

Training martian seismologists for InSight

Benjamin Fernando, Kasra Hosseini and Maria Tsekhmistrenko describe how blind testing is helping to train martian seismologists.

The past 100 years have seen enormous advances in our knowledge of the Earth's internal structure. Today's geophysicists have at their disposal many techniques with which to probe the dynamics and composition of our planet's interior: gravimetry, geomagnetics, magnetotellurics, seismology and more. But when it comes to Mars, or indeed any of the other planets in the solar system, we know remarkably little about their interiors.

There's a simple reason for this: most detailed geophysical investigations require high quality *in situ* measurements. Mars missions are both expensive (NASA's latest attempt, InSight, will end up costing over \$800 million) and challenging (historically, around half the spacecraft launched to the Red Planet have failed). Orbital missions are certainly useful, but the lack of a global magnetic field limits the amount that can be learnt about the core from space, and while gravimetric surveys are available they are at far lower resolution than for the Earth. Thus, there is no substitute for surface experiments capable of recording precision seismic and geodetic data.

While the Apollo Moon missions of the 1970s demonstrated the science returns possible from sending seismo-meter packages to other celestial bodies, no martian mission successfully achieved this. In 1976, NASA's twin Viking spacecraft landed the first seismic packages on another planet but, unfortunately, they returned little useful data. On Viking 1 the instrument failed to deploy properly. Viking 2's seismometer did deploy, but because it was mounted on the spacecraft rather than properly coupled to the ground, it mostly recorded the shaking of the lander in the martian wind. This noise contaminated the entire dataset, enabling an upper bound to be estimated for the frequency of large marsquakes, but precluding detailed investigation of individual seismic events (Anderson *et al.* 1977).

On 26 November this year, NASA's Insight Mission, launched in May atop an Atlas V rocket from Vandenberg Air Force Base, will attempt to land on Elysium Planitia, a broad open plain near the planet's equator. Once on the surface, it aims to investigate questions that are at present unanswered, such as what the state of Mars's core is. Mars has no intrinsic magnetic field, so there is no strongly convecting metallic fluid at depth, but it is not clear when this flow stopped and hence when Mars lost its magnetic field. The magnetosphere helps to shield the atmosphere from erosion by the solar wind and so the existence, then loss, of a magnetic field would have had important consequences for the evolution of the planet's atmosphere and, potentially, its habitability. Being able to constrain the planet's evolutionary history using detailed measurements of its core properties would be invaluable to planetary scientists.

Additionally, many models of the planet's structure present a dichotomy in the heat flow and expected temperature of Mars's crust at depth. The temperature of the crust influences the planet's tectonic behaviour, for example in determining the depth at which the transition from rocks deforming under stress in a brittle manner to a ductile manner takes place. Determining the depth at which quakes nucleate would be useful in constraining the point at which this transition occurs.

Other dilemmas include how long ago the planet lost its volcanism, why extinct volcanoes are so strongly clustered, and what differentiation there is within the planet's interior. Terrestrial geophysicists are also very curious about such matters, as they offer both context and comparison to what is observed on Earth.

The InSight spacecraft

To investigate these questions, InSight carries one of the most sophisticated science packages ever launched to the Red Planet. The Heat Flow and Physical Properties Package (HP³) aims to penetrate 3–5 m into the martian crust to measure the heat flux from the planet's interior to shed light on

its evolutionary dynamics. The Seismic Experiment for Interior Structure (SEIS) contains instruments capable of measuring seismic waves at frequencies of up to 50 Hz during normal operation, and up to 100 Hz for precision measurements of local rock properties during the insertion of the HP³ instrument into the ground. Initial commissioning data from SEIS, parts of which were switched on for testing as InSight travelled to Mars, are reported to look promising.

To ensure proper coupling to the ground and to avoid the problems experienced by

.....
"InSight carries one of the most sophisticated science packages ever sent to Mars"
 Viking 2, a small robotic arm mounted on the spacecraft will lift SEIS onto the ground after landing, at which point it will be covered by a wind and thermal shield.

