# THE OBSERVATORY

Vol. 146 2026 AUGUST No. 1313

## MEETING OF THE ROYAL ASTRONOMICAL SOCIETY

Friday 2026 January 9 at $16^{h}\ 00^{m}$
in the Geological Society Lecture Theatre, Burlington House

Mike Lockwood, *President*
in the Chair

*The President.* This is a hybrid meeting and those on Zoom should see a small green shield which means you are using the most up-to-date version of Zoom. Please put questions in the Q&A and they will be read out by Dr. Pam Rowell at the end of each talk.

I'm going to show one of the latest Christmas cards in history. It is a diagram produced by Annie and Edward Maunder which first appeared in *MN* in 1904. It's commonly called a Butterfly diagram and it shows solar latitude *versus* time, and you can see sunspots moving to lower latitude as each solar cycle progresses. If you extend it a few years and turn it through 90 degrees you get a Christmas tree. [Laughter.] That is to thank you for all the support that you give to the RAS. Although it's too late to wish you a happy Christmas it's not too late to hope that you have the very best of what is left in 2026.

Now for the best job that the President gets to do and that is to announce the annual awards. I'd like to thank the members of the Education Outreach panel, the Agnes Mary Clerke Medal panel, the Astronomy panel and the Geophysics panel, and also Nush Cole who managed much of the administrative support for us. Thanks also to those who put in nominations. The James Dungey Lecture will be given by Professor Andy Hillier (Exeter), the Harold Jeffreys Lecture by Professor Andy Biggin (Liverpool), the George Darwin Lecture by Professor Suzanne Aigrain (Oxford), and the Gerald Whitrow Lecture by Professor Richard Battye (Manchester). Early Career Awards go to Dr. Ke Zhu (Bristol, Geophysics) and Dr. Deaglan Bartlett (Oxford, Astronomy). The Fowler Award for Geophysics goes to Professor Paula Koelemeijer (Oxford) and that for Astronomy is awarded to Dr. Ingrid Pelisoli (Warwick). The Service Award (G) goes to Professor Mark Burchell (Kent). The Ian Robson Lifetime Achievement Award goes to Professor Raman Prinja (UCL). The Higher Education Award goes to Dr. Kurt Froebrich (Kent), whilst the Secondary Education Award goes to Mrs. Vicky Dean (The King's Academy, Middlesborough). The Agnes Mary Clerke Medal is awarded to Dr. Edward Gomez (Cardiff), the Jackson–Gwilt Medal to Professor Alistair Glasse (Edinburgh), the Chapman Medal is awarded to Professor Matthew Owen (Reading), and the Eddington Medal goes to Professor Debora Šijački (Cambridge). The Herschel Medal is awarded to Professor Andrew Bunker (Oxford). The Gold Medal for Geophysics is awarded to Professor Andrew Jackson (ETH, Zurich) and the Gold Medal for Astronomy goes to Professor Shrinivas Kulkarni (California Institute of Technology).

Our first speaker is Dr. Stephen Thorp. He is a postdoctoral researcher at the Institute of Astronomy and Kavli Institute for Cosmology, Cambridge, where he works with Professor

Hiranya Peiris. He spent the first three years of his postdoctoral position in the Oskar Klein Centre at Stockholm University, and relocated to Cambridge in 2025 September. From 2018 to 2022 he was a PhD student at the Institute of Astronomy, Cambridge, working with Professor Kaisey Mandel. Before that he studied physics at the University of Birmingham. Stephen has broad interests in statistics and machine learning applied to astronomy problems, particularly in the fields of galaxy evolution (which he'll be talking about today), strongly lensed transients, and supernova cosmology. The title of his talk is 'POP-COSMOS: Machine-learning the galaxy population for large cosmological surveys'.

*Dr. Stephen Thorp*. Thank you for having us here. We have had a really nice Specialist Discussion Meeting on Machine Learning Applied to Astronomy and I will be presenting a broader view of some work that I presented earlier.

This presentation introduces the 'POP-COSMOS' framework, a generative model for the galaxy population probed by large cosmological surveys. We will look at what it means to build a generative model, how we can build such a model for the galaxy population, the challenges involved, and how we can overcome those challenges. We will conclude by applying this framework to data from the Cosmic Evolution Survey (COSMOS).

Our motivation is cosmological surveys — particularly photometric ones — that seek to map the matter distribution and large-scale structure (LSS) in the Universe by observing billions of galaxies. The matter-density field in the Universe — and the underlying cosmology — is traced by the clustering of galaxies in 3D space, and/or distortion of their observed shapes due to weak gravitational lensing. Those tracers are the primary way today's large surveys probe the nature of dark energy.

