmospheric measurements are only very sparsely available in the second Martian year. Our detection method is only limited by the availability of 20 sps data set from the VBB sensor - but could also be modified to run on 10 sps data and recordings from the Short-Period sensor after retraining.

Sun and Tkalčić (2022) reported the discovery of 47 new LF events found during both the low- and high noise periods with template matching techniques using the S-wave waveform of 9 strong reference marsquakes. Neither MQNet nor a manual review, using the same criteria used to confirm MQNet detections described in this manuscript, can confirm the large majority of these reported detections. The exceptions are two reported events 24 s and 39 s after S0173a (MQS P-pick) that we interpret as surface-reflected arrivals of S0173a (Durán et al., 2022), and a third proposed detection on Sol 356 that coincides with our newly-identified HF family event (D0356c: but, we do not observe the reported LF energy).

We use a consistent approach for event detection across the mission, nevertheless, the event detectability remains strongly influenced by the varying noise conditions (Fig. 11) and only few events are found during periods with noise levels far above commonly observed event amplitudes. In addition, we observe the absence of high frequency events during relatively low noise periods, that should allow for the detection of at least strong events (around Sol 500, 800 and 1150, Fig. 13). This is in agreement with the results of Knapmeyer et al. (2021) who investigated the changes in HF event rate in the MQS catalogue, while considering varying detection efficiency over time, and proposed that the seasonal changes of solar illumination are the most likely driver of HF event activity. Their study extends up to Sol 627, but includes a forecast until Sol 1200 that is generally consistent with our observations on the onset and cease of HF activity. However, the strong increase in HF rate in the second season remains unexplained (Fig. 12). Since our detection method is fully automated, our analysis shows that the rate increase is not caused by different practices in the manual data review by MQS over time, but is a real feature of the data set. At this time, no explanation for such an interannual variation in seismicity is known to us and it remains a genuinely unexpected observation about Mars by SEIS. It remains to be seen how much more data the InSight seismometer will record and whether we will get observations in the next Martian year that would further constrain the variation in seismic activity rate.

7 Conclusion and Outlook

We have demonstrated how deep learning techniques can successfully be used to detect marsquakes in the challenging InSight seismic data set. The high performance of MQNet stems from the augmentation of the training data with synthetic event waveforms, encouraging a similar approach on other sparsely available labelled datasets. We use the network to automatically reproduce the MQS catalogue and extend it by 60%. As a result, we observe and confirm distinct changes in the HF event rate throughout the mission, while the LF event rate remains relatively constant and seems to be modulated by the background noise alone. Further, the network enables us to remove both atmospheric and instrumental noise contamination from the events and report an estimate of their amplitudes. Event denoising looks promising as it allows not only the removal of noise at frequencies outside of the event's bandwidth (as it is possible e.g. with a bandpass filter), but also for noise overlapping in time-frequency domain. Future work will focus on exploring how the denoising can further support phase picking and the estimation of magnitudes and back azimuths, which could advance the characterisation and location of a large number of events in the MQS and our extended catalogue.

Acknowledgments

We acknowledge NASA, CNES, their partner agencies and institutions (UKSA, SSO, DLR, JPL, IPGP-CNRS, ETHZ, IC, MPS-MPG) and the flight operations team at JPL, SISMOC, MSDS, IRIS-DMC and PDS for providing SEED SEIS data. We acknowledge funding from (1) Swiss State Secretariat for Education, Research and Innovation (SEFRI project “MarsQuake Service-Preparatory Phase”), (2) ETH Research grant ETH-06 17-02, and (3) ETH+02 19-1: Planet MARS. The Swiss contribution in implementation of the SEIS electronics was made possible through funding from the federal Swiss Space Office (SSO), the contractual and technical support of the ESA-PRODEX office, (4) A. E. Stott is supported by funding from the French Space Agency CNES and ANR (ANR-19-CE31-0008-08). This paper is InSight Contribution Number 257.

Open Research

The InSight seismic event catalogue version 12 (InSight Marsquake Service, 2022) (to be released 2022-10-01), waveform and station metadata data are available from IRIS-DMC and IPGP Datacenter (InSight Mars SEIS Data Service, 2019), as are previous catalogue versions.

The supporting Information contains additional material on the MQS catalogue, details on the network selection and optimisation, examples from the review process of newly-identified events, and temporal distribution of MQNet detections and LF family events across the mission. Our extended event catalogue, as well as our sol-wise solution for the complete processed data set will be provided on public repository (currently available:).

The UNet architecture was implemented in TensorFlow (Abadi et al., 2015) and the network figure (Fig. 3) was created with (Haris Iqbal., 2018). Synthetic envelopes were computed from the Instaseis (van Driel et al., 2015) based on Mars model InSight_KKS21GP (Stähler et al., 2021), accessible at http://instaseis.ethz.ch/marssynthetics/InSight_KKS21_GP/ (last accessed June 2022). The data were processed with ObsPy (Beyreuther et al., 2010), NumPy (Harris et al., 2020) and Scipy (Jones et al., 2001), and visualisations were created with Matplotlib (Hunter, 2007).