This should reduce the noise induced by dynamic shaking of the instrument in the wind, and by thermal heating and cooling of the instrument.

Additional wind, pressure and temperature sensors are also carried to record data that will prove useful in the post-processing removal of some of the atmospheric contamination. Other instruments being carried include two cameras and a radio experiment to track the planet's rotation, which will prove useful in more accurately determining its moment of inertia, and hence its internal structure.

Interpretation of martian seismic data

Even after satisfactory landing and deployment of the instruments, and collection of the seismic data, its interpretation presents further challenges. On Earth, after a seismic event, data from receiving stations around the world can rapidly be combined to provide estimates of the source parameters: magnitude, focal mechanism and depth. This requires analysis of the ground motions induced by the passage of different kinds of seismic waves past the receivers at different locations, combined with a global and local model of the interior of the Earth (see box "Understanding the Earth").

On Mars, however, the situation is very different and the inversion is much more challenging. InSight consists of only one receiving station, meaning that

multistation triangulation is impossible. Secondly, the *a priori* knowledge of the planet's structure is less well constrained than for the Earth, making the parameter space that must be searched when performing a fit much larger.

Additional difficulties are introduced by the operating environment: the dominant atmospheric and thermal noise on Mars is very different to the ocean-generated noise seismologists are experienced at dealing with on Earth. Additionally, the martian atmosphere is much more rarefied than ours, and as a result has a much lower heat capacity. This is expected to create extremely strong wind and thermal expansion signals during the large-scale temperature changes around sunrise, which may overprint any seismic signals present.

The nature of seismic excitation and attenuation may also be quite different, and hence analytical techniques developed for use on Earth may require some adjustment. In the case of the former, the thin atmosphere provides far less heating to incoming projectiles, enabling more meteorites to reach the surface than on Earth. This means that impactor events will be likely to generate a larger proportion of the seismicity on Mars than on Earth. They will need to be analysed differently, with the important source quantities being impactor size, angle of entry and explosion altitude rather than depth and rupture magnitude.

Additionally, a significant portion of the seismic attenuation on Earth comes about from water trapped along grain boundaries in the rock. The expected differences in the water content of the crust and mantle on Mars may change the planet's attenuation profile, and this would need to be investigated and accounted for to ensure robust inversion of the source parameters.

The challenges inherent in processing InSight's data, and the new techniques that may need to be developed, have led the InSight team to take an unusual route to finding the best approach. They decided to run a "blind test", in which members of the community would develop methods for analysing synthetic data and to make inference about Mars's structure, in preparation for the return of genuine data next year (Clinton *et al.* 2017).

Fourteen published and potential one-dimensional models of Mars's interior structure were chosen, and participants received travel-time curves for the different seismic phases for each possible model. From these models, one was chosen and used to generate an Earth year's worth of broadband synthetic martian seismic data, but the participants were not told which model had been used.

The aim was for each participating team to identify as many of these events within

Understanding the Earth

Over time, analysis of waves from many different seismic events leads to the development of a model of the Earth's interior. These models include information about both the planet's gross structure, such as the presence of a liquid outer core, and about local-scale heterogeneities such as magma plumes or sedimentary basins.

Although the signals from different events may look dissimilar at first sight, their fundamental structure is normally similar. A three-component seismometer – that is, one capable of recording vertical ground motion as well as in the horizontal plane – will first detect the arrival of body waves, first P waves (pressure) and later S waves (shear), followed by surface waves (Love waves, then Rayleigh waves). These different waves are referred to as seismic phases. These waves can be identified both by the relative timing of their arrivals and the components of the seismometer in which their signatures present

themselves.

In the source-centred "rtz" coordinate system commonly used in seismic analysis, the radial (*r*) direction is on a line pointing from the source to the receiver, the vertical (*z*) points outwards through the surface in the source-receiver plane, while the transverse component (*t*) points out of the source-receiver plane to complete the triad. P and Rayleigh waves oscillate almost exclusively in the *r*–*z* plane, while S and Love waves are most prominent in the transverse component.