Two of the largest photometric surveys ever undertaken are now in progress: the *Vera C. Rubin Observatory* Legacy Survey of Space and Time (LSST), and the European Space Agency's *Euclid* mission. The cosmological data sets produced will consist of four-band (*Euclid*) or six-band (LSST) photometry, and the measured shapes and positions of billions of galaxies over thousands of square degrees. We will not have 3D positions for all galaxies — only spectroscopy can provide that by measuring redshift precisely — so our cosmological analyses will be done *via* tomography. That means sorting galaxies into redshift bins based on their photometry; that is sufficient for cosmology, but relies on detailed modelling of the sorting process. We need to understand the true redshift distributions of galaxies sorted into each of our tomographic bins.

With that context in mind, we will switch gears to talk about generative modelling. The goal of any generative model is to design an end-to-end simulator that is able to replicate a target data set. We will work through what that means conceptually, how we can build a model for the galaxy population, and why that will help us with tomographic binning. The accompanying Fig. 1 shows three sketches of the model building process.

The first panel **(a)** shows the general model structure we will follow. At the deepest layer of the model, we have a population distribution. The intrinsic (physical) properties of any finite sample of objects will be a random draw from the underlying population distribution. The intrinsic properties will often not be directly observable; a physical model or process will be needed to translate those into observable features. Finally, to generate realistic simulated observations, we will require a model for the noise and selection effects in our target data set.

To 'train' our generative model, we will need to specify some measure of similarity between our simulated observations and the target data set. We will then simultaneously optimize any degrees of freedom or tuneable parameters in the layers of our model, until our simulated observations match the data. The largest degree of freedom will be the underlying population distribution; if we can learn that, we will have learned something about the true population underlying our data.

To adapt that general model structure to the galaxy population, there will be several challenges:

1. What set of physical properties is sufficient to parametrize a galaxy, and how do we map those to observables (photometric fluxes)?

2. Will our model be computationally scaleable for billion-galaxy data sets?

3. What is a sufficiently flexible form for the population distribution?

4. Can we realistically replicate any relevant observational effects?

To tackle the first challenge, we will use stellar-population synthesis (SPS). Pioneered by Beatrice Tinsley, that is a detailed framework for synthesizing galaxy spectral-energy distributions (SEDs). SPS models assemble SEDs from component stellar populations, incorporating effects such as dust, chemical evolution, nebular emission, and active galactic

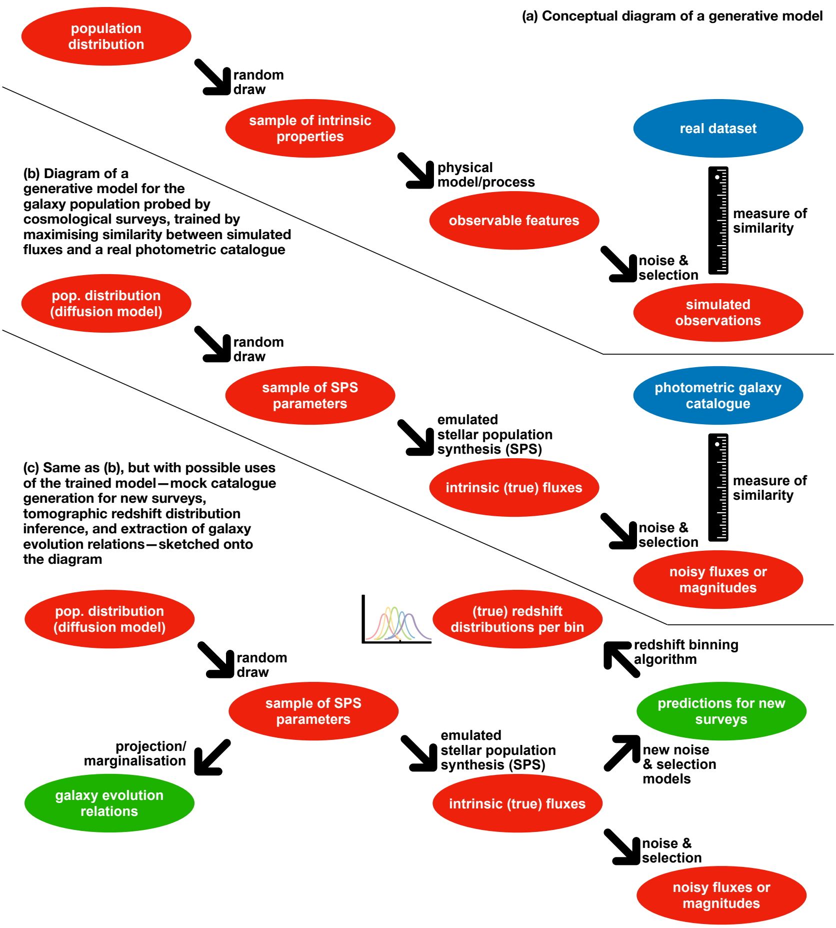


FIG. 1

Diagrams sketching a generative model for the galaxy population. The top panel **(a)** shows a problem-agnostic conceptual diagram. The second panel **(b)** shows our model design for photometric galaxy survey data. The final panel **(c)** shows possible uses.

