

QuakeMigrate: a Python Package for Automatic Earthquake Detection and Location Using Waveform Migration and Stacking

Tom Winder *,^{†,1,2}, Conor A. Bacon *,^{†,1,3}, Jonathan D. Smith ⁴, Tom Hudson ⁵, Robert S. White ¹

*Equally contributing authors, [†]Corresponding authors: tomwinder@hi.is, conor.bacon@cantab.net, ¹Department of Earth Sciences, University of Cambridge, Cambridge, United Kingdom, ²Institute of Earth Sciences, University of Iceland, Reykjavík, Iceland, ³Lamont-Doherty Earth Observatory, Columbia University, New York, United States, ⁴British Antarctic Survey, Cambridge, United Kingdom, ⁵Department of Earth and Planetary Sciences, ETH Zurich, Zurich, Switzerland

Author contributions: *Methodology:* Tom Winder, Conor A. Bacon. *Software:* Tom Winder, Conor A. Bacon, Jonathan D. Smith, Tom Hudson. *Formal Analysis:* Conor A. Bacon, Tom Winder. *Writing - Original draft:* Tom Winder, Conor A. Bacon. *Writing - Review & Editing:* Jonathan D. Smith, Tom Hudson, Robert S. White. *Visualization:* Tom Winder, Conor A. Bacon. *Supervision:* Robert S. White. *Funding acquisition:* Robert S. White.

Abstract QuakeMigrate is a modular, open-source Python package providing a framework to efficiently and automatically detect and locate microseismicity. From raw seismic waveforms and a modest number of physically meaningful inputs and tuneable parameters, it produces catalogues of earthquakes including hypocentres, their associated uncertainties, phase arrival times, and local magnitude estimates, all of which are compatible with other common software within the field of (micro-)seismic analysis. The waveform migration and stacking approach on which the software is founded enables phase arrivals with very low signal-to-noise ratios at individual stations to be successfully incorporated in the network-based detection and location. It also implicitly associates phase arrivals across a network even at very small inter-event times, which commonly poses difficulties for traditional pick-then-locate methods. Here, we use a synthetic test-case to outline the fundamental concepts that underpin QuakeMigrate and its implementation, before presenting its application to two real-world datasets with different characteristics to highlight its flexibility and performance.

Non-technical summary QuakeMigrate is a software package written in Python that has been designed to build catalogues of earthquakes from raw, continuous seismic data, with only a limited number of parameter choices from the user. The technique is particularly powerful when used in scenarios where the strength of the vibrations generated by the events of interest are small (e.g., microearthquakes in volcanic environments, icequakes generated by slip at the base of glaciers, or explosions) or where there are large numbers of events in a short space of time. Here, we use a synthetic test-case to outline the fundamental concepts that underpin QuakeMigrate and its implementation, before presenting its application to two real-world datasets with different characteristics to highlight its flexibility and performance.

Production Editor:
Yen Joe Tan
Handling Editor:
Hongyu Sun
Copy & Layout Editor:
Théa Ragon

Received:
July 25, 2025
Accepted:
April 8, 2026
Published:
May 7, 2026

1 Introduction

Detecting and locating microearthquakes from continuous waveform records is the fundamental step in microseismic processing. Increasing use of dense local seismic networks and the widespread adoption of storing continuous waveforms instead of triggered event data has led to a huge increase in the quantity of available data (e.g., Arrowsmith et al., 2022). With this abundance of data comes the theoretical possibility of detecting large numbers of far weaker events, but—when viewed on seismic records from individual stations—their phase arrivals are often difficult to distinguish from noise. Traditional workflows for event detection—based on manual inspection of

seismograms—become intractable in the face of such large quantities of data, particularly in settings with extremely high event rates (e.g., Cesca and Grigoli, 2015; Grigoli et al., 2013, 2017, 2018; Ágústsdóttir et al., 2019; Smith et al., 2020; Niemz et al., 2020). Consequently, automated techniques are fundamental to exploiting modern recordings and computing facilities so as to generate highly complete earthquake catalogues and maximise the new insights that can be gained from them.

The standard and most widely used approach for automated event detection entails first identifying discrete phase arrivals (making ‘picks’) by considering the seismic traces recorded at each station individually, before comparing and combining (‘associating’) the picks made at stations across the network (e.g., Grigoli et al., 2018). An event detection is declared if enough picks are successfully associated to a common source, the location of which can then be inverted for. However, there is a clear trade-off between lowering the detec-

*Equally contributing authors

[†]Corresponding author: tomwinder@hi.is, conor.bacon@cantab.net

tion threshold—in the desire to make more picks and produce a more complete catalogue—and the resulting increase in mis-picks. Even in low-noise conditions, and despite significant recent advances in phase association algorithms (e.g., Ross et al., 2019; Zhu et al., 2022; Münchmeyer, 2024), associating phase picks to individual seismic events is a challenging procedure, particularly so in the presence of mis-picks. In addition, microseismic datasets often contain large numbers of events with short inter-event times, which may be so closely spaced that their phase arrivals overlap in time. Incorrect phase identification or association can result in missed or duplicate detections and significant location errors.

More recently, a second class of waveform-based approaches has been developed, derived from the waveform migration and stacking techniques used in active-source seismology (Kao and Shan, 2004; Baker et al., 2005). Instead of reducing the information recorded at each seismometer to discrete time picks, the waveforms recorded across the entire network are combined to perform a grid-search over time and space for coherent sources of energy in the subsurface (Kao and Shan, 2004; Baker et al., 2005; Hansen and Schmandt, 2015; Grigoli et al., 2013; Shi et al., 2019; Poiata et al., 2016; Grigoli et al., 2016; Shi et al., 2022; Beaucé et al., 2023). The governing principle behind this procedure is to focus or reconstruct the source energy at the time and location that corresponds to the earthquake origin time and hypocentre, and is variably referred to as: source-scanning, Kirchhoff Reconstruction, waveform migration and stacking (used hereafter for clarity), diffraction stack imaging, backprojection imaging, beamforming, coalescence mapping or coherence scanning (e.g., Drew et al., 2013; Cesca and Grigoli, 2015; Li et al., 2020). In doing so, these techniques exploit the coherence between information recorded at stations across a network, significantly improving detection capability and robustness when considering low signal-to-noise ratio (SNR) microearthquakes (Grigoli et al., 2018). These benefits have been shown in a diverse range of microseismic monitoring applications, including in industrial contexts such as mining and geothermal exploitation, in the study of natural seismicity at local and regional scale, and at volcanoes and glaciers (e.g., Kao et al., 2006; Grigoli et al., 2017; Cesca and Grigoli, 2015, and references therein).

However, while the waveform migration and stacking approach provides great promise for automation, in existing implementations there remain several steps which require manual intervention, particularly at the detection stage and in distinguishing real events from artefacts in the final quality-control check. Furthermore, the benefit of utilising the otherwise redundant information contained in the full waveform may result in significantly increased computation time compared to pick-based routines. These limitations have thus far hindered the widespread application of this technique.

Here, we present a novel approach to the waveform migration and stacking workflow, which seeks to overcome these challenges. QuakeMigrate is a computationally efficient open-source software package for the

automatic detection and location of microseismicity. The package is written in Python to be accessible, easy-to-use, and fully cross-platform—relying on only a small number of widely used Python libraries: NumPy (Harris et al., 2020); SciPy (Virtanen et al., 2020); PyProj (Snow et al., 2025; Evenden et al., 2025); pandas (McKinnney, 2010; pandas development team, 2026); and Matplotlib (Hunter, 2007). The core migration and stacking routines have been compiled from C and wrapped in Python in order to maximise their execution speed. The ObsPy library (Beyreuther et al., 2010) provides parsers for all common seismic data formats, making it straightforward to apply QuakeMigrate to existing archives of waveform data. The principles of adaptability, extensibility, and compatibility have motivated much of the design and implementation of the package, which uses a modular architecture with well-defined interfaces to facilitate rapid application to seismic data from a wide range of settings. The user interacts with the software through simple high-level functions, but the code is thoroughly commented and documented throughout to promote extension or augmentation by the user-base, and to ensure that this code is not provided as a black box.

2 Background and theoretical approach

A thorough discussion of the variety of waveform-based location techniques developed so far is provided in Li et al. (2020). Here we briefly outline the relative merits and limitations of these differing approaches, and describe the theory underpinning QuakeMigrate.

Across all implementations of the waveform-based approach to seismic event detection and location, the representations of the waveforms that are migrated, and the stacking (or imaging) operator used to combine them at each grid point, govern the overall computational cost, sensitivity, noise robustness, and location resolution. At one extreme, full-waveform inversion (FWI) can be used to simultaneously obtain high-fidelity source locations and an estimate of the source moment tensor (e.g., Willacy et al., 2019), though at an often prohibitive computational cost, particularly for microseismic events with high source frequencies. Time-reverse imaging (TRI) represents a compromise, consisting of simply reverse-propagating the recorded waveform back into the subsurface, according to the velocity of the medium, necessitating only a single simulation of seismic wave propagation (e.g., Larmat et al., 2006; Werner and Saenger, 2018). However, unlike for the isotropic sources used in active-source seismology, the waveforms from natural earthquakes must be corrected for the source mechanism, to prevent destructive interference of waveforms originating in different quadrants of the focal sphere. Furthermore, a dense search grid and sensor array, and accurate velocity model, are still required to avoid aliasing, meaning that even this more streamlined approach remains extremely computationally demanding. To overcome these limitations, the concept of ‘partial waveform stacking’ (PWS) (Pesicek et al., 2014) was devel-




Figure 1 Cartoon illustrating the migration and stacking search algorithm for an incorrect (top) and correct (bottom) trialed source location and origin time. The columns, from left to right, show: a 2-D slice of the search volume, where the black triangles denote seismic stations, the pink star is the true location of the synthetic source, and the coloured squares represent the currently (blue) and previously (grey) trialed source locations; P- (red) and S-wave (blue) onset functions, time-shifted by the traveltimes corresponding to the currently trialed source location at a trialed origin time, denoted by the vertical purple dashed lines (the true origin time is denoted by the vertical grey dashed line); and a map view of the corresponding stacked 2-D coalescence function, the peak of which represents the computed source location. Onset functions are ordered by distance between the station and the currently trialed source location, so differ between the top and bottom rows. Figure design after Grigoli et al. (2018).

oped as a further simplification of FWI, and stands out as the most widely used and successful implementation of the waveform-based earthquake detection and location techniques introduced to date (Li et al., 2020).

PWS is a hybrid approach, where primary phases contained in the earthquake waveform (e.g., P- and S-waves) are considered through transformation of the seismic traces recorded at each station using a characteristic function (CF). Various CFs have been designed to distinguish the changes in seismogram amplitude, polarisation, or frequency content that mark the onset of P- and S-phase arrivals. Many of these algorithms share their foundation with those developed for automatic phase picking, including the short-term average to long-term average ratio (STA/LTA) of the seismogram amplitude (Drew et al., 2013; Grigoli et al., 2013; Hansen and Schmandt, 2015), higher-order statistics such as kurtosis (Langet et al., 2014), and the Hilbert transform (or envelope function). However, instead of solely extracting the timing of peaks that exceed a user-defined threshold (‘picks’), here all of the information in the transformed signal is retained. Hereafter referred to as onset functions, these continuous signals are designed to peak around the time of seismic phase (e.g., P and S) arrivals and to be always positive to ensure constructive stacking when migrated.

The detection and location process then comprises three key steps. First, the continuous waveforms

recorded at each seismic station are reduced to continuous onset functions for each seismic phase of interest, via the chosen characteristic function. These may be calculated based on individual components (typically the vertical sensor for P-waves), or a combination thereof (typically combining horizontal components for S-waves). Second, the onset functions are continuously migrated into a 3-D subsurface grid, where each grid point represents a potential hypocentre location, and each timestamp a candidate earthquake origin time. At each grid node, the onset functions from every station in the network are aligned according to a pre-calculated traveltime lookup table for each station and phase, and combined according to the selected stacking function. This is repeated for a range of potential origin times, resulting in a 4-D image (or “coalescence function”) representing an exhaustive search in space and time. Where onset functions from stations across the network stack coherently—or coalesce—at a particular grid point and origin time, this will produce a peak in the value of the 4-D coalescence function, indicating an event hypocentre. The third and final step comprises identifying these peaks through space and time, in order to produce an event catalogue.

The workflow introduced within QuakeMigrate differs from previous implementations in that the migration is performed twice, separating the detection and location workflow into three stages. First, in the ‘De-

tect' stage, an exhaustive scan is performed on a decimated search-grid, optimised for computational efficiency. In the 'Trigger' stage the resulting continuous time series of coalescence values is analysed to identify peaks corresponding to the times and locations of candidate events. Finally, in the 'Locate' stage, a short window of time around the timestamp of each triggered event is re-migrated on a higher resolution grid, and location uncertainties are reported along with several other optional outputs. This novel approach significantly increases the computational efficiency of the procedure, and care has been taken to ensure that considerable speed-ups can be achieved without any reduction in detection performance.

3 Method: Software functionality and workflow

The QuakeMigrate workflow has been designed to combine computational efficiency with user-friendliness, consisting of simple steps which are described in detail in the documentation and demonstrated in the example usage scripts available in the QuakeMigrate GitHub repository (<https://github.com/QuakeMigrate/QuakeMigrate>). The modular architecture of the package (Figure 2) promotes transparency in the inputs and outputs of each stage, and overall facilitates straightforward adaptation, customisation, and extension of the core workflow. Our intention is for the package to serve not only as an efficient implementation of the coalescence method outlined in [Drew et al. \(2013\)](#), but also to provide a common framework for the broader class of migration-based techniques. To this end, a number of the modules (e.g., the 'Onset' module) have been designed as plug-ins, with well-defined interfaces such that alternatives to the STA/LTA characteristic function can be incorporated and comparisons can be made (see Supplementary Figure S1 for an overview of the QuakeMigrate workflow).

In the following sections, we illustrate the principal components of the package using a synthetic case study. Details of how the input data are simulated are provided in the Supplementary Material.

3.1 Installation

The QuakeMigrate software package is available for download from the Python Package Index (PyPI), with builds available for a number of major operating systems and current Python versions. Standard maintenance, updates, and issues are performed and tracked in a public GitHub repository. Data, scripts, and additional commentary are available through this public repository for each of the examples presented below, along with several additional example use cases. In addition to this, we perform comprehensive testing and continuous integration to ensure that the software remains consistent across updates and improvements. Contributions, discussions, and collaborations are encouraged through this public platform. Community participation and support are vital for the success of

open-source projects like QuakeMigrate (see Acknowledgments).

3.2 Generation of traveltimes lookup tables

The pre-computation of traveltimes lookup tables (LUTs) is fundamental to all migration techniques. Within QuakeMigrate, the 3-D grid chosen for the lookup table represents the search grid; the nodes over which the migration and stacking will be performed (though this can later be decimated to give a larger node spacing). Pre-computing traveltimes for each station and phase ensures this computationally expensive step is only performed once.