This picture is complicated by the reflection and diffraction of seismic phases at interfaces between rocks of different properties. For example, PP is a P wave that has bounced once off the planet's surface before reaching the receiver, while PKP is a P wave that has travelled through the outer core. Additional modifications to a signal can be introduced by seismic anisotropy, whereby rocks respond differently to waves incident from different

directions – this can cause a fraction of the energy to leak from one component to another. Nonetheless, by measuring the components of oscillation, relative arrival times, amplitudes and phases of the different types of seismic wave it is possible to differentiate between them.

Further investigation can be conducted using synthetic modelling, whereby simulations are run for various different source and structural parameters, based on a reference model. A fitting can then be performed to find the closest match to the data. However, because the parameter space to be searched is so enormous and a unique inversion cannot necessarily be defined, this can be a computational challenge and on Earth is normally only done with some *a priori* knowledge of the likely input parameters. This is feasible on Earth because of the extensive catalogue of seismic signals and the well-developed models of the planet's interior.

the data as they could and, where possible, locate them on the planet and determine source parameters. The end result was expected to be correct identification of which of the 14 models was used in the data generation. Of course, when it comes to the analysis of the actual InSight data, the challenge will be that we will not have *a priori* theoretical travel time and dispersion curves for all possible models, rather these will have to be developed empirically.

Participants were told that the year's worth of data contained roughly 250 events with a variety of different source parameters, from locations with a random distribution. Magnitudes and event depths followed a skewed Gaussian profile, with peaks in magnitude at around 2.5 and at approximately 35 km depth. The largest events were stated to be no larger than $M_w 5$ and the deepest no more than 100 km.

A model of the dynamic response to the wind shaking and thermal heating of the probe was used to create synthetic atmospheric noise, which was superimposed on

the seismic data to simulate realistic data returns. The three-component seismometer data and wind pressure data were provided at a frequency of 2 Hz, the wind speed, direction and temperature at 0.1 Hz and the magnetic field at 0.2 Hz.


Ten teams eventually made submissions, using various different approaches. A full summary of the results will be published in a paper later this year, but here we present a short discussion of the

approach taken by one of the teams, consisting of PhD students and postdoctoral researchers from the Seismology Group in the Department of Earth Sciences at the University of Oxford.


Processing

Our team chose to analyse the data using a mixture of custom-written algorithms and those from the Python ObsPy seismic package. Here, all references to specific days of data are given in martian solar days ("sols", each equivalent to 24 hours and 40 minutes).


1 Waveforms raw (top) and processed (bottom) for the U component of the spacecraft seismometer for sols 052 and 053, in which no event is immediately obvious. The removal of the strong diurnal trend after processing is clear, though several hours of intense wind noise remain each day.



2 Spectra for an event on sol 195, the quiescent section noise on sol 11 and the high-amplitude wind noise which lasts a few hours, from sol 1. The vertical dashed lines indicate the passband used in filtering, which provided the best improvement in signal-to-noise ratio.



3 Detail of the large event on sol 257. The “start” of the event is picked manually to be the first point at which the wave amplitude rises significantly above the background. The P waves are clear approximately 50s in, followed by the S waves at approximately 250s. The largest amplitudes are formed from the Love waves (first) and Rayleigh waves (second, with long coda).



One component of the raw, unprocessed velocity data for two sols (approximately 180 000s in total) are shown in figure 1 (upper panel). The coordinate system is centred on the spacecraft and so not geographically meaningful, but several phenomena can be identified immediately. The first is the clear diurnal trend caused by thermally induced expansion and contraction of the ground, with the observed background amplitude rising and falling over the course of a sol.

There is an additional noise signal which is significant for several hours of each sol, presenting as sharp peaks which can be of similar amplitudes to the background. Analysis of the accompanying atmospheric data indicates that these peaks

correlate with variations in air pressure and wind speed, and hence are attributable to dynamic shaking. The wind noise throughout the rest of the day is over an order of magnitude lower in amplitude. We differentiate between the “wind noise” that overprints any clear seismic signal for part of each sol, and the “quiet noise” that is present throughout but does not necessarily obscure seismic signals.