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., . . . Zheng, X. (2015). *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Retrieved from <https://www.tensorflow.org/> (Software available from tensorflow.org)
- Allen, R. (1982). Automatic phase pickers: Their present use and future prospects. *Bulletin of the Seismological Society of America*, 72(6B), S225–S242.
- Anderson, D. L., Miller, W., Latham, G., Nakamura, Y., Toksöz, M., Dainty, A., . . . Knight, T. (1977). Seismology on Mars. *Journal of Geophysical Research*, 82(28), 4524–4546.
- Banerdt, W. B., Smrekar, S., Banfield, D., Giardini, D., Golombek, M., Johnson, C., . . . Wicczorek, M. (2020). Initial results from the InSight mission on Mars. *Nature Geoscience*, 13, 1–7. doi: 10.1038/s41561-020-0544-y
- Barkaoui, S., Lognonné, P., Kawamura, T., Stutzmann, É., Seydoux, L., de Hoop, M. V., . . . others (2021). Anatomy of Continuous Mars SEIS and Pressure Data from Unsupervised Learning. *Bulletin of the Seismological Society of America*, 111(6), 2964–2981.
- Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann, J. (2010). ObsPy: A Python toolbox for seismology. *Seismological Research Letters*, 81(3), 530–533.
- Boore, D. M. (2003). Simulation of ground motion using the stochastic method. *Pure and applied geophysics*, 160(3), 635–676.
- Ceylan, S., Clinton, J. F., Giardini, D., Böse, M., Charalambous, C., Van Driel, M., . . . others (2021). Companion guide to the marsquake catalog from InSight, Sols 0–478:

- Data content and non-seismic events. *Physics of the Earth and Planetary Interiors*, 310, 106597.
- Chandna, P., Miron, M., Janer, J., & Gómez, E. (2017). Monoaural audio source separation using deep convolutional neural networks. In *International conference on latent variable analysis and signal separation* (pp. 258–266).
- Charalambous, C., Stott, A. E., Pike, W., McClean, J. B., Warren, T., Spiga, A., ... others (2021). A comodulation analysis of atmospheric energy injection into the ground motion at InSight, Mars. *Journal of Geophysical Research: Planets*, 126(4), e2020JE006538.
- Chatain, A., Spiga, A., Banfield, D., Forget, F., & Murdoch, N. (2021). Seasonal variability of the daytime and nighttime atmospheric turbulence experienced by InSight on Mars. *Geophysical Research Letters*, 48(22), e2021GL095453.
- Civilini, F., Weber, R., Jiang, Z., Phillips, D., & Pan, W. D. (2021). Detecting moonquakes using convolutional neural networks, a non-local training set, and transfer learning. *Geophysical Journal International*, 225(3), 2120–2134.
- Clinton, J., Ceylan, S., van Driel, M., Giardini, D., Stähler, S. C., Böse, M., ... others (2021). The Marsquake catalogue from InSight, sols 0–478. *Physics of the Earth and Planetary Interiors*, 310, 106595.
- Clinton, J., Giardini, D., Böse, M., Ceylan, S., Van Driel, M., Euchner, F., ... others (2018). The Marsquake service: Securing daily analysis of SEIS data and building the Martian seismicity catalogue for InSight. *Space Science Reviews*, 214(8), 133.
- Compaire, N., Margerin, L., Garcia, R. F., Pinot, B., Calvet, M., Orhand-Mainsant, G., ... others (2021). Autocorrelation of the ground vibrations recorded by the SEIS-InSight seismometer on Mars. *Journal of Geophysical Research: Planets*, 126(4), e2020JE006498.
- Compaire, N., Margerin, L., Monnereau, M., Garcia, R. F., Lange, L., Calvet, M., ... others (2022). Seasonal variations of subsurface seismic velocities monitored by the SEIS-InSight seismometer on Mars. *Geophysical Journal International*, 229(2), 776–799.
- Dahmen, N., Clinton, J. F., Ceylan, S., van Driel, M., Giardini, D., Khan, A., ... others (2021). Super high frequency events: A new class of events recorded by the InSight seismometers on Mars. *Journal of Geophysical Research: Planets*, 126(2), e2020JE006599.
- Dahmen, N., Zenhäusern, G., Clinton, J. F., Giardini, D., Stähler, S. C., Ceylan, S., ... others (2021). Resonances and lander modes observed by insight on Mars (1–9 Hz). *Bulletin of the Seismological Society of America*, 111(6), 2924–2950.
- Durán, C., Khan, A., Ceylan, S., Zenhäusern, G., Stähler, S., Clinton, J., & Giardini, D. (2022). Seismology on Mars: An analysis of direct, reflected, and converted seismic body waves with implications for interior structure. *Physics of the Earth and Planetary Interiors*, 325, 106851.
- Giardini, D., Lognonné, P., Banerdt, W. B., Pike, W. T., Christensen, U., Ceylan, S., ... others (2020). The seismicity of Mars. *Nature Geoscience*, 13(3), 205–212.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press. (<http://www.deeplearningbook.org>)
- Haris Iqbal. (2018). *HarisIqbal88/PlotNeuralNet v1.0.0 (v1.0.0)*. Code. (<https://doi.org/10.5281/zenodo.2526396>)
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., ... Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825), 357–362. Retrieved from <https://doi.org/10.1038/s41586-020-2649-2> doi: 10.1038/s41586-020-2649-2
- Hobiger, M., Hallo, M., Schmelzbach, C., Stähler, S., Fäh, D., Giardini, D., ... others (2021). The shallow structure of Mars at the InSight landing site from inversion of ambient vibrations. *Nature communications*, 12(1), 1–13.
- Hunter, J. D. (2007). Matplotlib: A 2D graphics environment [Editorial Material]. *Comput. Sci. Eng.*, 9(3), 90–95.
- Hurst, K., Fayon, L., Knapmeyer-Endrun, B., Schmelzbach, C., van Driel, M., Ervin, J.,