Several methods for traveltimes calculation are provided within QuakeMigrate, including support for either uniform or 1-D velocity models. LUTs calculated by other means can be loaded via conversion to the format used by the NonLinLoc software package ([Lomax et al., 2000](#)). This provides flexibility for the user to use velocity models which include, for example, lateral velocity heterogeneity, or wavespeed anisotropy. The user can also choose to include any phase of interest: using at least P- and S-phases will result in a significant improvement in detection capability and hypocentre resolution compared to considering only P-wave arrivals (e.g., [Langet et al., 2014](#), Supplementary Figures S5 and S6).

For the synthetic example, we define a search volume spanning (0.3°, 0.3°, 30 km), centred on (0°, 0°, 15 km), and randomly generate a seismic network with 10 station locations randomly sampled from a uniform distribution in latitude, longitude, and elevation (Supplementary Figure S2). The traveltimes for each station and phase are computed for a simple 1-D velocity model using the eikonal solver of [Podvin and Lecomte \(1991\)](#) provided within the NonLinLoc software, and wrapped in the QuakeMigrate LUT module (Figure 2). Synthetic waveforms are generated for a source with a hypocentre of (0.0°, 0.0°, 15 km) (see the Supplementary Material for further details).

3.3 Reading continuous waveform data

Raw waveforms are parsed into the migration engine from a local archive of continuous waveform data (Figure 2, Supplementary Figure S1). Any regular archive structure can be accommodated, with explicit support for a number of standard formats (e.g., SeisComp Data Structure, SDS). All common seismic data formats are supported (including miniSEED, SAC, and SEG-Y), thanks to the parsers provided by the ObsPy library ([Beyreuther et al., 2010](#)).

Though the migration-based approach generally improves robustness to uncorrelated noise, assessment of data quality is still strongly recommended prior to any processing with QuakeMigrate. Sharp signal offsets—due to instrument failures, for example—may lead to peaks in the various onset functions used for migration. If sufficiently high in amplitude, these spurious peaks may still dominate the stack, potentially leading to false triggers or poorer quality locations. Inclusion

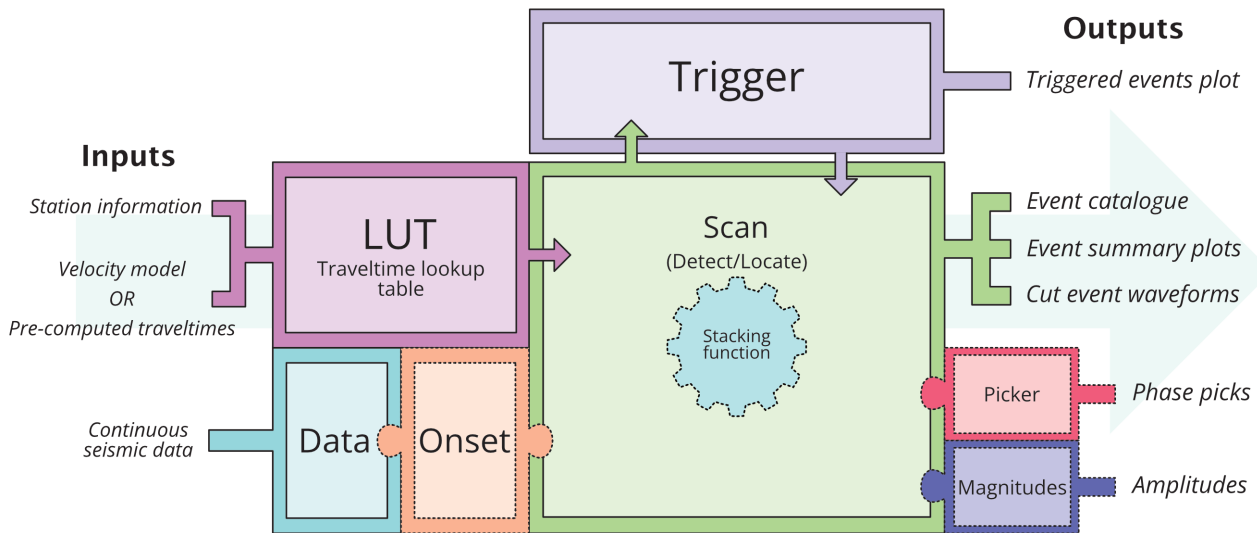


Figure 2 Schematic illustrating the structure of the QuakeMigrate software package. Sections with dashed outlines (Onset, Picker and Magnitudes) are plugin/extension modules of the code that can be customised or substituted for alternatives by the user; arrows indicate the sequence of steps in the processing workflow. LUT = traveltime lookup table (Section 3.2).

of data with timing errors will also inevitably be detrimental to performance.

3.4 Calculation of onset functions

The transformation of raw waveforms to an always-positive function which peaks at the onset of a seismic phase of interest (here referred to as ‘onset functions’) is core to the migration and stacking approach, as it facilitates constructive interference at the true source location. The default onset function in QuakeMigrate considers only the amplitude of the waveform, through taking the ratio of the short-term average to long-term average absolute amplitude, i.e., the STA/LTA function (Figure 3). Despite its simplicity, it has been widely used (e.g., Grigoli et al., 2013; Drew et al., 2013; Hansen and Schmandt, 2015; Cesca and Grigoli, 2015), and shown to provide good resolution, robustness in the presence of noise, and relatively low sensitivity to model errors (Beskardes et al., 2018; Cesca and Grigoli, 2015).

Furthermore, Drew et al. (2013) showed that where an appropriate choice of short- and long-term window lengths is made, the widths of peaks in the onset functions correspond to the timing uncertainty in the corresponding phase arrivals, while their amplitude corresponds to the signal-to-noise ratio (SNR). When combined through multiplication, the resulting 3-D coalescence maps can be interpreted as (non-normalised) probability density functions describing the source location uncertainty. The uncertainty in the returned event locations can thus be characterised without resorting to (often computationally intensive) statistical re-sampling techniques, such as bootstrapping (e.g., Grigoli et al., 2013).

Advances in phase identification are continuously being made, with ever more information from the raw signal (e.g., phase, frequency, and polarisation) being in-

corporated, both analytically (e.g., Grigoli et al., 2014; Cesca and Grigoli, 2015; Shi et al., 2019; Trojanowski and Eisner, 2017; Beskardes et al., 2018) and via machine learning approaches (e.g., Ross et al., 2018; Mousavi et al., 2020; Lapins et al., 2021; Woollam et al., 2019, 2022; Shi et al., 2022; Beaucé et al., 2023; Isken et al., 2025). In anticipation of continued future improvements, this module has been implemented as a ‘plug-in’ (Figure 2, Supplementary Figure S1) such that any such transformation might be used, and rigorously compared within a single framework. In addition, while for sensors deployed on the ground surface the vertical and horizontal component seismograms will dominantly record P- and S-wave arrivals, respectively, this might not be the case for, e.g., borehole deployments (Drew, 2010), or other unconventional sensors such as fibre-optic cables (Hudson et al., 2025), or when using nodal arrays comprised solely of vertical-component sensors. By default, the S onset function is calculated from the two horizontal seismograms, and P from the vertical, but the user may choose different channel mappings, and/or design an onset function that uses polarisation analysis to help distinguish between phases.

Though the strict relationships between arrival time uncertainty and arrival SNR to the width and height of the peaks will not necessarily be maintained with alternative onset functions, the theory may be extended so long as the mapping is approximate (Drew et al., 2013). QuakeMigrate will thus still provide robust estimates of the *relative* timing uncertainties and SNRs between phase arrivals, and thus between events.

3.5 The stacking operator

At each timestamp (candidate event origin time), t , an instantaneous 3-D coalescence map is calculated by migrating the onset functions according to the trav-

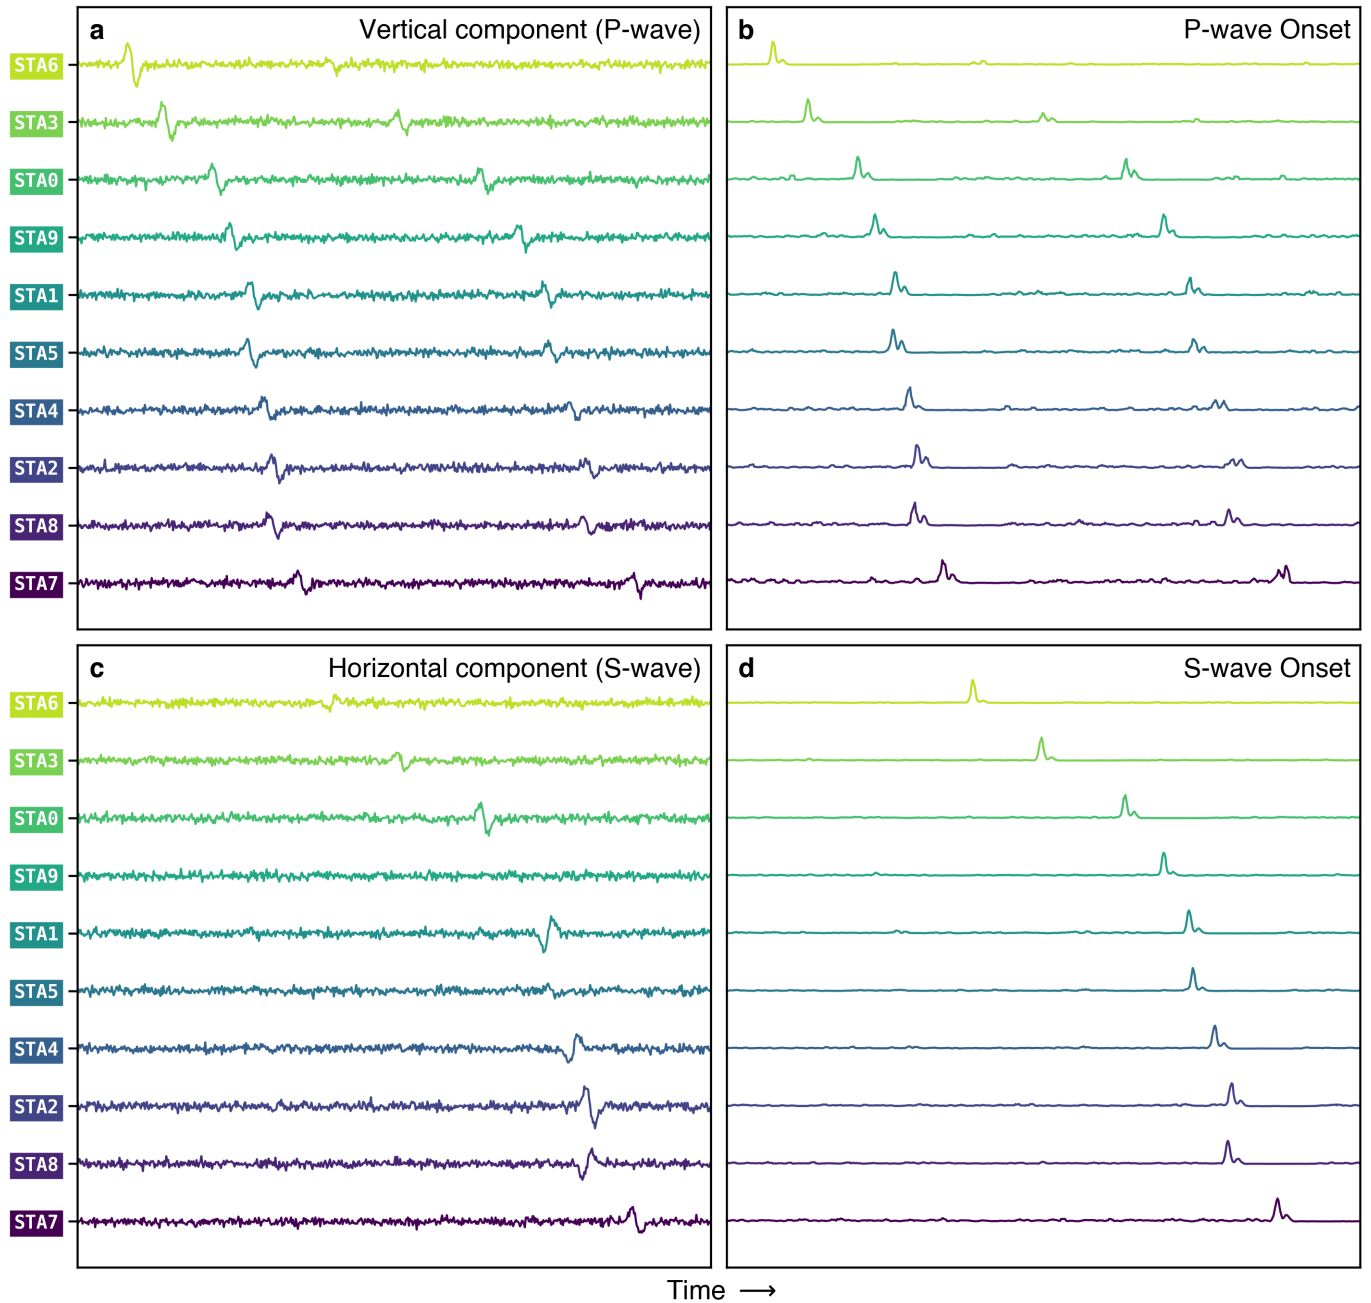


Figure 3 Schematic diagram illustrating the transformation of raw seismic data to onset functions. Here pre-processed seismic data (left; bandpass filtered between 1–10 Hz) are transformed into onset functions (right), sensitive to P- and S-wave phase arrivals, respectively. The P-wave onset functions (panel b) are calculated from the short-term average/long-term average ratio (STA/LTA) of the amplitude of the vertical component seismograms (panel a); S-wave onset functions (panel d) are calculated in the same manner from each of the horizontal components (panel c) with a single combined S-wave onset function calculated as the root-mean-square of the north and east component onset functions, where the default channel mapping is used. Note that for clarity only one of the two horizontal component seismograms is shown in (c); in cases where a stronger arrival is measured on the other component (e.g., STA5) a large combined S-wave onset function peak is still seen in (d).

eltimes, τ , stored in the traveltimes lookup table, and stacking them at each grid point, x_j (Figure 1). Drew et al. (2013) determined that where the onset functions $OF_i(t)$ approximate continuous (non-normalised) probability density functions describing the likelihood and timing of a phase arrival, taking the root of their product allows us to calculate a coalescence function, f_C , as

$$f_C(t, x_j) = \sqrt[n]{\prod_{i=1}^n OF_i(t + \tau_i(x_j))}, \quad (1)$$

the calculation of which can equivalently be expressed as the exponent of the arithmetic mean of their logarithms,

$$f_C(t, x_j) = \exp \left\{ \frac{\sum_{i=1}^n \ln(O F_i(t + \tau_i(x_j)))}{n} \right\}, \quad (2)$$

where n is the number of input onset functions, which is a more computationally robust approach, through avoiding overflow. This calculation is repeated over a range of timestamps to produce a 4-D coalescence function which describes the focussing of seismic energy in the subsurface through time (Figure 1).