Our first step was to detrend the signal and remove the instrument response, which was done using ObsPy; we successfully removed most of the diurnal variation from the signal. We also compute on-the-fly integration of the signal, such that we move from a velocity representation of the ground motion to a displacement one (figure 1).

Rather than making assumptions about the expected spectral characteristics of a marsquake, we first looked through the data to find a large, unambiguous event. One such example was found on sol 195. Assuming that the spectra of marsquakes are more similar to each other than to the noise spectra, this gave us a baseline by which we could choose which spectral bands to perform our searches in.


A comparison between the spectrum of the quake and of two arbitrarily chosen noise sections in which no event is obviously present is shown in figure 2. The noise is much larger in amplitude at frequencies up to 0.08 Hz, and again at frequencies close to 1 Hz (although this high-frequency noise band may come from aliasing arising from the removal of the instrument response). Armed with this knowledge, we tested different filters to try to optimize the signal-to-noise ratio of the observed events. A four-corner bandpass filter with a range of 0.08–0.5 Hz was found to be best and we used this filter for all further data analysis. The resulting filtered seismogram for another large event on sol 257 is shown in figure 3.

Detecting an event


Terrestrial seismologists normally use automated algorithms to detect events; we thought it prudent to try applying some of these to the InSight data because the potential search efficiency is much higher. We note, however, that with data that is potentially very different from terrestrial spectra, visual inspection and investigation are also valuable. Terrestrial seismology commonly uses short-term-averaging/long-term-averaging (STA/LTA), an algorithm that is triggered when the local average value of a signal exceeds the longer-term background. We used this as a first step towards automating the detection procedure.

However, we found that this procedure yielded too many false positives, because it was being triggered by the pulses in the seismogram caused by gusts of wind. To correct for this, we experimented with using correlation between the seismometer signal and the wind pressure data to exclude wind-induced peaks. This showed promise, but because of the limited time available for the blind test we were unable to apply it over the entire dataset. As a result, our events were identified using both visual inspection and STA/LTA, but were only selected for further analysis if the presented waveform looked clearly seismic rather than atmospheric in origin.

In order to make a realistic estimation of magnitude and event type, we needed to constrain the epicentral distance (distance from earthquake epicentre to receiver) and azimuth/bearing of the source. The



4 Integrated energy over P wave duration in radial and tangential components. Energy in tangential component is clearly minimized at $\sim 65^\circ$ and $\sim 245^\circ$, defining the bearing from grid north to an axis along which the source lies.



5 The same waveform, but in source-centred coordinates. Compared to figure 3, note how the phases now present in different components: P and Love waves in r and z , s and Love waves mostly in t .

first step was to transform the data from a spacecraft-centred coordinate system into one useful for seismic analysis.

The first coordinate transform applied was thus to rotate the data into an xyz coordinate system, with x as the east–west (east positive) axis, y is the north–south axis and z radially outwards from the surface. This can be done by applying a simple rotation matrix to the three-component data, where the elements of the matrix take account of the spacecraft’s bearing from north and inclination to the horizontal. Once this rotation is complete, different parts of the event waveform are clearly identifiable, as in the data for the sol 257 event shown in figure 3.

The second coordinate transform is to move from the geographical system to the source-centred (rtz) system used in terrestrial seismology and described in the box “Understanding the Earth”. This requires finding the azimuth of the source relative to the receiver, and we illustrate two separate methods for doing this.

First, we expect the P waves to oscillate almost entirely within the radial–vertical plane, with little to no transverse component of motion. Thus, we can isolate the P waves using a short time window and perform a grid search across all possible azimuth angles to find the axis that minimizes the energy of the wave contained within the transverse component, which is equivalent to maximizing the energy contained in the radial–vertical plane. Results for this method are shown in figure 4 and constrain the event to lie somewhere on a great circle passing through the spacecraft. The seismogram is then rotated such that the radial direction coincides with the source–receiver line, with the results shown in figure 5.

An alternative method is to plot the

three-dimensional point cloud representing the particle motion induced by the P waves during their entire duration. By finding the eigenvectors of the covariance matrix of the point cloud, we are able to identify the dominant axes of oscillation. We found that the two methods are in good agreement for our test cases.