nuclei. Modern SPS models represent a diversity of galaxy SEDs with a handful of parameters such as stellar mass, metallicity, star-formation rate (SFR), redshift, and dust optical depth. We will use a 16-parameter SPS model.

We now face a computational challenge. Detailed SPS models are inherently complicated, including many components, slow underlying calculations, and time-consuming I/O operations (*e.g.*, loading spectral templates). Additionally, many 'industry standard' SPS codes are not optimally suited for GPU-based high-performance computing (HPC) systems. To scale SPS to billion-galaxy data sets, we can use 'neural emulation'. That involves training an artificial neural network to learn the mapping between parameters and SEDs (or photometric fluxes) that is provided by SPS. A neural emulator will be able to replicate the output of an SPS code to very high accuracy (fluxes approximated within 1%), with a fraction of the computational cost (≈10 000× faster). Our emulator code, 'SPECULATOR', is publicly available on-line.*

Our third challenge is to find a sufficiently flexible way of representing the underlying population distribution over our 16 SPS parameters. Two possible routes are: (*i*) to take a classical hierarchical Bayesian approach and make assumptions to write down an explicit form; or (*ii*) to use a flexible density estimator as in generative artificial intelligence (AI). We will take the second route here, using a class of models called 'diffusion models'. Those represent a complex 'target' distribution as a continuous deformation of a simple 'base' distribution (often an isotropic Gaussian). Rather than explicitly learning a functional form for the target distribution, one instead learns the transformation. Those kinds of models have been workhorses for tasks such as AI image generation. We have a lightweight code, 'FLOWFUSION', available for building diffusion models.†

Our final challenge is applying realistic observational effects: transforming our intrinsic model fluxes or SEDs into simulated observations by adding appropriate noise and selection effects. Adding realistic noise can be framed as another density-estimation task — learning the distribution of flux errors conditional on true fluxes — which can be solved using a diffusion model (or similar). Sample selection criteria (*e.g.*, magnitude limits) will usually be known from the survey design, but can in principle be represented probabilistically and learned during training.

The second panel **(b)** of Fig. 1 shows our full generative model design. That is the recipe we implement in the 'POP-COSMOS' framework.‡ We train the model — optimizing the diffusion model that represents the underlying population distribution — by maximizing the similarity between simulated noisy fluxes and a deep photometric galaxy catalogue.

The final panel **(c)** of Fig. 1 shows two possible uses of our trained model. For example, we can simulate realistic photometric galaxy catalogues for wide-area surveys (*e.g.*, LSST or *Euclid*), and can subsequently apply a tomographic redshift-binning algorithm to our simulated photometry. That lets us infer the underlying redshift distribution in each tomographic bin, necessary for accurate cosmology. Estimating those distributions *via* generative modelling sidesteps challenges associated with stacking per-galaxy redshift estimates, or using unrepresentative spectroscopic calibration samples. The second use-case we highlight is the inference of galaxy-evolution relationships. Since the learned population distribution will encompass the evolution of galaxy properties across cosmic history, we can 'slice' it in different ways — by projecting or marginalizing over parameters — to extract key relationships.

We have trained the POP-COSMOS model on data from the ≈2 deg$^2$ COSMOS field; specifically the *COSMOS2020* catalogue's 26-band ultraviolet–infrared photometry. We have in our training set ≈420 000 *COSMOS2020* galaxies brighter than ≈26 mag in *Spitzer IRAC*

*https://github.com/Cosmo-Pop/speculator

†https://github.com/Cosmo-Pop/flowfusion

‡https://github.com/Cosmo-Pop/pop-cosmos

[*Infrared Array Camera*] Channel 1 (≈3·6 μm). We validate the trained model by comparing simulated colours and magnitudes to the real catalogue, and by comparing the learned galaxy evolution relations to literature results from other surveys or tracers. Some key results we extract include:

- colour–redshift relations;
- the redshift distribution of galaxies with $r < 25$ mag;
- the SFR *vs.* stellar-mass relation (star-forming sequence);
- the cosmic SFR density; and
- the stellar-mass and redshift-dependent fraction of quiescent galaxies.