As with the onset function calculation, the modular implementation of the stacking function makes it possible in principle to substitute it with an alternative (Figure 2); for example to explicitly calculate the coherence between P and S onset functions by taking their dot-product, as in [Grigoli et al. \(2013\)](#). This might improve robustness to false detections due to under-migration of true events (the “shooting star” artefact described by [Beskardes et al. \(2018\)](#)), but precludes the detection of atypical events, for example due to landslides, or exotic sources in volcanic environments, where the P or S phase arrivals are highly emergent, indistinct, or even absent altogether (see Sections 4.2 and 5.1).

3.6 Detect—exhaustive search in space and time

It is important to note that the Detect stage is not dependent on any particular theoretical basis for the combination of onset functions, as it is, in essence, a purely pragmatic exercise in identifying candidate events from noise, rather than attempting to locate them accurately. Recognition of this distinction provides the motivation to separate the Detect from the Locate stage of QuakeMigrate, allowing each routine to be optimised individually for its respective goals. This unique approach provides significant benefits in computational efficiency and in modularity, increasing both transparency to the user and flexibility to optimise the workflow to adapt to the differing challenges and research goals associated with different datasets (Supplementary Figure S1).

The default stacking operator is equivalent to taking the geometric mean of the onset functions. In comparison to the arithmetic mean, the geometric mean is significantly less sensitive to outliers, and—in contrast to taking the median of the migrated onset functions (e.g., [Hansen and Schmandt, 2015](#))—is still sensitive to their coherence. It is these characteristics that make this operator suitable for Detect, as well as for its originally intended purpose in Locate.

The search for the location and amplitude of the maximum coalescence value in the grid at each timestamp, $\hat{f}_C(t)$, is also performed by the core compiled C migration library, and for the purposes of Detect, a ‘normalised’ maximum coalescence value, \hat{f}_{NC} as

$$\hat{f}_{NC}(t) = \frac{\hat{f}_C(t)}{\left(\frac{\sum_{j=1}^N f_C(t, x_j)}{N} \right)}, \quad (3)$$

is also returned (where N is the total number of nodes in the search grid), defined as the instantaneous maximum coalescence value in the grid divided by the mean coalescence value across the 3-D grid at that timestamp. The normalised coalescence value is comparatively less sensitive to global changes in the coalescence value across the search grid, usually due to changes in station availability, or a drop in SNR across the network during the passing of coda waves after large events. This is useful in maximising the extent to which the Trigger stage can be automated, even in challenging scenarios such as during intense seismic swarms.

3.7 Trigger—identifying candidate events

The next challenge is to identify, from the continuous migration outputs, a catalogue of candidate events (Supplementary Figure S1). While in essence this is a task of peak finding, the behaviour of the coalescence time-series is considerably complicated by a range of factors that may influence the detectability of coherent phase arrivals within the background wavefield. Beyond the size and timing of earthquakes that occur within the search grid, other influences include changes in noise amplitude or noise source properties, the number of stations with usable data during a given time-step, the source properties of earthquakes which occur (their frequency content, duration, impulsiveness, and the network geometry with respect to their location), and the inter-event time. These factors will differ strongly from network to network, and may also vary considerably through time for any given network.

A simple, yet effective, approach is to use a static trigger threshold. Whenever the 1-D coalescence function exceeds this threshold, a candidate event is recorded in the triggered events catalogue. A minimum temporal separation between consecutive events (a minimum event interval, or MEI) can also be specified, which helps prevent multiple triggers from a single event as the coalescence converges to a maximum (Figure 4). The trade-off here is clear: lowering the threshold will increase the sensitivity of the method to smaller events, but may also result in spurious triggers.

This basic concept can be further developed to combat some of the challenges posed by real datasets. Dynamic thresholding, where the threshold for some period of time is assessed based on the available data, can help with handling changes in the coalescence function baseline over the course of hours (perhaps related to an uptick in seismic activity caused by an earthquake swarm, or intermittent telemetry dropouts), days (diurnal variations in the ambient noise field), or even longer (changing network geometry). A range of options are already available within the package—including based on the Median Absolute Deviation (MAD), and a multiplier of the median value of the signal—but this module has been designed to be extensible, such that future improvements can be straightforwardly incorporated (Supplementary Figure S1).

Since each timestamp in the 1-D coalescence function corresponds to a point in space (from the 4-D search), each peak is associated with a source location. These




Figure 4 Annotated triggered events plot illustrating the process by which candidate earthquakes are triggered from the normalised maximum coalescence time series. Periods where the coalescence value (black line) exceeds the detection threshold (dashed blue line) are identified as candidate events (green shaded regions), with a minimum duration determined by the user-specified Marginal Window (MW). Event windows separated by a shorter interval than the specified Minimum Event Interval (MEI), shown by grey shading, are merged, with the largest amplitude peak retained as the single candidate event within that period.

trigger locations allow the user to filter for candidate events within some sub-region of the search grid, either for separate analysis in Locate, or to exclude artefacts caused by mis-located regional events (with true locations outside the search area) which tend to cluster at the grid boundaries and can therefore easily be recognised and removed.

3.8 Locate—re-migration of candidate events on a high-resolution grid

With a set of candidate events in hand from the Detect and Trigger stages (Supplementary Figure S1), the goal of the Locate stage is to calculate more accurate and precise hypocentre locations, and to provide a robust assessment of the location uncertainty, among other outputs. These accompanying statistics may then be used to efficiently and automatically distinguish genuine earthquakes from false triggers, allowing the user to both further tune their choice of trigger threshold if necessary, and to filter the located events to produce a robust earthquake catalogue (e.g., [Winder, 2022](#); [Bacon, 2022](#)).

The migration and stacking process is identical to that used in Detect, though for a much shorter time-step, which makes it possible to perform the migration on a higher resolution grid while maintaining acceptable memory usage and compute time (see Supplementary Material Section S3 for more details). This time-step is taken to encapsulate the marginal window length, defined as the period of time over which the 4-D probability density function is integrated (or ‘marginalised’) to recover the 3-D posterior spatial posterior probability density function (PDF) for the event hypocentre.

The resulting 3-D map can then be treated as a PDF describing the best estimate of the source location and its

spatial uncertainty (Figure 5). A sub-gridded estimate for the location of the peak in the PDF is returned from a sub-sampled 3-D spline function fitted to the coalescence map locally around its maximum. Separately, a 3-D Gaussian function is also fitted locally, after applying a modest Gaussian smoothing, in order to characterise the shape of the peak and a (parameterised) quantitative estimate of the location uncertainty (Figure 5).

The ‘export’ module provides interfaces between the outputs of Locate and several other widely used analysis packages, in order to streamline the incorporation of QuakeMigrate into existing workflows. These are implemented as individual plug-ins, which can readily be added to.

3.9 Computational cost and memory profiling

Detect is generally the most computationally expensive stage of a QuakeMigrate run, and a number of parameters may be tuned to optimise the balance between runtime, memory usage, and detection performance. The effect of these tools is discussed in the Supplementary Material (Section S3) and illustrated in Supplementary Figure S7, and the memory usage for each stage of a QuakeMigrate run is profiled in Supplementary Figures S8–S11. See also Section 4 for discussion of parameter choices for the two real-world example use cases, as well as further guidance in the Documentation that accompanies the software.

4 Application to real-world datasets

We demonstrate the performance and flexibility of QuakeMigrate by applying it to two real-world datasets.




Figure 5 QuakeMigrate event location plot for the synthetic example. a) shows a map view of a horizontal slice through the 3-D coalescence map at the earthquake hypocentre, coloured by normalised coalescence value. b) shows a longitude-depth cross-section. White dashed lines indicate the maximum probability location; black ellipse delineates a Gaussian estimate of the event location uncertainty; white triangles show station locations.

Data, scripts, and instructions for how to run these examples are archived in [Bacon and Winder \(2025\)](#).

4.1 Basal icequakes at the Rutford Ice Stream, Antarctica

Icequakes, a broad class of seismic events associated with the movement or fracture of ice, can provide important insights into the dynamics of glacier flow and deformation. This example has been chosen to illustrate the performance of QuakeMigrate in dealing with exceptionally high event rates, and in a context where the expected distribution of event locations is well understood ([Smith et al., 2015](#); [Hudson et al., 2019](#)), allowing for intuitive assessment of the location performance. The seismic velocity of the ice layer is well-approximated by a constant (homogeneous) model, and the icequakes exhibit high SNR phase arrivals, making this real-world dataset a natural next step from the synthetic example examined previously. Icequakes at the Rutford Ice Stream are generated by a mechanism generally referred to as basal stick-slip, associated with sliding of the ice over the underlying till or bedrock ([Winberry et al., 2009](#)). This leads to them occurring

predominantly within tens of metres of the ice-bed interface, here in distinct spatial clusters ([Smith et al., 2015](#); [Hudson et al., 2019](#); [Kufner et al., 2021](#)). The location performance of QuakeMigrate can thus be easily visualised by comparing the automatically generated QuakeMigrate hypocentres with this expected distribution, particularly in depth.

The critical parameter choices for icequake detection are the bandpass filter frequencies, and the window lengths used to calculate onset functions using the STA/LTA algorithm. Icequakes in this region typically have peak P-wave frequencies of approximately 130 Hz ([Smith et al., 2015](#)), however we find that detection and location performance are unaffected by using a slightly lower frequency band of 20–124 Hz, which allows us to reduce the onset function and scan sampling rate to 250 samples per second (sps), which significantly reduces compute time. The P-wave short-term window length is set at 0.01 s, which is $1.3 \times$ the dominant period of the signal, as this was found to produce high signal-to-noise ratios, likely due to the simple, impulsive nature of P and S phase arrivals typically observed in Antarctic icequakes. The long-term window is set to 0.25 s, adequate for capturing the background noise amplitude, but short enough not to overly smear the onset function peak associated with the extremely sharp phase arrivals. The S-wave filter and windows are set at 10–124 Hz and 0.05 s and 0.5 s, corresponding to their lower frequencies and longer durations. Traveltimes were computed for constant phase velocities of $v_P = 3,841 \text{ ms}^{-1}$ and $v_S = 1,970 \text{ ms}^{-1}$, as in [Smith et al. \(2015\)](#).

For Detect, it was necessary to retain a relatively dense search grid node spacing of 100 m in order to successfully capture the sharp coalescence peaks produced by the extremely impulsive, high frequency phase arrivals produced by these icequakes—experiments with a decimated lookup table resulted in significantly fewer detections. A static trigger threshold was used due to the stability of the network and seismicity rate over the short two-day study period, which made it straightforward to manually select an optimal value. In total, 49,249 candidate events were triggered in 48 hours. This high event rate is illustrated in [Figure 6](#).

Two filters were used to separate real events from false triggers in the preliminary catalogue of 49,249 events, of which 47,354 were successfully located. Both filters have clear physical meanings, which is important in order to allow the effects of applying them to be understood. Firstly, a filter was applied based on the Global Covariance values reported by Locate. This statistic is measured as the covariance in the marginalised 3-D coalescence map after removing values below the 90th percentile, and is designed to characterise the profile of the coalescence peak. Real events typically stand out from the background as sharp spikes ([Figure 5](#)), while artefacts of various forms generally exhibit much more flat-topped or gently sloping maxima, or consist of multiple separate peaks. A representative value over all three spatial dimensions is calculated as the geometric mean, which provides a good representation of their coherence, as well as the “average” magnitude of the three values. In order to make a more in-




Figure 6 5-minute snapshot of the normalised maximum coalescence trace (black line), illustrating the high rate of events at the Rutford Ice Stream. There are 74 triggered events that exceed the static trigger threshold of 3.0 (dashed blue line), corresponding to an event every 4.0 s, slightly below the mean rate of $\sim 25,000$ events per day. This is significantly lower than the cross-grid traveltimes; a scenario in which traditional pick-then-associate methods might struggle.

formed choice of filter value, we perform a statistical analysis of a representative subset of manually labelled events—comprising the first 10.5 minutes of the dataset, containing 149 / 47,354 events—categorised based on inspection of the Locate summary plots. Based on this analysis, we choose to remove events with a geometric mean Global Covariance value of greater than 150 m (Supplementary Material Section S4, and Supplementary Figure S12).

A second filter was applied based on the coalescence value, only selecting events with a coalescence value (roughly equivalent to the network-averaged SNR) of greater than 5.5. This helped to remove a small proportion of spurious scattered events resulting from onset function peaks that may align on a small subset of channels, but not on the majority of stations in the network. This coalescence filter value is higher than would be recommended for most other environments, but is appropriate here due to the exceptionally high SNRs of the icequake phase arrivals, and consequent coalescence values. Conservative values have been selected, in order to highlight the location performance of QuakeMigrate for the best-located events; for these filter values, the comparison to the ground-truth subset of manually labelled events indicates that the expected proportion of true positives is more than 99%, with a false negative rate of approximately 40% (Supplementary Figure S12). Nevertheless, 20,015 events remain in the catalogue, over a period of two days.

The filtered catalogue is displayed in Figure 7, and clearly demonstrates the capability of QuakeMigrate to detect and locate icequakes, and to efficiently produce a clean catalogue of seismicity with few false positives. The total of 20,015 events over two days corresponds to a $\sim 100 \times$ increase in event detections over the same period analysed by Smith et al. (2015), using the CMM algorithm (Drew et al., 2013). Clear spatial clustering

of seismicity is evident, both laterally—indicating sticky spots at the bed (Smith et al., 2015)—and tightly about the ice-bed interface, as expected from the basal stick-slip mechanism (Hudson et al., 2020).