- ... others (2021). Resonances of the InSight seismometer on Mars. *Bulletin of the Seismological Society of America*, 111(6), 2951–2963.
- InSight Mars SEIS Data Service. (2019). *SEIS raw data, Insight Mission. IGP, JPL, CNES, ETHZ, ICL, MPS, ISAE-Supaero, LPG, MFSC.* (<https://doi.org/10.18715/SEIS.INSIGHT.XB.2016>)
- InSight Marsquake Service. (2022). *Mars Seismic Catalogue, InSight Mission; V11 2022-07-01. ETHZ, IGP, JPL, ICL, Univ Bristol. Dataset.* (<https://doi.org/10.12686/a17>)
- Jansson, A., Humphrey, E., Montecchio, N., Bittner, R., Kumar, A., & Weyde, T. (2017). Singing voice separation with deep u-net convolutional networks.
- Jones, E., Oliphant, T., Peterson, P., & Others. (2001). *SciPy: Open source scientific tools for Python.* (<http://www.scipy.org/>)
- Khan, A., Ceylan, S., van Driel, M., Giardini, D., Lognonné, P., Samuel, H., ... others (2021). Upper mantle structure of Mars from InSight seismic data. *Science*, 373(6553), 434–438.
- Khan, A., Sossi, P., Liebske, C., Rivoldini, A., & Giardini, D. (2022). Geophysical and cosmochemical evidence for a volatile-rich Mars. *Earth and Planetary Science Letters*, 578, 117330.
- Kim, D., Davis, P., Lekić, V., Maguire, R., Compaire, N., Schimmel, M., ... others (2021). Potential pitfalls in the analysis and structural interpretation of seismic data from the Mars InSight mission. *Bulletin of the Seismological Society of America*, 111(6), 2982–3002.
- Kim, D., Lekić, V., Irving, J. C., Schmerr, N., Knapmeyer-Endrun, B., Joshi, R., ... others (2021). Improving constraints on planetary interiors with PPS receiver functions. *Journal of Geophysical Research: Planets*, 126(11), e2021JE006983.
- Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980.*
- Knapmeyer, M., Stähler, S. C., Daubar, I., Forget, F., Spiga, A., Pierron, T., ... others (2021). Seasonal seismic activity on Mars. *Earth and Planetary Science Letters*, 576, 117171.
- Knapmeyer-Endrun, B., & Hammer, C. (2015). Identification of new events in Apollo 16 lunar seismic data by Hidden Markov Model-based event detection and classification. *Journal of Geophysical Research: Planets*, 120(10), 1620–1645.
- Knapmeyer-Endrun, B., Panning, M. P., Bissig, F., Joshi, R., Khan, A., Kim, D., ... others (2021). Thickness and structure of the martian crust from InSight seismic data. *Science*, 373(6553), 438–443.
- Ksanfomaliti, L., Zubkova, V., Morozov, N., & Petrova, E. (1982). Microseisms at the VENERA-13 and VENERA-14 Landing Sites. *Soviet Astronomy Letters*, 8, 241.
- Latham, G. V., Ewing, M., Press, F., Sutton, G., Dorman, J., Nakamura, Y., ... Duennebier, F. (1970). Passive seismic experiment. *Science*, 167(3918), 455–457.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Li, J., Beghein, C., Wookey, J., Davis, P., Lognonné, P., Schimmel, M., ... Banerdt, W. B. (2022). Evidence for crustal seismic anisotropy at the InSight lander site. *Earth and Planetary Science Letters*, 593, 117654.
- Lognonné, P., Banerdt, W. B., Giardini, D., Pike, W. T., Christensen, U., Laudet, P., ... Wookey, J. (2019). SEIS: Insight’s seismic experiment for internal structure of Mars. *Space Science Reviews*, 215(1), 12.
- Lognonné, P., Banerdt, W. B., Pike, W. T., Giardini, D., Christensen, U., Garcia, R. F., ... Zweifel, P. (2020, February). Constraints on the shallow elastic and anelastic structure of Mars from InSight seismic data. *Nature Geoscience*. doi: 10.1038/s41561-020-0536-y
- Mimoun, D., Murdoch, N., Lognonné, P., Hurst, K., Pike, W. T., Hurley, J., ... Banerdt, W. B. (2017). The noise model of the SEIS seismometer of the InSight mission to Mars. *Space Science Reviews*, 211(1), 383–428.
- Mousavi, S. M., Ellsworth, W. L., Zhu, W., Chuang, L. Y., & Beroza, G. C. (2020).

- Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking. *Nature communications*, *11*(1), 1–12.
- Mousavi, S. M., Zhu, W., Sheng, Y., & Beroza, G. C. (2019). CRED: A deep residual network of convolutional and recurrent units for earthquake signal detection. *Scientific reports*, *9*(1), 1–14.
- Murdoch, N., Alazard, D., Knapmeyer-Endrun, B., Teanby, N., & Myhill, R. (2018). Flexible mode modelling of the InSight lander and consequences for the SEIS instrument. *Space Science Reviews*, *214*(8), 1–24.
- Perol, T., Gharbi, M., & Denolle, M. (2018). Convolutional neural network for earthquake detection and location. *Science Advances*, *4*(2), e1700578.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *International conference on medical image computing and computer-assisted intervention* (pp. 234–241).
- Ross, Z. E., Meier, M.-A., Hauksson, E., & Heaton, T. H. (2018). Generalized Seismic Phase Detection with Deep Learning. *Bulletin of the Seismological Society of America*, *108*(5A), 2894–2901.
- Scholz, J.-R., Widmer-Schmidrig, R., Davis, P., Lognonné, P., Pinot, B., Garcia, R. F., ... others (2020). Detection, analysis, and removal of glitches from InSight’s seismic data from Mars. *Earth and Space Science*, *7*(11), e2020EA001317.
- Stähler, S. C., Khan, A., Banerdt, W. B., Lognonné, P., Giardini, D., Ceylan, S., ... others (2021). Seismic detection of the martian core. *Science*, *373*(6553), 443–448.
- Stott, A. E., Charalambous, C., Warren, T. J., Pike, W. T., Myhill, R., Murdoch, N., ... others (2021). The Site Tilt and Lander Transfer Function from the Short-Period Seismometer of InSight on Mars. *Bulletin of the Seismological Society of America*, *111*(6), 2889–2908.
- Sun, W., & Tkalčić, H. (2022). Repetitive marsquakes in Martian upper mantle. *Nature communications*, *13*(1), 1–9.
- Tibi, R., Hammond, P., Brogan, R., Young, C. J., & Koper, K. (2021). Deep Learning Denoising Applied to Regional Distance Seismic Data in Utah. *Bulletin of the Seismological Society of America*, *111*(2), 775–790.
- van Driel, M., Ceylan, S., Clinton, J. F., Giardini, D., Horleston, A., Margerin, L., ... others (2021). High-frequency seismic events on Mars observed by InSight. *Journal of Geophysical Research: Planets*, *126*(2), e2020JE006670.
- van Driel, M., Krischer, L., Stähler, S. C., Hosseini, K., & Nissen-Meyer, T. (2015). Instaseis: Instant global seismograms based on a broadband waveform database. *Solid Earth*, *6*(2), 701–717.
- Yang, L., Liu, X., Zhu, W., Zhao, L., & Beroza, G. C. (2022). Toward improved urban earthquake monitoring through deep-learning-based noise suppression. *Science advances*, *8*(15), eabl3564.
- Zhu, W., & Beroza, G. C. (2019). PhaseNet: a deep-neural-network-based seismic arrival-time picking method. *Geophysical Journal International*, *216*(1), 261–273.
- Zhu, W., Mousavi, S. M., & Beroza, G. C. (2019). Seismic signal denoising and decomposition using deep neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, *57*(11), 9476–9488.
- Zweifel, P., Mance, D., ten Pierick, J., Giardini, D., Schmelzbach, C., Haag, T., ... others (2021). Seismic high-resolution acquisition electronics for the NASA InSight Mission on Mars. *Bulletin of the Seismological Society of America*, *111*(6), 2909–2923.