To conclude, we have designed a recipe for interpretable generative models, combining machine-learning tools with a physically motivated parameterization. We have applied that approach to building a generative model for the galaxy population probed by wide and deep surveys, and have examined the galaxy evolution results one can extract from the model (detailed fully by Thorp *et al.* in *ApJ*, **993**, 240).

*Professor Ofer Lahav*. Experience from surveys shows that it really depends on what you do with it rather than the application. Specifically, some groups just use point estimation of the redshifts and others use the whole PDF or distribution function. I just wonder about the application of your method.

*Dr. Thorp*. This is attempting to be a third way of getting redshift distributions, rather than stacking or averaging point estimates or per-galaxy posteriors — which we know isn't strictly speaking correct in a Bayesian sense. We're trying to get the redshift distributions in tomographic bins in a third way by building a forward simulator and applying our real-observations noise model and observational selection effects along with whatever point estimator is being used to do the target on the real galaxies. When we apply that to the model galaxies we obviously know their true underlying $z$ so we can see what was the redshift distribution of the galaxies that made it into our 'top-hat' bins as it were.

*The President.* You glossed over random selection of the stuff when you were showing that iteration in the training; do you do the selection every time or is it for that random selection?

*Dr. Thorp*. When you are matching the real data you just want to draw enough model galaxies from the underlying distribution that it's capturing the whole structure.

*The President.* Are you using the same selection through the whole of that process?

*Dr. Thorp*. We are using the same magnitude limit because we are trying to model a survey with one magnitude limit although, in principle, if you have reason to believe that there is some varying depth across the sky you could fold that in.

*The President.* For the computer time that you are using you probably don't want to be doing a sensitivity study to check that your selection is random enough and large enough.

*Dr. Thorp*. Most of that is done with stochastic optimization. There is a degree of randomness anyway. Every time you adjust something in your population model you take a fresh draw and push it through the whole thing and there is a stochasticity there anyway.

*The President.* Any other questions? Thank you very much, Stephen. [Applause.]

The next speaker, Dr. Francesco Scotto di Uccio, is currently a postdoctoral researcher at the Department of Physics of the University of Naples Federico II, where he obtained his Master's and Bachelor's degrees in Physics, followed by a PhD in Structural, Geotechnical, and Seismic-Risk Engineering (awarded on 2025 February 21), with a thesis entitled 'Detection and characterization of microseismicity using advanced techniques' (supervisors: Prof. Gaetano Festa and Prof. Matteo Picozzi). During his PhD programme, he carried out research periods as a Visiting Student at Stanford University and at the GFZ German Research Centre for Geosciences in Potsdam, where he had the opportunity to develop and apply advanced techniques for microseismicity detection, also using fibre-optic sensors for seismic monitoring. The title of his talk is 'From background seismicity to seismic crises: What can we learn from machine-learning-enhanced catalogues?'

*Dr. Francesco Scotto di Uccio.* Microseismicity continuously occurs along the seismogenic structures that can also host larger, destructive earthquakes. Therefore, its characterization can potentially provide crucial information on the geometry and mechanical state of the underlying faults, before the occurrence of notable events. However, conventional catalogues are limited in size, since many small events are hidden in the noise. Therefore, discovering such events calls for advances in both earthquake-detection techniques and monitoring infrastructures.

We showcase how the integration of machine learning and similarity-based detection techniques can increase the content of seismic catalogues both for background seismicity and seismic sequences, up to one order of magnitude as compared to conventional manual catalogues in Southern Italy. In that framework, machine-learning models provide an enhanced set of events as compared to the standard catalogue, which can be used as templates to detect effectively lower-magnitude, colocated events, also with a short interevent time. Following that strategy, seismic catalogues can also be enriched by nearly one order of magnitude. The location of the earthquakes in the enhanced catalogue revealed the activated fault patches, while the determination of source properties enabled the definition of evolutive models for the seismic sequences. In the case of seismic sequences, relocated seismicity clearly identifies local patches on kilometre-scale structures, featuring an orientation coherent with the main faults of the area. To monitor background seismicity, we integrated the standard seismic network with 200 stations deployed in dense arrays for one year, demonstrating the possibility of consistently detecting small-magnitude earthquakes with the use of short-term dense arrays and established machine-learning models. Moreover, that configuration enables the downscaling of the seismicity characteristics to small, decametre-size events, achieving resolution as from multiple years of conventional monitoring. We estimated focal mechanisms using a CNN [convolutional neural network] for evaluating P-wave polarities and we inverted the data set of fault planes for determining the characteristic of the stress field. The orientation of the stress field resulting from the inversion of the fault planes agrees with previous estimates, based on six years of microseismic events. The stress field obtained has a vertical most compressive stress $\sigma_1$ and horizontal intermediate and least compressive stresses, $\sigma_2$ and $\sigma_3$, respectively, the latter azimuth lying in the direction of the anti-Apenninic extensional stress, featuring coherent values of the relative stress magnitude as compared to former studies. From a spatio-temporal analysis, we observed a spatial homogeneity of stress field in the area on ≈20-km scale, while the time distribution of the stress ratio reports lower values observed up to the end of January and higher values up to June, modulated by the occurrence of the oblique faulting. That result correlates with seasonal variations in the regional strain.