This two orders of magnitude increase in event detections was facilitated in large part by the separation of the QuakeMigrate workflow into three distinct stages (Figure 2, Supplementary Figure S1). This allowed for straightforward identification of a short, representative period of data on which to efficiently experiment, refine, and test both onset function parameters and grid decimations, in order to make an informed choice of settings. This is particularly important in temporary icequake studies, where important factors such as the source frequencies and attenuation properties of the medium are often poorly known, and which would otherwise be both challenging and potentially extremely time consuming to explore. For example, icequakes detected at the base of the Skeiðarárjökull outlet glacier, Iceland, have significantly different characteristics than those presented here, and required significant effort to detect and characterise successfully (Hudson, 2019; Hudson et al., 2019). The comprehensive set of statistics output by QuakeMigrate alongside each earthquake location also plays a pivotal role (Supplementary Material Section S4). This applies in particular to the Global Covariance statistic, which is innovative in being explicitly designed to distinguish real events from artefacts, rather than attempting to describe the location quality under the implicit *a priori* assumption that the statistical measure in question is describing a real event. Its use allows the coalescence value filter threshold to be reduced while still maintaining an exceptionally high level of true- to false-positive detections (Supplementary Figures S12–S13). For further discussion of considerations and challenges in detecting and locating basal icequakes using QuakeMigrate, the reader is referred to

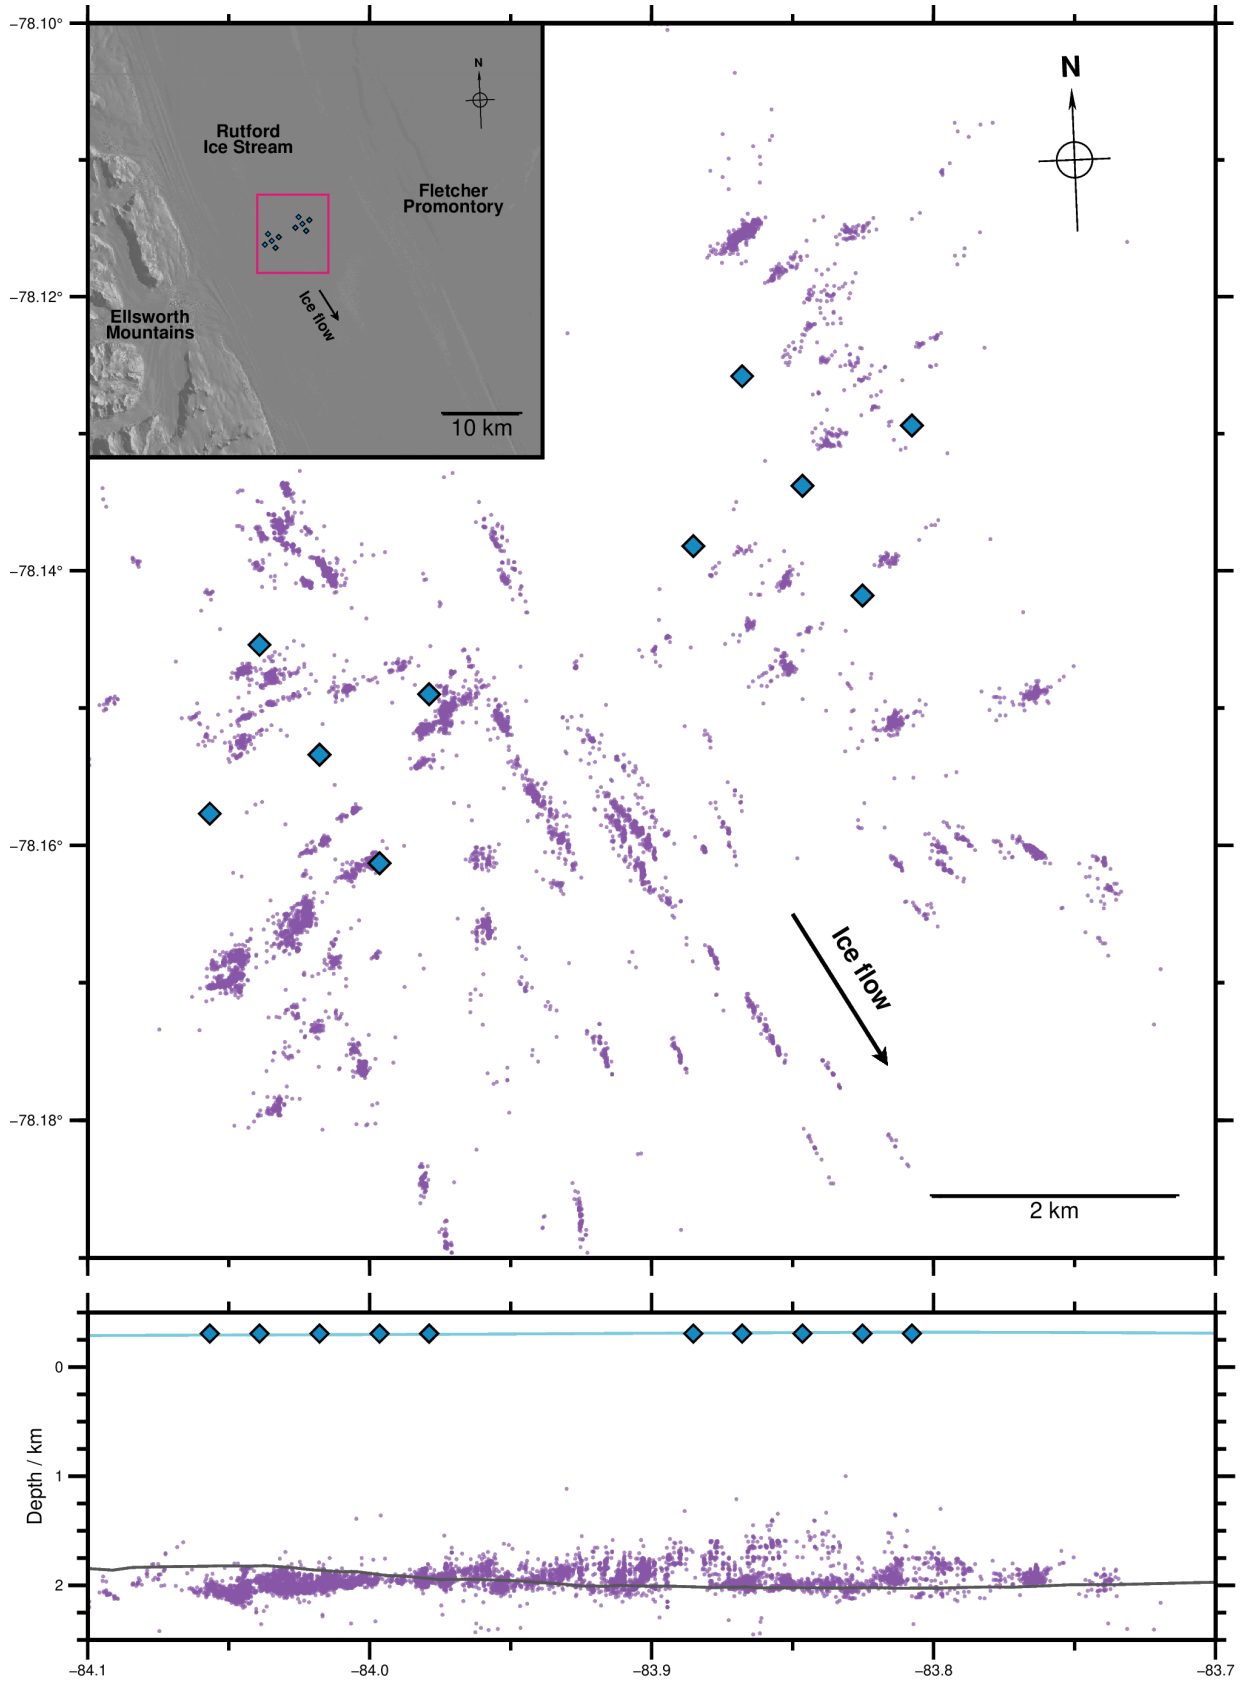


Figure 7 Icequakes detected at the Rutford Ice Stream, Antarctica, from 20–21 January 2009. Scatter points represent 20,015 icequake hypocentres. The ice surface and ice-bed interfaces are indicated by the blue and grey lines, respectively (King et al., 2016); blue diamonds show the locations of seismic stations. Inset shows location within the Rutford Ice Stream. Further details on this dataset can be found in Hudson et al. (2019).

Hudson et al. (2019).

4.2 Tectonic & volcanic microseismicity at Askja volcano, Iceland

The flexibility in defining the characteristic function used to transform raw waveforms into onset functions allows the user to tune QuakeMigrate to perform well for the wide diversity of seismic events found in different settings. This example demonstrates how, using the default STA/LTA onset function, QuakeMigrate provides a powerful tool to detect extremely small microseismic events with a broad range of characteristics. We present earthquakes detected and located in the vicinity of Askja volcano, in the Northern Volcanic Zone of Iceland (Figure 8a, inset), over a 24 hour period on 26th October 2011. Askja is located in the active rift zone at the spreading centre between the North American and European plates, and consequently is highly seismically active (Winder, 2022; Greenfield et al., 2020; Soosalu et al., 2006). Earthquakes in this area can be broadly classified into three categories: tectonic earthquakes related to fault movements along the plate boundary; volcano-tectonic events related to magmatic and hydrothermal systems in the brittle shallow crust; and volcanic “long-period” (LP) earthquakes thought to be directly associated with the movement of volcanic fluids (magma or exsolved volatiles), marking their ascent path through the ductile lower crust (Winder, 2022; Greenfield and White, 2015; Key et al., 2011b,a).

A significantly less dense search grid is used in this case—with 1 km node spacing—reflecting the larger station spacing (~10 km) and lower frequency content (~10 Hz) of the targeted seismicity, which results in a larger expected location uncertainty (or width of the coalescence peak). For Detect this was further decimated to 2 km node spacing, which provides a $\sim 5 \times$ speed-up and memory usage reduction without significantly reducing detection performance (Supplementary Material Section S3 and Supplementary Figure S7). The search grid was designed with a 10 km buffer around each side in map view, and a 5 km padding at the base in depth, to allow candidate events triggered within these regions to be filtered out. This is important in this setting where significant seismic activity occurs all along the Northern Volcanic Zone, which—though not located within the network—will still produce clear phase arrivals on most stations, and sufficiently large coalescence peaks to be triggered at the grid edges. This padding provides a buffer for these events with true locations outside the region of interest to be removed by a simple spatial filter. The 1-D velocity model from Greenfield et al. (2016) is used to calculate the travel-time lookup table, which provides a good approximation despite the significant lateral variation in velocity structure beneath the edifice of Askja itself. Band-pass filters of 2–16 Hz and 2–14 Hz and STA/LTA windows of 0.2/1.0 s were used for the P- and S-phase onset functions, respectively, reflecting microearthquake peak frequencies of around 6–10 Hz (e.g., Greenfield and White, 2015). For Detect, the “overlapping” STA/LTA function is used, which provides a more smoothed on-

set function that is less sensitive to sharp signal offsets (that are most commonly caused by instrument failures), while the “centred” STA/LTA function is used for Locate to provide a more precise location and a better assessment of the location uncertainty. A marginal window of 1.0 s is used in Trigger and Locate, which reflects the expected spatial uncertainty of event locations of around 3–5 km, and seismic velocities at seismogenic depths of around 5 km s^{-1} . A 1-D smoothing is applied to the coalescence timeseries in Trigger before identifying peaks that exceed a static threshold of 1.45. This smoothing (using a gaussian function with a 0.25 s kernel) helps remove short-duration spikes caused by random stacking of incoherent noise, while retaining broader peaks corresponding to true event detections. Within Locate, magnitudes are calculated with the built-in local magnitudes sub-module (which implements the method of Keir et al., 2006), using the attenuation function of Greenfield et al. (2020).

The deep long-period (DLP) seismicity represents a particular challenge for event detection and location. The DLP waveforms consist mostly of low frequencies (~1–3 Hz; Figure 8e), which overlap significantly with the powerful oceanic microseism observed in Iceland. In addition, the phase arrivals are generally emergent, and P waves are anomalously weak, or even absent (Figure 9c). However, the depth of the DLP events (relative to the aperture of the seismic network around Askja), means they are often detected, if weakly, on a large number of stations—even if the SNR of the individual phase arrivals is close to or below the noise level (Figure 9c). Conventional “pick-then-associate” algorithms would likely fail to detect these events, whereas this represents a perfect use-case for network-based detection and location algorithms such as QuakeMigrate. By exploiting the coherence of phase arrivals detected across the network, seismic sources can be successfully identified amongst the incoherent noise (Figure 9c). Importantly, QuakeMigrate is capable of detecting and locating these exotic volcanic LP events without reducing its ability to successfully catalogue the more prevalent shallow, high-frequency tectonic microearthquakes (Figure 8). This demonstrates that the algorithm promises to provide a comprehensive solution to earthquake detection and location in a broad range of settings, including those where events with a broad range of source characteristics occur within a small area.

4.3 Location benchmark: comparison with manually picked earthquakes at Askja

The Askja dataset provides a further opportunity to “benchmark” QuakeMigrate’s location performance, by comparing it to a rich catalogue of manually refined earthquake locations (Winder, 2022). Manual earthquake analysis and location refinement is generally regarded as the gold standard in studying microseismicity with local seismic networks. Over the course of 13 years of study, more than 2,000 shallow earthquakes in the region have been manually analysed. Here we use a subset of the best-located events to compare with locations