Note that neither of these methods in isolation constrain the direction along the great circle that the seismic waves arrived from – to do this, we must also calculate the angle of incidence from the particle motion, which lifts the degeneracy in arrival angle.

Epicentral distance

The next step is to calculate the source–receiver distance. The first method we use is to measure the travel-time difference between the first arrivals of the different body waves. Because they propagate at different speeds, the greater the difference in arrival time between P and S waves, the further that the waves have propagated. However, unlike the determination of source azimuth, this parameter is model-dependent: the ratio between P wave and S wave speed is not constant across the 14 different models. As a result, we first need to calculate the expected travel-time differentials for a number of different source depths and distances for each of the 14 models, using the TauP software. We then overlay our measured travel-time difference (20° uncertainty in azimuth and up to 20 s uncertainty in arrival time to account for error in isolating the first arrivals from the background), to see which region of the depth–distance parameter space provides a match.


As expected, most events in isolation yield a range of possible depth–distance combinations, but not necessarily a unique

solution, indicating that the particular model being tested is a potential candidate for that chosen to generate the data. The greater number of individual events that a particular model can fit, the more likely it is to have been the one used. An example of a good fit is shown in figure 6, for the EH45TCold structural model.

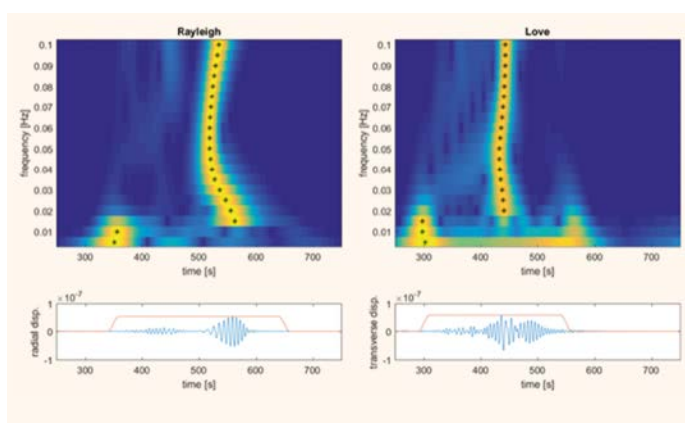
Some of the 14 models, however, do not offer a solution, meaning that no combination of possible source depth and distance yielded the expected travel time. This can occur either because the source depth would have had to be greater than the maximum (100 km in this test), or because the waves would have had to come from within a “shadow zone”, a range of angles from which diffraction and refraction return no incident waves at the receiver location.

Of the 14 models, we found a variant of the EH45TCold model to be the best, fitting 8 of the 26 events located with an acceptable uncertainty range. The other 17 events were also potentially locatable using this model, but the low signal-to-noise ratio made determining the exact first arrival time of the P waves difficult and hence an accurate measurement of travel time differential could not be made. For three of the events, surface waves were well enough resolved to enable verification of the derived epicentral distances. In addition to measuring the travel-time differences between the Love and Rayleigh waves and comparing these to the travel times provided, we could also make intra-waveform measurements. This is possible because the surface waves all show weak dispersion, such that the higher frequency portions of the waveforms arrive earlier than the lower frequency ones. By comparing the observed degree of dispersion (see figure 7 for an example) to that

6 Possible travel times for different depths and distances in the EH45T-Cold model. Gradation runs from red (longer travel times, more distant events) to blue (shorter, closer). The red bar shows the range of possible origins for the sol 257 event based on its measured travel-time difference. The white wedge is a shadow zone where no waves are received. Note the greater sensitivity of this method to distance than to depth – lines of constant travel time are nearly parallel to the depth axis.



7 Dispersion spectra for the Rayleigh and Love waves (upper panels), showing the difference in arrival times of the different frequencies. Data below 0.02 Hz are excluded from the analysis as they are not sufficiently resolved. The windows (red) and waveforms (blue) used in each plot are shown in the lower panels.



expected for the different models, we can make additional estimates of the propagation distance. As with the body-wave case, such measurements are model-dependent.