We finally focussed on the characterization of the 2025 Santorini seismic sequence, revealing the intense level of seismic activity during the crisis and demonstrating how deep-learning approaches can support near-real-time tracking of the evolution and dynamics of the sequence. Advanced detection techniques can increase standard catalogues by ≈12× (≈48k events), identifying five stages of the crisis according to hourly rate (≈70 earthquakes per hour) on February 6 when an emergency was declared. Deep catalogues enable the investigation of moment tensors and tomographic models, potentially identifying magmatic sources.

*Dr. Pam Rowell.* This question has come in from Kathy Whaler. "Do you think there is scope to use machine learning in attempts to infer temporal changes in tomographic models, *e.g.*, associated with magma movement?"

*Dr. Scotto di Uccio.* That tomographic model was not one of my main activities in this work but actually we are still not so much involved in the tomographic models in that way, but given the huge seismic activity that we were able to recognize by machine learning we can use more traditional approaches to try and get some variation in the velocity in the vents in the area.

*Dr. Quentin Stanley*. That was absolutely fascinating. It puts me in mind of a discussion that we had some years ago in a Discussion Meeting looking at the difficulty of understanding earthquakes. Two comments: are you able to use your machine learning to put it on edge compute so that your seismic stations are able to say what is going on rather than post analysis? Would that be a possibility in the future?

*Dr. Scotto di Uccio*. Simple techniques as algorithms that identify earthquakes based on the energy evaluation of the seismic signals are running, in real-time, on dedicated servers that receive the collected data, with a latency lower than one second. In their standard configuration, machine-learning models for earthquake detection are based on more extended windows (*e.g.*, 30 s), therefore are not fully applicable to real-time monitoring. In the future, we can eventually retrain the model with shorter windows.

*Dr. Stanley*. That would be a tremendous advantage going forward and the other little bit is because you have that amount of data coming through are you able to analyze it or is it coming through too fast?

*Dr. Scotto di Uccio*. To get information we integrated results from seismic modelling but also from GNS stations and, of course, seismic information can provide the tools for understanding the loading.

*Professor Mark Lester*. Thank you, Francesco. Could I just ask: in the Santorini data towards the end of that period in the final phase, the distribution of the magnitudes seems to start to become quite banded so you have values at the highest level and then at two other levels. Is there a physical reason for that compared with the distribution that you see earlier?

*Dr. Scotto di Uccio*. Actually it does not report the magnitude but that is the distance along a profile, *e.g.*, the lower bands are closer to Santorini's area and the higher bands closer to the Amalfi coast. That has a physical meaning because, typically, after an earthquake we expect lower-magnitude events to be more frequent than larger-magnitude events associated with the same earthquake. Here everything was quite complex because we had large earthquakes, hundreds of events. It was difficult to disentangle which was the decay rate for which earthquake in the case of Santorini; it is very overlapping. It means that the seismicity is occurring in some particular spots on the basin.

*The President.* I'd like to ask you about the situation after the Florence earthquake but we should carry on with the next speaker. Zahra Zali is a postdoctoral researcher in seismology at GFZ Helmholtz Centre for Geosciences. Her research focusses on the analysis of continuous seismic and strain data to study transient processes in faults and volcanoes, with an emphasis on developing machine-based methods. She has worked on volcanic eruptions, slow deformation, and earthquake-related signals across a range of tectonic settings. Her current research aims to detect and characterize the preparatory processes preceding earthquakes and volcanic eruptions.

*Dr. Zahra Zali*. The Earth is constantly active, even when nothing dramatic appears to be happening. Subtle ground motions, slow deformation, and weak vibrations carry information about how volcanoes and faults evolve over time. Modern geophysical networks are designed to monitor that activity continuously, recording the planet's behaviour day and night with increasing sensitivity and coverage. Seismometers, strainmeters, and other geophysical sensors therefore produce vast streams of data. Yet much of our traditional analysis still treats the Earth as if it only speaks in short, discrete events: earthquakes, explosions, eruptions. Anything slower, weaker, or more ambiguous is often filtered out, averaged away, or simply ignored.