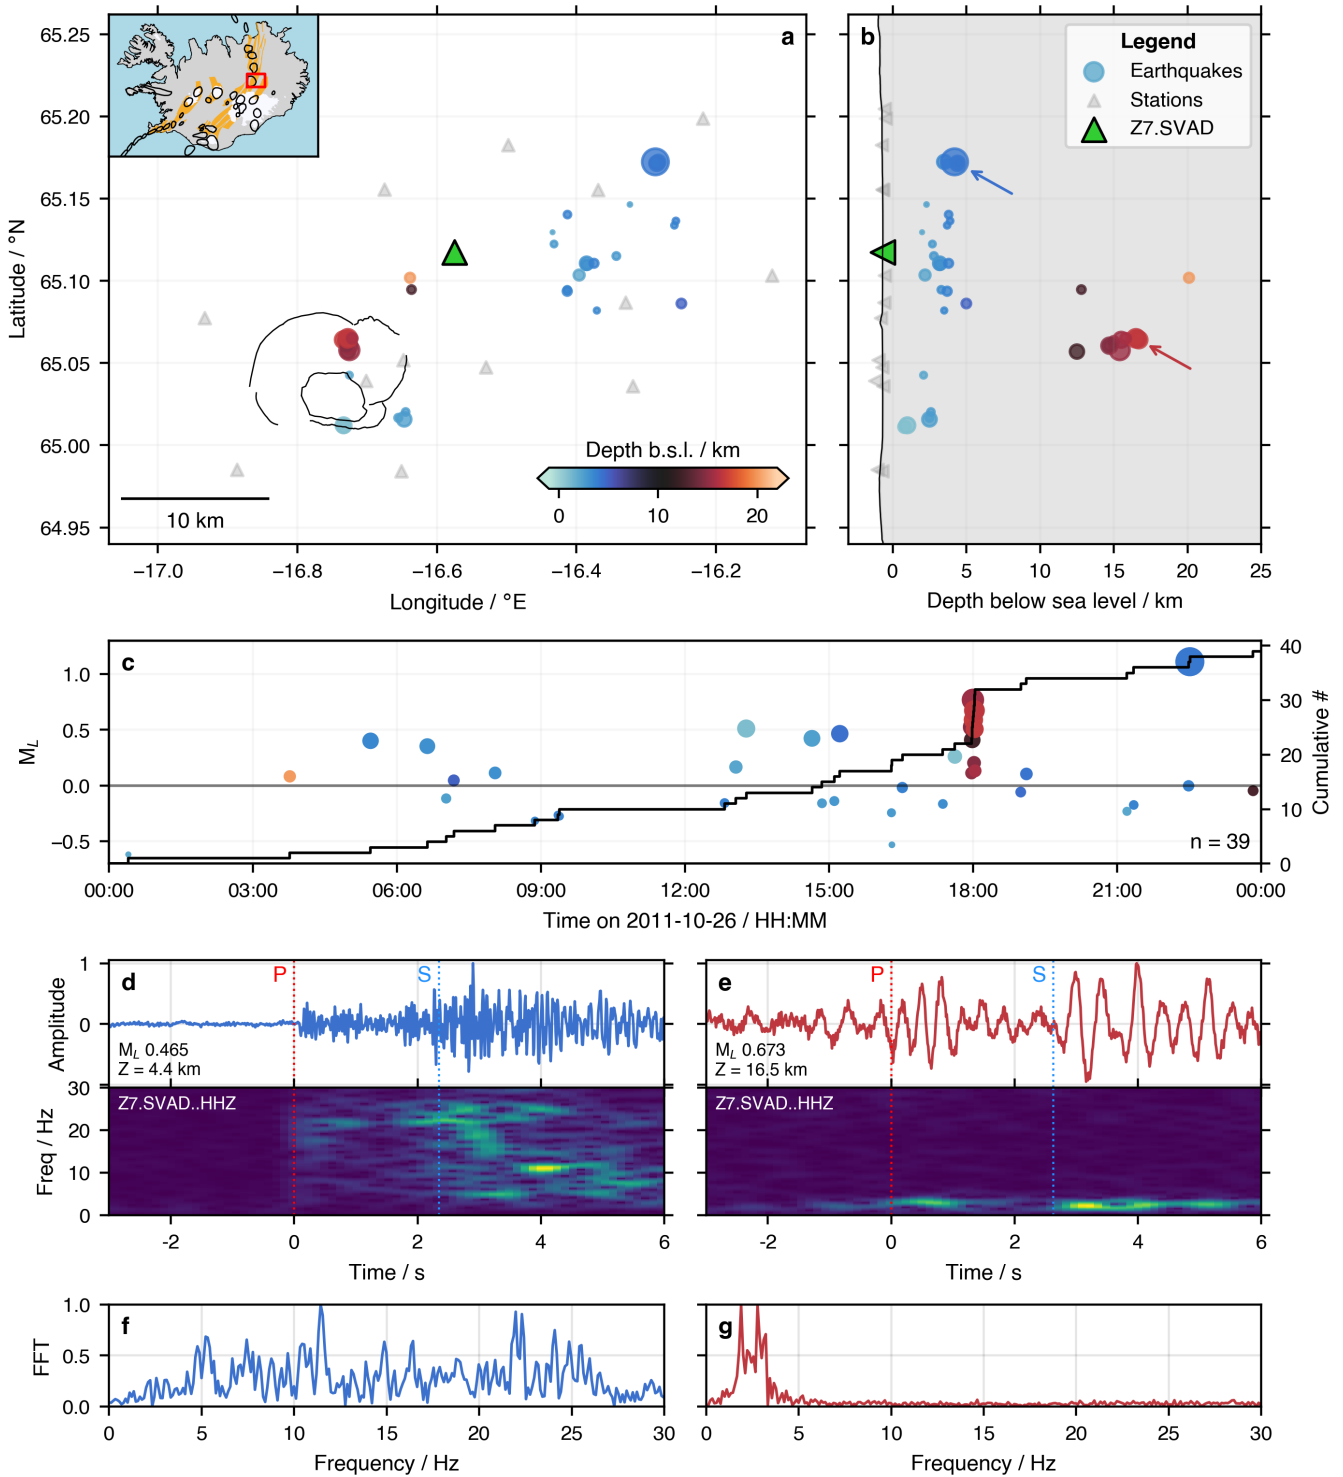


Figure 8 Seismicity around Askja volcano on 26th October 2011. **a**, **b** and **c** show map, cross-section, and timeline views of the QuakeMigrate automatic earthquake catalogue, with circles coloured by event depth and scaled by magnitude. Inset of **a** shows location within Iceland; lines show nested caldera boundaries of Askja volcano (Hjartardóttir et al., 2009), and triangles show stations used by QuakeMigrate. Line in **b** shows ground surface; plotted with no vertical exaggeration. **d**, **e**, **f** and **g** show waveforms, spectrograms, and power spectra for a shallow tectonic event (left, blue) and a deep long-period event (right, red)—indicated by arrows in panel **b**—as measured on the vertical component of station Z7.SVAD, which is indicated by the green triangle in panel **a**. Blue and red dotted lines show modelled P and S phase arrival times. Waveforms are bandpass filtered between 1.5–30 Hz to remove the oceanic microseism and high frequency environmental noise.

calculated automatically with QuakeMigrate.

Manually picked phase arrival times for shallow earthquakes in the region around Askja volcano are inverted for hypocentre locations using NonLinLoc—a

non-linear earthquake location algorithm provided as part of the NonLinLoc software package (Lomax et al., 2000). The velocity model and eikonal solver (Podvin and Lecomte, 1991) used to generate the lookup ta-

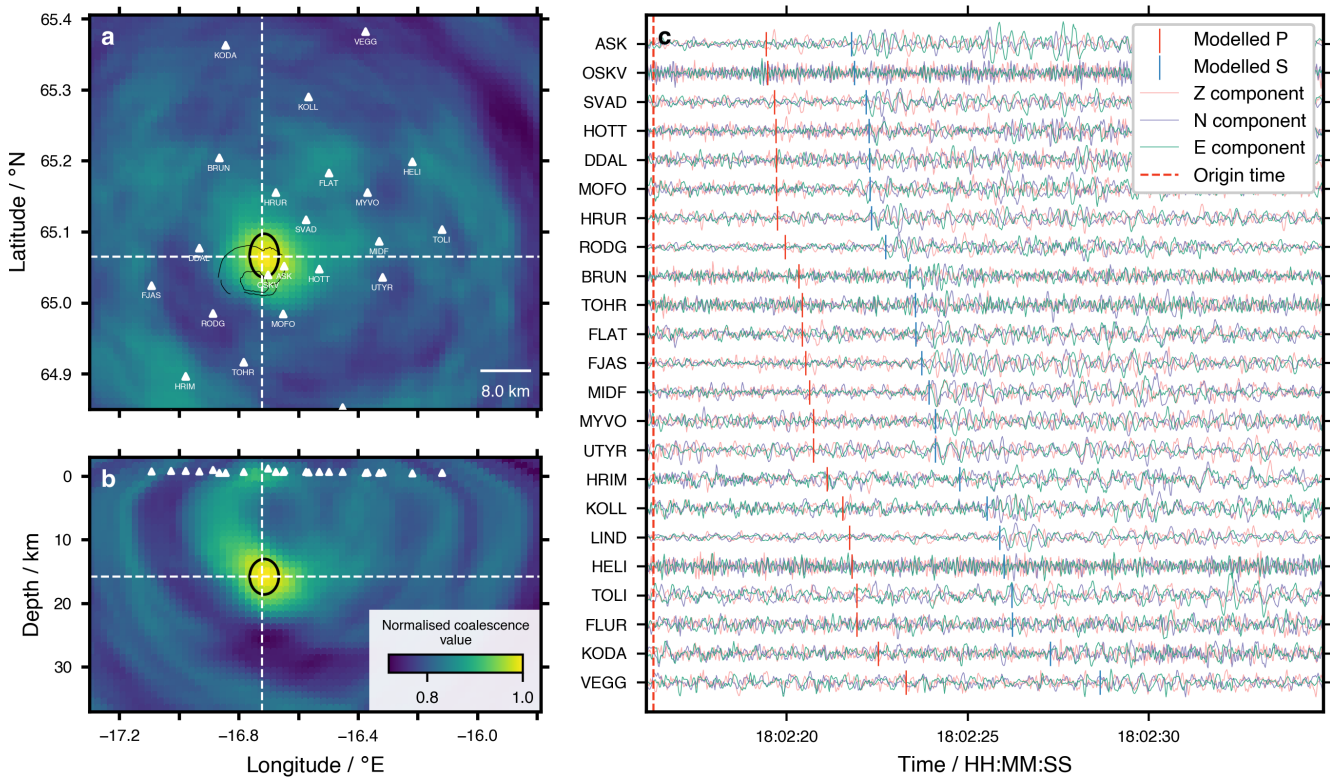


Figure 9 QuakeMigrate summary plot for a magnitude M_L 0.13 deep long-period microearthquake (EventID: 20111026180216660). **a** and **b** show the marginalised 3-D coalescence map in map and cross-section view, representing the event location uncertainty. White dashed lines show the maximum probability location (and the location of the cross-section); white triangles show station coordinates; black ellipse shows the gaussian approximation of the location uncertainty; black lines show the nested calderas of Askja volcano. Panel **c** shows bandpass-filtered waveforms (2–16 Hz for Z, and 2–14 Hz for N&E components), with modelled arrival times for P and S phases (red and blue ticks, respectively) calculated from the best-fit hypocentre location.

ble used by NonLinLoc are identical to those used by QuakeMigrate, facilitating a direct comparison between the two algorithms: the only differences are the input data (analyst traveltime picks vs. continuous waveforms) and location algorithm (NonLinLoc vs migration), as described in Winder (2022).

1,706 manually analysed earthquakes with magnitudes ranging from $-0.5 < M_L < 4.3$ were selected after filtering for the best located events (root-mean-square residual < 0.14 s, geometric mean location error < 1.5 km, at least 12 phase picks, and minimum 4 P- and 4 S-phase picks), and their origin times used as trigger times to calculate automatic locations with QuakeMigrate. The locations derived from these two independent analysis workflows are displayed for comparison in Figure 10. This forms a rigorous benchmark of the location performance of QuakeMigrate, using earthquakes with a wide range of magnitudes, located with varying network geometry, and crucially using identically generated traveltime lookup tables—the one common input required by the two techniques.

The locations for individual events are compared quantitatively in Figures 10e-g. The range of location differences along each axis is smaller than the mean location uncertainties reported for the manually refined earthquake locations (0.3 km in X and Y, and 0.7 km in depth). The mean (systematic) differences are also

small, particularly in epicentral location. This echoes the impression from visual comparison of the two sets of locations (10a-d), indicating an excellent agreement between the automatically calculated QuakeMigrate locations and the best-located hypocentres derived from manual picking of event phase arrivals. Comparison between the location differences and quoted location uncertainties for individual events shows that substantially more than 68% of event locations agree to within the 1σ uncertainties reported by both QuakeMigrate and NonLinLoc (Supplementary Figure S14). This further supports the interpretation that the small location differences can be explained by the inherent absolute uncertainty in individual event locations.

This comparison demonstrates that the locations calculated automatically using QuakeMigrate agree within uncertainty with a sample of nearly 2,000 manually refined earthquake locations in the region around Askja volcano. This indicates that for these events, QuakeMigrate, a fully automatic location algorithm, achieves at least the same level of location accuracy as this enormously more labour intensive manual analysis. The runtime for QuakeMigrate—including plotting an optional PDF summary for each event—was equivalent to ~ 12.5 s per event, using 4 threads for the migration step. On the same system, NonLinLoc ran slightly faster, equivalent to ~ 8.5 s per event. However, using

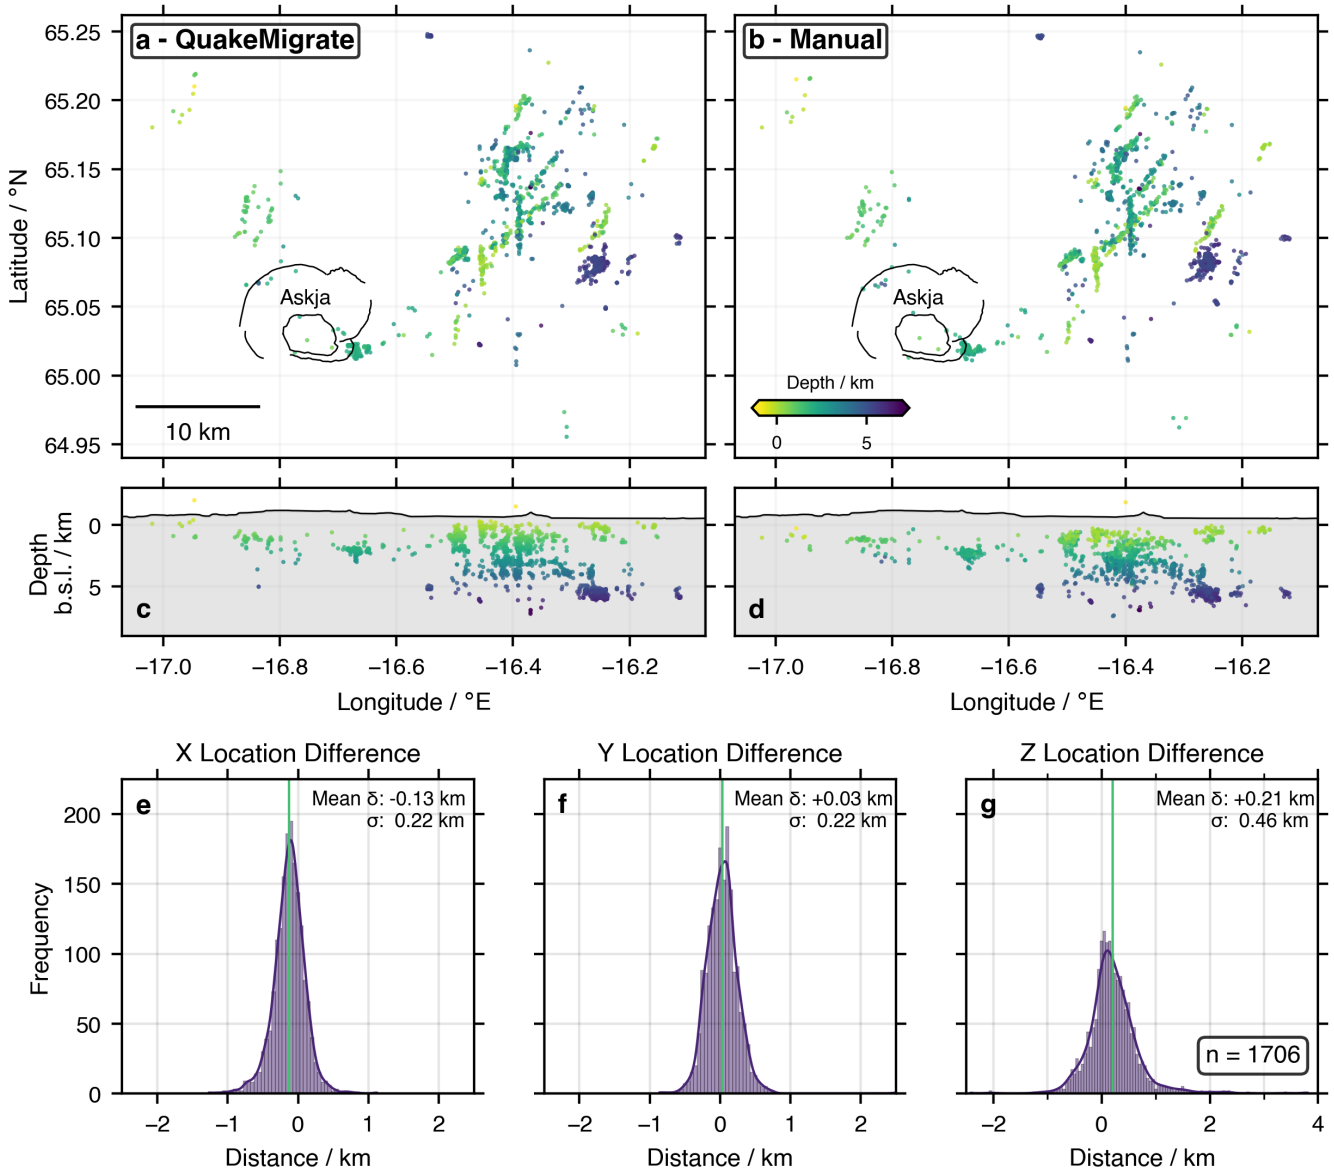


Figure 10 Location comparison for manually analysed events around Askja between 2007–2015. **a** and **b** show map views for QuakeMigrate and NonLinLoc locations, respectively, with events coloured by depth; map area is the same displayed in Figure 9. **c** and **d** show cross-section views. **e**, **f** and **g** show location differences between QuakeMigrate and NonLinLoc locations across 1,706 events.

a conservative estimate of 20 minutes per event for an analyst to annotate phase picks, the QuakeMigrate runtime represents a $\sim 100 \times$ speed-up.