The combination of body wave and surface wave estimations is a useful one because they have different sensitivities to structures and sound speed profiles at different depths – the latter are most tuned to those in the crust as they are constrained to propagate there, while the former also display sensitivity at greater depths in the mantle.

Blind test results

In total, 84 teams from 20 countries registered for the test, with 10 groups from 5 countries eventually making catalogue submissions; two of these teams came from the Mars Quake Service teams themselves. A project involving a number of French high schools (coordinated by the Geoazur Laboratory in France) was the most successful at event location, detecting some 102 of the roughly 250 events in the catalogue.

The Mars Quake Service was the most successful at identifying and geographically locating events, submitting some 63 of which all were located. Of the eight quake hypocenters the Oxford team located, two

were within the L1-norm requirements for computed distance (10% error) and azimuth (20°). As expected, all of the submitting teams had developed procedures that were best at detecting the closer or larger events.

There are several other techniques and procedures that may be implemented to glean further information about the events from their seismic signals, which we considered but were unable to implement in the available time. Many of these are likely to be usable and implemented by those analysing InSight's data. The use of modelling synthetics could be significantly expanded, in order to search a larger region of parameter space. Additionally, inversions (to derive the model from the observations) could be made more comprehensive by using more sophisticated multiparameter fitting techniques.

In terms of the specifics for processing the signals from possible events, it may be possible to constrain the event azimuth using the surface wave arrivals. This idea is not well explored in terrestrial seismology because it is normally not needed; it may require some development to be useful.

We were aware that we discarded some information by examining only the “minor arc” wave phases, that is, the set of waves

that reach the probe by travelling the shorter of the two distances from the source to receiver around the great circle. There are clear hints of surface waves travelling along the major arc as well, which would enable us to make estimates of how the different paths change the observed signatures of the quakes, and hence to potentially infer the presence of localized subsurface heterogeneities. Similarly, if phases which are also sensitive to the structure of the core can be detected, or indeed if we could find the planet's normal modes, this would give us valuable information about the deep interior.

We anticipate that constraining event magnitude and source type (impactor or quake) will be much harder than on Earth, but again synthetic modelling should be of significant help. Additionally, several other authors have proposed empirical relations that may be tested and then used to make such estimations.

After the test, the organizers revealed that the model used in the data generation was in fact EH45TColdCrust2. This was a member of the model family that our team concluded was most likely to be correct, with a cold crust and liquid inner core (Sanloup *et al.* 1999).

As a result of the blind test, the Mars Quake Service has concluded that there are several methods developed by teams during the blind test that should be included in their analyses, including use of the surface wave phase information and of body wave polarization. They also identify discrimination between impacts and quakes to be an area where further work and modelling is needed, because only one team successfully identified an impact.

InSight is due to land on Mars on 26 November 2018. The first results from the seismometer will be expected sometime next year, with the results of the blind test published beforehand. After that, we may be able to make the first data-led deductions about the interior of another planet. ●

AUTHOR

Benjamin Fernando and Maria Tsekhmistrenko are DPhil students at the University of Oxford, UK. Kasra Hosseini is a postdoctoral research assistant at the University of Oxford.

ACKNOWLEDGMENTS

Figures and data analysis by Benjamin Fernando, Kasra Hosseini, Maria Tsekhmistrenko, Harriet Godwin, Claudia Haindl, Thomas Garth, Alexandre Szenicer and Martin van Driel. With thanks to the InSight Blind Test Team.

MORE INFORMATION

InSight Mission <http://blindtest.mars.ethz.ch/#download>
Blind Test <http://blindtest.mars.ethz.ch/#download>

REFERENCES

- Anderson DL *et al.* 1977 *J. Geophys. Res.* **82**(28) 4524
- Clinton JF *et al.* 2017 *Seismol. Res. Letts.* **88**(5) 1290
- Sanloup C *et al.* 1999 *Phys. Earth Planet. Interiors* **112**(1-2) 43