In this talk, I explore how machine learning, particularly unsupervised methods, can help us listen differently. Instead of asking only *when* an event happened, we can ask *how* signals evolve through time, and whether systematic patterns emerge before major changes in the system.

The challenge is not a lack of data. Modern seismic and geodetic networks generate terabytes of continuous recordings. The difficulty is that those signals are noisy, overlapping, and non-stationary. Important processes, such as magma migration or slow fault slip, may unfold gradually and leave only subtle fingerprints. Event-based approaches, by design, are not well suited to capturing such behaviour.

Machine learning offers a way forward, not as a replacement for physical understanding, but as a tool for organizing complexity. In particular, unsupervised learning allows us to explore continuous data without predefined labels or templates, something that is often unavailable in volcanic or tectonic settings.

A central idea in my work is to transform continuous signals into a form where patterns can be compared objectively. Seismic or strain data are first converted into frequency representations, using wavelet transforms or spectrograms. Those representations are high-dimensional and difficult to analyze directly. Autoencoders, a class of neural networks, address that problem by learning compact, nonlinear representations of the data. In effect, they learn how to describe complex signals using a small number of latent variables. Once the data are mapped into that latent space, clustering becomes meaningful. Similar signal patterns group together, and their temporal evolution can be tracked. That process is automated and does not rely on human-defined signal categories.

I first illustrated that approach using volcanic data from the 2021 Geldingadalir eruption in Iceland, a well-monitored event that produced a rich mixture of earthquakes, continuous tremor, episodic tremor, and changes in eruptive style. Fig. 2 shows the evolution of clustered seismic signals during that eruption. Each point represents a short window of continuous seismic data, mapped into a low-dimensional space using an autoencoder and then grouped according to signal similarity.

Several distinct regimes emerge naturally from the data. Before the eruption onset, the signal space is dominated by earthquake-related activity. After the eruption begins, a cluster associated with continuous tremor becomes prominent. That is followed by a separate cluster linked to episodic tremor during phases of lava fountaining, and later by another form of continuous tremor accompanying sustained lava outflow.

What is particularly striking is that a coherent tremor cluster appears several days before the eruption starts. Those signals are subtle and are often missed by classical analyses. By using a deep-based representation of the data, subtle changes and recurring patterns in the signals can be captured and tracked more systematically through time. The timing of that transition is consistent with independent geodetic evidence for shallow magma emplacement prior to eruption onset. Later in the eruption, the same clustering framework captures a distinct change in signal behaviour in late April, before the onset of lava fountaining on May 2. In other words, the machine-learning model identifies a weak but coherent pattern in the seismic data that precedes the observable transition in eruptive style.

The same philosophy can be applied to tectonic systems. I next showed results from the 2023 7·8-Mw Kahramanmaraş earthquake in Turkey, where months of elevated, tremor-like activity were observed before the main shock. Clustering reveals episodes dominated by low-frequency signals that increase about six months prior to failure. Those episodes consist of many small, repetitive pulses.

Source locations show a spatial association with industrial sites, indicating that that activity is related to anthropogenic processes. That observation is important from several perspectives. First, it demonstrates that such activities can generate tremor-like signals that are recorded by seismometers. Second, it raises questions about the potential influence of those processes on the local stress state of the region. Third, it serves as a reminder that patterns identified by machine-learning methods require careful physical analysis and independent investigation before being interpreted.