5 Discussion

5.1 Alternative algorithms for onset function calculation

The default onset function algorithm included in the QuakeMigrate package is based on the short-term average to long-term average ratio (STA/LTA) of the seismogram amplitude. The core migration and stacking routine and the onset function module are implemented with a well-defined interface, however, enabling alternative onset functions to be incorporated as plug-ins and directly compared (Figure 2, Supplementary Figure S1). The STA/LTA algorithm is lightweight and re-

quires minimal pre-processing, provides good resolution and performance even for small earthquakes (e.g., Beskardes et al., 2018; Cesca and Grigoli, 2015), and, importantly, can be applied to any seismic dataset without, e.g., requiring pre-training a machine-learning model or the assembly of template events for matched-filter processing. Some degree of hyperparameter tuning is required to optimise performance for a given signal frequency content (Sections 4.1 and 4.2); however, the parameters are intuitively linked to signal properties (Hudson et al., 2019). Furthermore, it has been shown that the STA/LTA parameters used in Section 4.2 perform well across a broad range of tectonic and volcanic microseismicity, covering magnitudes ~ -2 to 5 (e.g., Greenfield et al., 2022; Supendi et al., 2023; Hudson et al., 2022; Yemane et al., 2025)—including locating both the DLP and shallow tectonic earthquakes present

in the Askja dataset shown here, which have significantly different frequency content (Figure 8).

Recently, significant progress has been made in developing a large family of phase detection algorithms based on machine-learning techniques (e.g., Ross et al., 2018; Mousavi et al., 2020; Zhu and Beroza, 2018). Unlike the STA/LTA algorithm, these can be designed to exploit information beyond the sharp increases in signal amplitude associated with primary seismic phase arrivals. This can help overcome the reduced sensitivity of STA/LTA where arrivals are emergent, for S-arrivals masked by the coda from earlier arriving P-waves, or for closely spaced events where the SNR is again reduced by the preceding coda (Schaff and Beroza, 2004). Furthermore, many of these techniques analyse 3-component data simultaneously, instead of using the simple mapping of vertical-component data to the P-wave onset function and horizontals to S described here (Section 3.4). This reduces the likelihood of errant peaks in the S-wave onset function caused by the presence of P-wave energy on the horizontal components (due to non-vertical incidence and/or scattering), and vice versa (e.g., Figure 3b). As a result, artefacts caused by incorrect phase identification are reduced—though this can also be mitigated by utilising polarisation information (e.g., Grigoli et al., 2014). It has also been shown that, where suitably trained, machine-learning phase detection algorithms can provide significantly improved robustness to data errors, including spikes caused by broken equipment, data gaps, and cultural noise which has significantly different time-frequency characteristics to earthquake phase arrivals, and avoid the tendency of STA/LTA algorithms to systematically pick slightly early (e.g., Mousavi et al., 2020).

Direct performance comparisons between STA/LTA and machine-learning methods typically focus on their application for phase picking, rather than onset function calculation for migration. This is a fundamentally different application, where observations are only counted if they exceed a specified “pick threshold”, whereas within the migration framework any peak that stacks coherently across the network contributes to the final detection/location (see Grigoli et al., 2018, for a more complete discussion). However, the general insight applies that—where appropriate hyperparameter choices are made, including the pick threshold—STA/LTA typically provides high recall (identifying most true earthquake phase arrivals) at the cost of lower precision when compared to the best-performing machine-learning methods, i.e., exhibiting more false positives (e.g., Mousavi et al., 2020). However, this comparative disadvantage is strongly mitigated by the robust and implicit phase association capability of the migration and stacking approach implemented within QuakeMigrate (Sections 2 and 3.5). Furthermore, the recall capability of machine-learning algorithms is dependent on both their architecture and the data used for their training. Though built to generalise, they typically face challenges when applied to datasets that differ strongly from the catalogues on which they were trained (e.g., Lapins et al., 2021; Shi et al., 2024; Ho et al., 2024; Mousavi et al., 2020). This is particularly relevant in

applications such as studying icequakes (Section 4.1), which represent an “out-of-distribution” generalisation, because most models are trained on datasets comprised dominantly of tectonic earthquakes (e.g., Ross et al., 2018; Zhu and Beroza, 2018; Mousavi et al., 2020; Shi et al., 2024) which have distinct characteristics. There are ongoing developments to address this, but so far these come at significantly increased computational cost (e.g., Shi et al., 2024), or require transfer learning on a labelled sample dataset (e.g., Lapins et al., 2021; Ho et al., 2024). Neither of these developments fully addresses the fact that, particularly in volcanic settings, new phases of unrest can produce previously unseen seismic signals that differ from existing earthquake templates (Lapins et al., 2021), which may therefore be missed by machine-learning or matched-filter algorithms, but still be successfully identified by simpler and less specific statistical algorithms such as STA/LTA, kurtosis, or the Hilbert transform (e.g., Shi et al., 2019). It is likely that combining the complementary strengths of different onset function algorithms—or systematically comparing their performance for a given dataset, to assess which best suits the user’s needs—will provide the best outcome. By allowing alternative onset functions to be incorporated as plug-ins in place of the default STA/LTA algorithm, QuakeMigrate provides a common framework to achieve this with both presently available onset function algorithms and those yet to be developed.

5.2 Computational cost

The exhaustive 4-D grid search executed in the Detect stage of a QuakeMigrate run is a continued focus for efficiency improvements, to enhance usability across a wide range of possible applications: from centimetre-scale acoustic-emission experiments to regional seismic arrays. A detailed overview of the computational cost for the Askja usage example (Section 4.2), and its scaling with parameter choice, is provided in the Supplementary Material (Section S3). Here we do not consider the Trigger stage (which generally runs in \ll 1 minute on a single thread for one day of data; Supplementary Figure S8) or the Locate stage, which is similar in runtime to alternatives such as NonLinLoc (Section 4.3), and significantly shorter than the Detect step other than where the event count is exceptionally high—and which is also required in a pick-then-locate workflow. For the Askja example, the optimal parameter choice for Detect results in a runtime of 38 seconds for one hour of data using 4 threads, or approximately 15 minutes for one full day of data, and \sim 92 hours (3.8 days) for one year. Computational cost is strongly dependent on the number of nodes in the search grid, which is controlled by the grid extent (in x, y, and z), node spacing (Supplementary Figure S7), as well as the sampling rate. The Askja example provides a representative use-case for a local seismic network, covering 71 x 62 x 40 km³. When decimated by a factor of 2 for Detect, this gives a node count of 36 x 32 x 21, or 24,192 grid nodes, with potential origin times sampled at 50 sps. Searching at the same sampling rate for a search grid spanning 150

$x 150 \times 40 \text{ km}^3$ would result in a Detect runtime of ~ 1 hour for a full day of data.

Where denser node spacing is required to achieve the level of resolution appropriate for very high frequency signals and dense network spacing, such as in the application to basal icequakes at Rutford ice stream (Section 4.1), computational cost is significantly higher, even for a much smaller grid extent. Here the LUT spans $9.2 \times 10.05 \times 1.5 \text{ km}^3$, but with a node spacing of 0.1 km for Detect, giving 150,288 nodes. Scanning at 250 sps, together this results in a compute time $\sim 30 \times$ higher. The memory requirement (which would otherwise scale similarly) is mitigated by completing the scan in shorter timesteps.

The icequake example therefore runs $\sim 3 \times$ faster than real-time, while the Askja example runs $\sim 100 \times$ faster. These runtimes can be compared to the now widely used family of machine-learning pickers, which generally run in a traditional pick, associate, then locate workflow. Where a suitable graphics processing unit (GPU) is available and a pre-trained model is used, phase identification and picking is relatively fast: of order 1–5 seconds per station per 24h of 3-component data (excluding GPD, which runs much more slowly), with most towards the bottom of that range (Münchmeyer et al., 2022; Lapins et al., 2021), or ~ 2.5 –12 hours for the 23 stations used in the Askja example, for 1 year of data. In this workflow, however, the association step must also be considered for a representative comparison to Detect, and—particularly where many events are detected—this can be a computationally intensive task. In the Askja region, 32 events/day were located between 2007–2020 using QuakeMigrate (Winder, 2022). This totals $\sim 12,000$ events per year. We refer to the runtime comparisons in Münchmeyer (2024)—which are indicative, rather than a full-scale benchmark of seismic phase associators—to provide indicative runtimes for this use case. For their shallow seismicity synthetic benchmark, with 100 events per day (\sim double that at Askja), runtimes range from 0.7–13.0 s / day, or 4–80 minutes for 1 year of data. For higher event rates, performance is significantly worse; the maximum event rate tested was 2,000 / day—significantly below the 10,000 events per day in even the filtered Rutford icequake catalogue (Section 4.2)—and in this case several associators did not complete the task within 48 hours for one day of data, though PyOcto and REAL ran within 0.3–20 minutes, depending on the number of mis-picks on noise (Münchmeyer, 2024). Indicative overall runtimes are therefore ~ 2.5 –14 hours for the machine-learning pipelines, which is a significant factor faster than QuakeMigrate Detect, at ~ 92 hours.

It is important to put these runtimes in context of typical research workflows, where differences of hours to several tens of hours between catalogue generation jobs are generally not prohibitive, and to acknowledge that it is common for users to have access to significantly more CPU resources than used here, which can be exploited through parallelisation for considerable speed-ups. Notably, the association benchmarks from Münchmeyer (2024) were run on 16 threads, while

the QuakeMigrate runtimes are from processing on 4 threads only. Speed-ups for the migration do not continue to scale linearly above this due to I/O and other overheads (Supplementary Figure S7), but running multiple Detect jobs in parallel—e.g., by splitting a year of data into several batches—means one can continue to achieve near-linear speed-up, reducing the runtimes by a further factor of 4, and reaching similar computational cost to the machine-learning pipeline. On a server or cluster with 100 CPUs or more, QuakeMigrate would run significantly faster. This, too, is without requiring access to a machine with a dedicated GPU. In the current implementation, for dense earthquake sequences—at small spatial scale, and with short inter-event intervals—migration is an equivalent or more efficient option, however in sparse cases—with low spatial density of events, and/or long inter-event times—the pick-then-locate workflow strongly benefits from earlier data reduction, whereas retaining continuous onset functions and combining them in the migration search uses significant computational resources evaluating non-hypocentres. These two very different approaches are therefore complementary depending on the user’s priorities and the application, and particularly in cases where a suitable pre-trained model is not available (Sections 4.1 and 5.1). Furthermore, there remain several opportunities to improve the efficiency of QuakeMigrate, including by sub-sampling the search grid with an Oct-Tree algorithm (Drew et al., 2013; Isken et al., 2025), implementing a boxcar filter to enable more aggressive decimation of the search grid without compromising detection performance (Drew et al., 2013), and decoupling the scan sampling rate from the onset function sampling rate, to search over possible origin times at a coarser interval (Shi et al., 2019). Together, these approaches could realistically reduce runtime by an order of magnitude.

5.3 Incorporation into automated workflows

The automation, speed-up, and demonstrated levels of detection capability and location accuracy are crucially important in the context of automated workflows where catalogue creation represents the foundational step. One limitation apparent in the QuakeMigrate locations (Figure 10a,c) is the remnant imprint of the underlying search grid, which is absent from the NonLinLoc results (Figure 10b,d) thanks to its implementation of an Oct-Tree sub-sampling approach to reduce the node spacing in the search grid (Lomax et al., 2000). This represents a clear avenue for future improvement (as implemented by, e.g., Isken et al., 2025), but where the QuakeMigrate absolute locations are to be refined by relative re-location (e.g., using GrowClust, Trugman and Shearer, 2017), it is their demonstrated accuracy that is important, as the relative location refinement will dramatically improve the precision even compared to the initial NonLinLoc locations (e.g., Winder, 2022). Further automated catalogue enhancement may be achieved through match-filter analysis (e.g., Chamberlain et al., 2017), but though attempts have been made to generate synthetic template events, the best results are achieved

where real events are used as a starting point (Chamberlain and Townend, 2018). This underlines the importance of generating an initial catalogue with the best possible detection performance across a variety of different event types (Figure 8), which is particularly relevant to volcanic environments (e.g., Thelen et al., 2022; Lapins et al., 2021). The ability to successfully filter out artefacts without removing real, small events which—if part of a significantly different cluster of seismicity—may otherwise not be detected by matched-filter routines, is also crucial (Section 4.1).

5.4 Ongoing and future developments in seismic data acquisition

In recent years, new technologies have been introduced which allow far denser sampling of the seismic wavefield. These include low-cost and highly portable geophones (seismic ‘nodes’) which may be used in stand-alone temporary arrays, or to densify relatively sparse permanent networks (e.g., Hudson et al., 2024), and the use of fibre-optic cables as distributed strain sensors (commonly referred to as Distributed Acoustic Sensing, or DAS) (e.g., Hudson et al., 2025). These advancements favour the network-based approach to earthquake detection and location outlined here, because the challenge of extracting arrival time picks and phase association is exacerbated by the increase to thousands of channels, yet these additional channels also enhance the benefit of exploiting coherency information from nearby receivers. The modular architecture of the QuakeMigrate package will allow for development, adoption and benchmarking of novel pre-processing and onset function algorithms designed to best harness these new data types (Section 5.1), and accommodate new challenges related to altered noise characteristics and sensitivity to different seismic phases (Hudson et al., 2025).