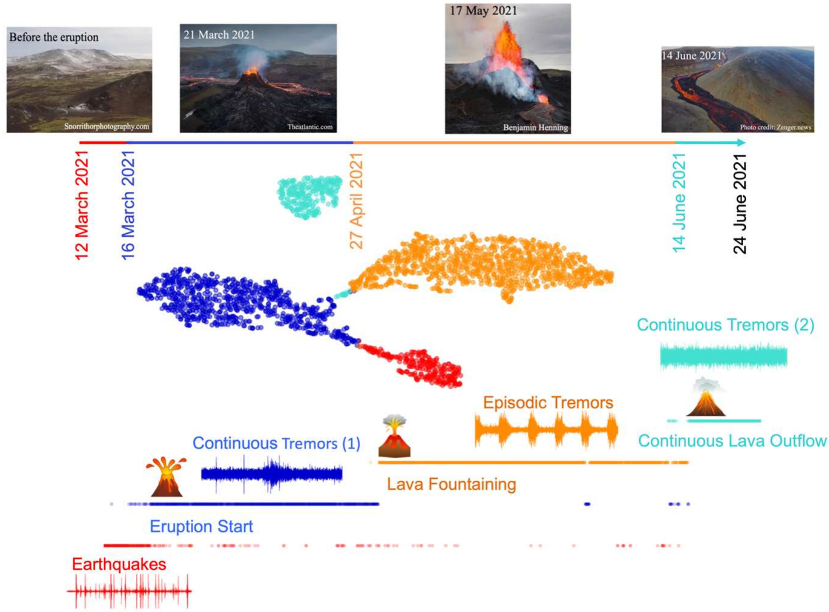


Fig. 2

Evolution of clustered continuous seismic signals during the 2021 Geldingadalir eruption, Iceland. The central panel shows a low-dimensional representation of continuous seismic data produced by an autoencoder, with colours indicating distinct clusters. Those clusters correspond to different eruptive regimes, including pre-eruptive seismicity, continuous tremor, episodic tremor during lava fountaining, and a later phase of sustained lava outflow. Photographs and dates above illustrate the close correspondence between cluster transitions and observed changes in eruptive behaviour.

The final part of the talk focusses on fault-slip processes at Parkfield on the San Andreas Fault, a long-studied transition zone between creeping and locked fault segments. Here, dense borehole strainmeter networks allow us to observe deformation with extremely high sensitivity. Using a deep-learning workflow adapted for strain data, we can automatically detect transient strain signals associated with short-duration slow-slip events.

Those slow-slip events are subtle, often lasting only minutes, and are easily masked by tides, atmospheric-pressure changes, and instrumental effects. After careful preprocessing, the machine-learning model identifies coherent strain transients that appear consistently across multiple stations. Their scaling behaviour extends known earthquake relationships into the aseismic domain, suggesting that slow-slip and earthquakes may be governed by related physical principles. Moreover, these slow-slip events modulate nearby low-frequency seismicity, possibly transferring stress along the fault. That highlights an important point: continuous deformation processes can actively shape the seismic cycle.

Across both volcanic and tectonic examples, a common picture emerges. Machine learning provides a unifying framework to detect, organize, and track transient processes in continuous Earth data. It excels at revealing structure where human intuition struggles, especially in high-dimensional, noisy signals. But interpretation remains essential. The clusters themselves are not explanations; they are tools for asking sharper physical questions.

In closing, listening to the Earth requires both sensitive instruments and flexible analysis methods. As our data sets continue to grow, approaches that embrace continuity rather than discretization will become increasingly important. Used carefully, machine learning allows us to hear parts of the Earth's story that were previously lost in the noise.

*The President.* Questions for Zahra?

*Dr. Rowell.* This question is also from Kathy Whaler who asks "You mentioned in the introduction also the large quantities of geodetic data that are now available. Are you currently analyzing them or have plans to analyze them with ML methods?"

*Dr. Zali.* I'm not currently working on that. I'm currently working on strain data. I might work in that direction to try and develop an algorithm for automatic analysis of GNS data as well.

*The President.* Any more questions? Quentin?

*Dr. Stanley.* With the clustering that you have done which has revealed a lot of that, how sensistive is this analysis to the clustering if you do it in a slightly different way?

*Dr. Zali.* The model is very sensitive in terms of small changes to the signal. We hope that what the model will do will be to find everything in the data. In terms of sensitivity to the data the model is seeing, it is very powerful. If you think about unseen data, that is a challenge and and we are trying to develop an algorithm for clustering done close to real time. We have not done that yet but we are hoping to.

*Dr. Stanley.* We look forward to the next exciting episode!

*The President.* Thank you very much, Zahra. [Applause.] As we have had a late speaker cancellation I have prepared a talk on the train on the way here today. I can promise that there is no machine learning involved in this talk. Here ML stands for Mike Lockwood. [Laughter.] The title of my talk is 'Semi-annual variation and the Russell–McPherron effect in geomagnetic activity'.

Geomagnetic activity peaks at the two equinoxes: we see it in all geomagnetic indices, we see it in aurorae and space weather. I am going to talk about the Russell–McPherron effect. We can break it down by looking at the time of calendar year *versus* the time of day because the Earth's magnetic field is spinning and we see a UT variation. The pattern is in the shape of an hour glass which we see in all global geomagnetic indices.

There have been three proposals for explaining the semi-annual variation. One is the so-called axial effect which was first suggested back in 1912, the equinoctial effect was first suggested by Bartels, also called the McIntosh effect, and then there is the Russell–McPherron effect. The solar magnetic equator is inclined to the ecliptic plane by seven degrees so that the Earth travels up and down in heliographic latitude and so the pattern depends on the time of year.

*Ulysses* has plotted the magnetic field and the velocity of the solar wind and you can see that the velocity is larger over the poles at sunspot minimum but relatively slow down the streamer belt which is close to the ecliptic but actually orientated to the magnetic equator of the Sun. At solar minimum we can go up and down in terms of solar-wind speed — that is a very sharp gradient. As Earth goes up and down in heliographic latitude we might see an effect that is known as the axial effect.

The McIntosh effect, first proposed by Bartels, says that the Earth's magnetic axis is rotating around the rotational axes because the poles are offset. Our magnetic field is not actually geocentric which is used in almost all models. It does not pass through the centre of the Earth so that adds a complication. McIntosh and Bartels proposed that the angle $\psi$ between the direction towards the Sun and our magnetic axis which has both a time-of-year and time-of-day variation gives a pattern which is named after McIntosh.

The last one is called the Russell–McPherron effect and we have a rotating Earth and a rotating magnetic axis which varies. That works on the idea that the magnetic field we receive from the Sun is just about in the solar magnetic equatorial plane — any north or south field as it left the Sun is completely wiped out by the time it reaches us. We are interested in what component of Earth's magnetic field points southwards in Earth's magnetic frame, because we may get magnetic reconnection and energy coupling within the magnetosphere when it points southwards in a magnetic frame and that gives a semi-annual variation pattern, but the UT of the two peaks is quite different.

Looking at the three patterns that we observe appear equinoctial, the big surprise is that, although there is an equinoctial effect, the most important factor is the Russell–McPherron [RM] effect. That has caused years of misunderstanding. How do we know it is the most important? Looking from the Sun at the Earth the field is normally in this equatorial plane and at the September equinox the magnetic axis of the Earth is tilted over so we actually require negative $B_z$ and to get that we require positive $B_y$ and GSEQ [geocentric solar equatorial coordinate system], whilst in March it is the other way round. Is the solar field E or W? The big surprise is that that really matters.

We can see that the RM effect really matters if we look at the southward field in the magnetic frame against the southward field in the solar frame. What gives the large southward field in the Earth's magnetic frame is the large southward field in GSEQ. It's actually the large southward field from the Sun and coronal mass ejection that is giving us the southward field in Earth's frame and the RM effect is very small compared to the effect of the transient events.

Why has that led to confusion? There are indices such as the DST [disturbance storm time] index that show an axial effect to add to the confusion. What that is about is time-scale, taking the RM pattern over one-hour data and averaging over six hours it still looks RM-like. Average over 12 hours and it turns into an axial pattern; DST has a response time of more than 12 hours and that is why the axial effect sometimes crops up. It's just a matter of averaging the data makes it look axial, you have averaged out the UT variation.

Where does the equinoctial pattern come from? If we take the pressure of the solar wind and we look at the equinoctial pattern in our geomagnetic data, we find that it emerges as the pressure of the solar wind goes up; the pattern amplitude is a function of pressure. It is something to do with squeezing the magnetosphere and what we think happens is the component of the solar-wind pressure which squeezes the near-Earth tail and you squeeze it more if the solar-wind dynamical pressure is higher. That gives us a clue about where that equinoctial pattern is coming from. It is to do with the tail. What we find from models is that if we get a bending of the tail in the northern summer, it affects the stability of the tail and that is where the equinoctial path is coming from. It's nothing to do with transfer of energy into the magnetosphere — that is all RM. The modulation that confused us is how that energy which is stored as the magnetic field in the tail and how that is released and deposited — that is where the equinoctial pattern comes from and we know that because of the effect of squeezing.

Recently I have been working on the density of the thermosphere at a satellite altitude of 400 km put together from a time series of orbital decay measurements from spacecraft such as *SWARM*. Looking at the annual variations in that density, you see the equinoctial peaks and at the times you expect for the RM effect. There is an annual variation because it depends on UV illumination, so if I subtract the annual variation out, I start to see an equinoctial peak which has nothing to do with the sense of $B_y$ so it's not RM but there is an additional RM effect on top of it. The bending of the tail affects the stability of the tail and that is really where the equinoctial pattern is coming from. If I take the threshold for sorting the IMF $B_y$ component you see that that difference starts to grow so it depends on the strength of the IMF. People have proposed all sorts of reasons for the equinoctial peaks in photospheric densities and they have not really looked at the space-weather element which might be more important than we thought. [Applause.]

*The President.* I meant to say at the beginning that we were the only Society not to have a licence but Geological Society has one so we will have a drinks reception and it will be next door. The next meeting will be on Friday, February 13. What can possibly go wrong?