6 Conclusions

QuakeMigrate provides a powerful and efficient method for building an earthquake catalogue from continuous seismic data. It is distributed as a cross-platform, open-source Python package, making it widely accessible for application to increasingly large quantities of continuous waveform data collected using dense (local) seismic networks.

Migration-based approaches to seismic event detection and location promise significant advantages in robustness to noise, detection capability amongst rapidly occurring swarms of microearthquakes, and in the ability to detect events with phase arrivals close to or below the signal-to-noise ratio at individual stations. However, computational cost remains a challenge. Here, through separating the detection, triggering and location steps, significant improvements are made in both the efficiency and adaptability of the method.

QuakeMigrate is designed with a modular architecture, providing transparency to the user, and the flexibility to adapt to the specific requirements and challenges of individual datasets. This also provides the ex-

citing opportunity to rapidly create and rigorously compare and benchmark new approaches and techniques within a robust framework. For example, the plug-in nature of the Onset module allows direct benchmarking of existing algorithms, as well as the prospect of harnessing new approaches. This includes exploiting the potential to use the continuous output from machine-learning algorithms for phase arrival detection.

Locate outputs a suite of statistics alongside each event location. These include metrics that are explicitly designed to distinguish real events from artefacts, as opposed to describing detections under the implicit assumption that they are true earthquakes. This promises a significant improvement in the capability to quality-control a preliminary catalogue.

The performance of QuakeMigrate is demonstrated with application to an example dataset from the Rutford Ice Stream, Antarctica, where we achieve a two orders of magnitude increase in event detections compared to previous studies. The locations attained are tightly constrained about the ice-bed interface, corresponding to the expected distribution from basal stick-slip seismicity, and underlining the excellent location performance.

Application to data from Askja volcano, Iceland, demonstrates the capability of QuakeMigrate to successfully detect and locate small earthquakes with a range of characteristics, including deep long-period earthquakes with emergent waveforms that are difficult to identify on recordings from individual stations. This is achieved while maintaining excellent detection and location performance for more typical, high-frequency, tectonic earthquakes in the shallow crust. Quantitative comparison with locations derived from manual phase picking demonstrates that QuakeMigrate achieves the same level of accuracy as this widely held “gold standard” in microseismic analysis.

Together, these example use cases demonstrate that QuakeMigrate is capable of providing the detection, location and filtering performance required to provide the best possible foundation for further automated processing, for example with template matching and/or relative relocation. As seismic networks become ever more dense, the performance advantage of migration-based methods like QuakeMigrate will only continue to improve. With further improvements in efficiency, we hope that such methods will become a standard tool for earthquake detection and location.

Acknowledgements

We would like to thank all those who have been involved with the development of QuakeMigrate. In particular, Miriam Reiss, Tim Greenfield, Miriam Gauntlett, and Ian Lee put earlier versions of this package through its paces and reported back to us any and all bugs, issues, and suggestions for improvement. This software builds on the theory and original implementation devised by Julian Drew, whose Python adaptation of GMM sparked the idea for this project. The authors also thank him for help and advice in getting to grips with the concepts that underpin this approach to earthquake loca-

tion, and the subtleties involved in its application. We are grateful to Thomas Lecocq and Corentin Caudron for enthusiastic and encouraging chats along the way, and for hosting a visit to Brussels during which many of the figures in this article were drafted. Fieldwork to collect the waveform data shown in Section 4 were supported by NERC awards NE/B502287/1, NE/F011407/1, and NE/H025006/1 and instrument loans 842, 857, 914, 968, 980 and 1022 from SEIS-UK. We thank the Icelandic Meteorological Office for sharing waveform data from 9 of their stations which were incorporated in the manual-pick location benchmark presented in Section 4.3. We are also deeply grateful for the efforts and contributions of all those who maintain the free and open-source packages on which QuakeMigrate relies. Figures were created using Matplotlib (v3.9.2 Hunter, 2007), seaborn (v0.13.2 Waskom, 2021) and GMT (v6.5.0 Wessel et al., 2019). We thank the editor Hongyu Sun and two anonymous reviewers for their feedback, which has considerably improved the manuscript.

Data and code availability

All scripts and data required to reproduce the figures in this paper are available at <https://github.com/QuakeMigrate/manuscript> and archived in Bacon and Winder (2025). Continuous waveform data used in Sections 4.1 and 4.2 are openly available via the EarthScope Data Management Center (previously IRIS DMC) under network codes YG_2009, 4F_2007 (White, 2007), and Z7_2010 (White, 2010) and cut waveforms and manually picked phase arrival times used in Section 4.3 are archived in Winder et al. (2025). An example of how these data can be downloaded is available in the GitHub repository that accompanies this article. The source code for QuakeMigrate is hosted on GitHub (<https://github.com/QuakeMigrate/QuakeMigrate>), with version snapshots archived under Winder et al. (2021). Figures were prepared using QuakeMigrate v1.2.2 (Bacon et al., 2025) and NonLinLoc v7.1.04 (05Sep2024, Lomax et al., 2023).

Competing interests

The authors declare no competing interests.

References

- Ágústsdóttir, T., Winder, T., Woods, J., White, R. S., Greenfield, T., and Brandsdóttir, B. Intense Seismicity During the 2014–2015 Bárðarbunga–Holuhraun Rifting Event, Iceland, Reveals the Nature of Dike-Induced Earthquakes and Caldera Collapse Mechanisms. *Journal of Geophysical Research: Solid Earth*, 124(8): 8331–8357, 2019. doi: 10.1029/2018JB016010.
- Arrowsmith, S. J., Trugman, D. T., MacCarthy, J., Bergen, K. J., Lumley, D., and Magnani, M. B. Big Data Seismology. *Reviews of Geophysics*, 60(2):e2021RG000769, 2022. doi: 10.1029/2021RG000769.
- Bacon, C. *Seismic anisotropy and microseismicity: from crustal formation to subduction termination*. PhD thesis, University of Cambridge, 2022. doi: 10.17863/CAM.82196.
- Bacon, C. and Winder, T. QuakeMigrate/manuscript: Initial release of analysis and visualisation code accompanying the QuakeMigrate manuscript. Zenodo, 2025. doi: 10.5281/zenodo.16419323.
- Bacon, C., Winder, T., Greenfield, T., and Bot, S. QuakeMigrate/QuakeMigrate: QuakeMigrate v1.2.2, 2025. doi: 10.5281/zenodo.17882838.
- Baker, T., Granat, R., and Clayton, R. W. Real-time Earthquake Location Using Kirchhoff Reconstruction. *Bulletin of the Seismological Society of America*, 95(2):699–707, 2005. doi: 10.1785/0120040123.
- Beaucé, E., Frank, W. B., Seydoux, L., Poli, P., Groebner, N., van der Hilst, R. D., and Campillo, M. BPF: A Backprojection and Matched-Filtering Workflow for Automated Earthquake Detection and Location. *Seismological Research Letters*, 95(2A): 1030–1042, 2023. doi: 10.1785/0220230230.
- Besekardes, G. D., Hole, J. A., Wang, K., Michaelides, M., Wu, Q., Chapman, M. C., Davenport, K. K., Brown, L. D., and Quiros, D. A. A comparison of earthquake backprojection imaging methods for dense local arrays. *Geophysical Journal International*, 212(3):1986–2002, 2018. doi: 10.1093/gji/ggx520.
- Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., and Wassermann, J. ObsPy: A Python Toolbox for Seismology. *Seismological Research Letters*, 81(3):530–533, 2010. doi: 10.1785/gssrl.81.3.530.
- Cesca, S. and Grigoli, F. Chapter Two - Full Waveform Seismological Advances for Microseismic Monitoring. In Dmowska, R., editor, *Advances in Geophysics*, volume 56, pages 169–228. Elsevier, 2015. doi: 10.1016/bs.agph.2014.12.002.
- Chamberlain, C. J. and Townend, J. Detecting Real Earthquakes Using Artificial Earthquakes: On the Use of Synthetic Waveforms in Matched-Filter Earthquake Detection. *Geophysical Research Letters*, 45(21):11,641–11,649, 2018. doi: 10.1029/2018GL079872.
- Chamberlain, C. J., Hopp, C. J., Boese, C. M., Warren-Smith, E., Chambers, D., Chu, S. X., Michailos, K., and Townend, J. EQ-corrscan: Repeating and Near-Repeating Earthquake Detection and Analysis in Python. *Seismological Research Letters*, 89(1): 173–181, 2017. doi: 10.1785/0220170151.
- Drew, J. *Coalescence microseismic mapping: an imaging method for the detection and location of seismic events*. PhD thesis, University of Cambridge, 2010.
- Drew, J., White, R. S., Tilmann, F., and Tarasewicz, J. Coalescence microseismic mapping. *Geophysical Journal International*, 195(3):1773–1785, 2013. doi: 10.1093/gji/ggt331.
- Evenden, G. I., Rouault, E., Warmerdam, F., Evers, K., Knudsen, T., Butler, H., Taves, M. W., Schwehr, K., Sales de Andrade, E., Karney, C., Couwenberg, S., Dawson, N., Snow, A. D., and Jimenez Shaw, J. PROJ. Zenodo, 2025. doi: 10.5281/zenodo.15030996.
- Greenfield, T. and White, R. S. Building Icelandic igneous crust by repeated melt injections. *Journal of Geophysical Research: Solid Earth*, 120(11):7771–7788, 2015. doi: 10.1002/2015JB012009.
- Greenfield, T., White, R. S., and Roecker, S. The magmatic plumbing system of the Askja central volcano, Iceland, as imaged by seismic tomography. *Journal of Geophysical Research: Solid Earth*, 121(10):7211–7229, 2016. doi: 10.1002/2016JB013163.
- Greenfield, T., White, R. S., Winder, T., and Ágústsdóttir, T. Seismicity of the Askja and Bárðarbunga volcanic systems of Iceland, 2009–2015. *Journal of Volcanology and Geothermal Research*, 391:106432, 2020. doi: 10.1016/j.jvolgeores.2018.08.010.
- Greenfield, T., Winder, T., Rawlinson, N., Maclennan, J., White, R. S., Ágústsdóttir, T., Bacon, C. A., Brandsdóttir, B., Eibl, E. P.,

- Glastonbury-Southern, E., et al. Deep long period seismicity preceding and during the 2021 Fagradalsfjall eruption, Iceland. *Bulletin of Volcanology*, 84(12):101, 2022. doi: 10.1007/s00445-022-01603-2.
- Grigoli, F., Cesca, S., Vassallo, M., and Dahm, T. Automated Seismic Event Location by Travel-Time Stacking: An Application to Mining Induced Seismicity. *Seismological Research Letters*, 84(4):666–677, 2013. doi: 10.1785/0220120191.
- Grigoli, F., Cesca, S., Amoroso, O., Emolo, A., Zollo, A., and Dahm, T. Automated seismic event location by waveform coherence analysis. *Geophysical Journal International*, 196(3):1742–1753, 2014. doi: 10.1093/gji/ggt477.
- Grigoli, F., Cesca, S., Krieger, L., Kriegerowski, M., Gammaldi, S., Horalek, J., Priolo, E., and Dahm, T. Automated microseismic event location using Master-Event Waveform Stacking. *Scientific Reports*, 6(1):25744, 2016. doi: 10.1038/srep25744.
- Grigoli, F., Cesca, S., Priolo, E., Rinaldi, A. P., Clinton, J. F., Stabile, T. A., Dost, B., Fernandez, M. G., Wiemer, S., and Dahm, T. Current challenges in monitoring, discrimination, and management of induced seismicity related to underground industrial activities: A European perspective. *Reviews of Geophysics*, 55(2):310–340, 2017. doi: 10.1002/2016RG000542.
- Grigoli, F., Scarabello, L., Böse, M., Weber, B., Wiemer, S., and Clinton, J. F. Pick- and waveform-based techniques for real-time detection of induced seismicity. *Geophysical Journal International*, 213(2):868–884, 2018. doi: 10.1093/gji/ggy019.
- Hansen, S. M. and Schmandt, B. Automated detection and location of microseismicity at Mount St. Helens with a large-N geophone array. *Geophysical Research Letters*, 42(18):7390–7397, 2015. doi: 10.1002/2015GL064848.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020. doi: 10.1038/s41586-020-2649-2.
- Hjartardóttir, Á. R., Einarsson, P., and Sigurdsson, H. The fissure swarm of the Askja volcanic system along the divergent plate boundary of N Iceland. *Bulletin of Volcanology*, 71(9):961–975, 2009. doi: 10.1007/s00445-009-0282-x.
- Ho, L. M., Walter, J. I., Hansen, S. E., Sánchez-Roldán, J. L., and Peng, Z. Evaluating Automated Seismic Event Detection Approaches: An Application to Victoria Land, East Antarctica. *Journal of Geophysical Research: Machine Learning and Computation*, 1(3):e2024JH000185, 2024. doi: 10.1029/2024JH000185.
- Hudson, T. S. *Investigating Volcanic and Glacial Processes Using Microseismicity*. PhD thesis, University of Cambridge, 2019. doi: 10.17863/CAM.45965.
- Hudson, T. S., Smith, J., Brisbourne, A. M., and White, R. S. Automated detection of basal icequakes and discrimination from surface crevassing. *Annals of Glaciology*, 60(79):167–181, 2019. doi: 10.1017/aog.2019.18.
- Hudson, T. S., Brisbourne, A. M., Walter, F., Gräff, D., White, R. S., and Smith, A. M. Icequake Source Mechanisms for Studying Glacial Sliding. *Journal of Geophysical Research: Earth Surface*, 125(11):e2020JF005627, 2020. doi: 10.1029/2020JF005627.
- Hudson, T. S., Kendall, J.-M., Pritchard, M. E., Blundy, J. D., and Gottsmann, J. H. From slab to surface: Earthquake evidence for fluid migration at Uturuncu volcano, Bolivia. *Earth and Planetary Science Letters*, 577:117268, 2022. doi: 10.1016/j.epsl.2021.117268.
- Hudson, T. S., Kettlety, T., Kendall, J.-M., O’Toole, T., Jupe, A., Shail, R. K., and Grand, A. Seismic Node Arrays for Enhanced Understanding and Monitoring of Geothermal Systems. *The Seismic Record*, 4(3):161–171, 2024. doi: 10.1785/0320240019.
- Hudson, T. S., Klaasen, S., Fontaine, O., Bacon, C. A., Jónsdóttir, K., and Fichtner, A. Towards a widely applicable earthquake detection algorithm for fibreoptic and hybrid fibreoptic-seismometer networks. *Geophysical Journal International*, 240(3):1965–1985, 2025. doi: 10.1093/gji/ggae459.
- Hunter, J. D. Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering*, 9(03):90–95, 2007. doi: 10.1109/MCSE.2007.55.
- Isken, M., Niemz, P., Münchmeyer, J., Büyükakpınar, P., Heimann, S., Cesca, S., Vasyura-Bathke, H., and Dahm, T. Qseek: A data-driven Framework for Automated Earthquake Detection, Localization and Characterization. *Seismica*, 4(1), 2025. doi: 10.26443/seismica.v4i1.1283.
- Kao, H. and Shan, S.-J. The Source-Scanning Algorithm: mapping the distribution of seismic sources in time and space. *Geophysical Journal International*, 157(2):589–594, 2004. doi: 10.1111/j.1365-246X.2004.02276.x.
- Kao, H., Shan, S.-J., Dragert, H., Rogers, G., Cassidy, J. F., Wang, K., James, T. S., and Ramachandran, K. Spatial-temporal patterns of seismic tremors in northern Cascadia. *Journal of Geophysical Research: Solid Earth*, 111(B3), 2006. doi: 10.1029/2005JB003727.
- Keir, D., Stuart, G. W., Jackson, A., and Ayele, A. Local Earthquake Magnitude Scale and Seismicity Rate for the Ethiopian Rift. *Bulletin of the Seismological Society of America*, 96(6):2221–2230, 2006. doi: 10.1785/0120060051.
- Key, J., White, R. S., Soosalu, H., and Jakobsdóttir, S. S. Correction to “Multiple melt injection along a spreading segment at Askja, Iceland”. *Geophysical Research Letters*, 38(10), 2011a. doi: 10.1029/2011GL047491.
- Key, J., White, R. S., Soosalu, H., and Jakobsdóttir, S. S. Multiple melt injection along a spreading segment at Askja, Iceland. *Geophysical Research Letters*, 38(5), 2011b. doi: 10.1029/2010GL046264.
- King, E. C., Pritchard, H. D., and Smith, A. M. Subglacial landforms beneath Rutford Ice Stream, Antarctica: detailed bed topography from ice-penetrating radar. *Earth System Science Data*, 8(1):151–158, 2016. doi: 10.5194/essd-8-151-2016.
- Kufner, S.-K., Brisbourne, A. M., Smith, A. M., Hudson, T. S., Murray, T., Schlegel, R., Kendall, J. M., Anandakrishnan, S., and Lee, I. Not all Icequakes are Created Equal: Basal Icequakes Suggest Diverse Bed Deformation Mechanisms at Rutford Ice Stream, West Antarctica. *Journal of Geophysical Research: Earth Surface*, 126(3):e2020JF006001, 2021. doi: 10.1029/2020JF006001.
- Langet, N., Maggi, A., Michelini, A., and Brenguier, F. Continuous Kurtosis-Based Migration for Seismic Event Detection and Location, with Application to Piton de la Fournaise Volcano, La Réunion. *Bulletin of the Seismological Society of America*, 104(1):229–246, 2014. doi: 10.1785/0120130107.
- Lapins, S., Goitom, B., Kendall, J.-M., Werner, M. J., Cashman, K. V., and Hammond, J. O. S. A Little Data Goes a Long Way: Automating Seismic Phase Arrival Picking at Nabro Volcano With Transfer Learning. *Journal of Geophysical Research: Solid Earth*, 126(7):e2021JB021910, 2021. doi: 10.1029/2021JB021910.
- Larmat, C., Montagner, J.-P., Fink, M., Capdeville, Y., Tourin, A., and Clévédy, E. Time-reversal imaging of seismic sources and application to the great Sumatra earthquake. *Geophysical Research Letters*, 33(19), 2006. doi: 10.1029/2006GL026336.
- Li, L., Tan, J., Schwarz, B., Staněk, F., Poiata, N., Shi, P., Diekmann, L., Eisner, L., and Gajewski, D. Recent Advances and Challenges of Waveform-Based Seismic Location Methods at Multi-

- ple Scales. *Reviews of Geophysics*, 58(1):e2019RG000667, 2020. doi: 10.1029/2019RG000667.
- Lomax, A., Virieux, J., Volant, P., and Berge-Thierry, C. Probabilistic Earthquake Location in 3D and Layered Models. In Thurber, C. H. and Rabinowitz, N., editors, *Advances in Seismic Event Location*, Modern Approaches in Geophysics, pages 101–134. Springer Netherlands, Dordrecht, 2000. doi: 10.1007/978-94-015-9536-0_5.
- Lomax, A., Megies, T., Ho, S., SAUREL, J.-M., and luca-s. `aloxmax/NonLinLoc`: Intermittent release to enable citing and reference of NonLinLoc through Zenodo. Zenodo, 2023. doi: 10.5281/zenodo.8046170.
- McKinney, W. Data Structures for Statistical Computing in Python. pages 56–61, Austin, Texas, 2010. doi: 10.25080/Majora-92bf1922-00a.
- Mousavi, S. M., Ellsworth, W. L., Zhu, W., Chuang, L. Y., and Beroza, G. C. Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking. *Nature Communications*, 11(1):3952, 2020. doi: 10.1038/s41467-020-17591-w.
- Münchmeyer, J. PyOcto: A high-throughput seismic phase associator. *Seismica*, 3(1), 2024. doi: 10.26443/seismica.v3i1.1130.
- Münchmeyer, J., Woollam, J., Rietbrock, A., Tilmann, F., Lange, D., Bornstein, T., Diehl, T., Giunchi, C., Haslinger, F., Jozinović, D., Michelini, A., Saul, J., and Soto, H. Which Picker Fits My Data? A Quantitative Evaluation of Deep Learning Based Seismic Pickers. *Journal of Geophysical Research: Solid Earth*, 127(1):e2021JB023499, 2022. doi: 10.1029/2021JB023499.
- Niemz, P., Cesca, S., Heimann, S., Grigoli, F., von Specht, S., Hammer, C., Zang, A., and Dahm, T. Full-waveform-based characterization of acoustic emission activity in a mine-scale experiment: a comparison of conventional and advanced hydraulic fracturing schemes. *Geophysical Journal International*, 222(1): 189–206, 2020. doi: 10.1093/gji/ggaa127.
- pandas development team, T. `pandas-dev/pandas`: Pandas, 2026. doi: 10.5281/zenodo.19340003.
- Pesicek, J. D., Child, D., Artman, B., and Cieřlik, K. Picking versus stacking in a modern microearthquake location: Comparison of results from a surface passive seismic monitoring array in Oklahoma. *Geophysics*, 79(6):KS61–KS68, 2014. doi: 10.1190/geo2013-0404.1.
- Podvin, P. and Lecomte, I. Finite difference computation of traveltimes in very contrasted velocity models: a massively parallel approach and its associated tools. *Geophysical Journal International*, 105(1):271–284, 1991. doi: 10.1111/j.1365-246X.1991.tb03461.x.
- Poiata, N., Satriano, C., Vilotte, J.-P., Bernard, P., and Obara, K. Multiband array detection and location of seismic sources recorded by dense seismic networks. *Geophysical Journal International*, 205(3):1548–1573, 2016. doi: 10.1093/gji/ggw071.
- Ross, Z. E., Meier, M., Hauksson, E., and Heaton, T. H. Generalized Seismic Phase Detection with Deep Learning. *Bulletin of the Seismological Society of America*, 108(5A):2894–2901, 2018. doi: 10.1785/0120180080.
- Ross, Z. E., Yue, Y., Meier, M.-A., Hauksson, E., and Heaton, T. H. PhaseLink: A Deep Learning Approach to Seismic Phase Association. *Journal of Geophysical Research: Solid Earth*, 124(1): 856–869, 2019. doi: 10.1029/2018JB016674.
- Schaff, D. P. and Beroza, G. C. Coseismic and postseismic velocity changes measured by repeating earthquakes. *Journal of Geophysical Research: Solid Earth*, 109(B10), 2004. doi: 10.1029/2004JB003011.
- Shi, P., Angus, D., Rost, S., Nowacki, A., and Yuan, S. Automated seismic waveform location using multichannel coherency migration (MCM)–I: theory. *Geophysical Journal International*, 216(3):1842–1866, 2019. doi: 10.1093/gji/ggy132.
- Shi, P., Grigoli, F., Lanza, F., Beroza, G. C., Scarabello, L., and Wiemer, S. MALMI: An Automated Earthquake Detection and Location Workflow Based on Machine Learning and Waveform Migration. *Seismological Research Letters*, 93(5):2467–2483, 2022. doi: 10.1785/0220220071.
- Shi, P., Meier, M.-A., Villiger, L., Tuinstra, K., Selvadurai, P. A., Lanza, F., Yuan, S., Obermann, A., Mesimeri, M., Münchmeyer, J., Bianchi, P., and Wiemer, S. From Labquakes to Megathrusts: Scaling Deep Learning Based Pickers Over 15 Orders of Magnitude. *Journal of Geophysical Research: Machine Learning and Computation*, 1(4):e2024JH000220, 2024. doi: 10.1029/2024JH000220.
- Smith, E. C., Smith, A. M., White, R. S., Brisbourne, A. M., and Pritchard, H. D. Mapping the ice-bed interface characteristics of Rutford Ice Stream, West Antarctica, using microseismicity. *Journal of Geophysical Research: Earth Surface*, 120(9): 1881–1894, 2015. doi: 10.1002/2015JF003587.
- Smith, J. D., White, R. S., Avouac, J.-P., and Bourne, S. Probabilistic earthquake locations of induced seismicity in the Groningen region, the Netherlands. *Geophysical Journal International*, 222(1):507–516, 2020. doi: 10.1093/gji/ggaa179.
- Snow, A. D., Cochran, M., Hoese, D., Miara, I., den Bossche, J. V., Mayo, C., Lucas, G., Cochran, P., de Kloe, J., Karney, C., Shaw, J. J., Anh, T. Q., Filipe, Ouzounoudis, G., Couwenberg, B., Lostis, G., Dearing, J., Jurd, B., Gohlke, C., Schneck, C., McDonald, D., Taves, M., Itkin, M., May, R., Stewart, A. J., de Bittencourt, H., Little, B., Hugonnet, R., and Rahul, P. S. `pyproj4/pyproj`: 3.7.2 Release. Zenodo, 2025. doi: 10.5281/zenodo.16874869.
- Soosalu, H., Jónsdóttir, K., and Einarsson, P. Seismicity crisis at the Katla volcano, Iceland—signs of a cryptodome? *Journal of Volcanology and Geothermal Research*, 153(3):177–186, 2006. doi: 10.1016/j.jvolgeores.2005.10.013.
- Supendi, P., Winder, T., Rawlinson, N., Bacon, C. A., Palgunadi, K. H., Simanjuntak, A., Kurniawan, A., Widiyantoro, S., Nugraha, A. D., Shiddiqi, H. A., Ardianto, Daryono, Adi, S. P., Karnawati, D., Priyobudi, Marliyani, G. I., Imran, I., and Jatnika, J. A conjugate fault revealed by the destructive Mw 5.6 (November 21, 2022) Cianjur earthquake, West Java, Indonesia. *Journal of Asian Earth Sciences*, 257:105830, 2023. doi: 10.1016/j.jseaes.2023.105830.
- Thelen, W. A., Matoza, R. S., and Hotovec-Ellis, A. J. Trends in volcano seismology: 2010 to 2020 and beyond. *Bulletin of Volcanology*, 84(3):26, 2022. doi: 10.1007/s00445-022-01586-0.
- Trojanowski, J. and Eisner, L. Comparison of migration-based location and detection methods for microseismic events. *Geophysical Prospecting*, 65(1):47–63, 2017. doi: 10.1111/1365-2478.12366.
- Trugman, D. T. and Shearer, P. M. GrowClust: A Hierarchical Clustering Algorithm for Relative Earthquake Relocation, with Application to the Spanish Springs and Sheldon, Nevada, Earthquake Sequences. *Seismological Research Letters*, 88(2A):379–391, 2017. doi: 10.1785/0220160188.
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, I., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., and van Mulbregt, P. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, 2020. doi:

10.1038/s41592-019-0686-2.

Waskom, M. L. Seaborn: statistical data visualization. *Journal of Open Source Software*, 6(60):3021, 2021. doi: 10.21105/joss.03021.

Werner, C. and Saenger, E. H. Obtaining reliable source locations with time reverse imaging: limits to array design, velocity models and signal-to-noise ratios. *Solid Earth*, 9(6):1487–1505, 2018. doi: 10.5194/se-9-1487-2018.

Wessel, P., Luis, J. F., Uieda, L. a., Scharroo, R., Wobbe, F., Smith, W. H., and Tian, D. The generic mapping tools version 6. *Geochemistry, Geophysics, Geosystems*, 20(11):5556–5564, 2019. doi: 10.1029/2019GC008515.

White, R. S. Askja 2007, International Federation of Digital Seismograph Networks., 2007. doi: 10.7914/SN/4F_2007.

White, R. S. Northern Volcanic Zone, International Federation of Digital Seismograph Networks, 2010. doi: 10.7914/SN/Z7_2010.

Willacy, C., van Dedem, E., Minisini, S., Li, J., Blokland, J.-W., Das, I., and Droujinine, A. Full-waveform event location and moment tensor inversion for induced seismicity. *Geophysics*, 84(2): KS39–KS57, 2019. doi: 10.1190/geo2018-0212.1.

Winberry, J. P., Anandkrishnan, S., Alley, R. B., Bindschadler, R. A., and King, M. A. Basal mechanics of ice streams: Insights from the stick-slip motion of Whillans Ice Stream, West Antarctica. *Journal of Geophysical Research: Earth Surface*, 114(F1), 2009. doi: 10.1029/2008JF001035.

Winder, T. *Tectonic earthquake swarms in the Northern Volcanic Zone, Iceland*. PhD thesis, University of Cambridge, 2022. doi: 10.17863/CAM.82505.

Winder, T., Bacon, C., Smith, J. D., Hudson, T. S., Drew, J., and White, R. S. QuakeMigrate, 2021. doi: 10.5281/zenodo.4442748.

Winder, T., Bacon, C. A., Smith, J. D., Hudson, T. S., and White, R. S. Datafiles for "QuakeMigrate: a Python Package for Automatic Earthquake Detection and Location Using Waveform Migration and Stacking", 2025. doi: 10.5281/zenodo.15236744.

Woollam, J., Rietbrock, A., Bueno, A., and De Angelis, S. Convolutional Neural Network for Seismic Phase Classification, Performance Demonstration over a Local Seismic Network. *Seismological Research Letters*, 90(2A):491–502, 2019. doi: 10.1785/0220180312.

Woollam, J., Münchmeyer, J., Tilman, F., Rietbrock, A., Lange, D., Bornstein, T., Diehl, T., Giunchi, C., Haslinger, F., Jozinović, D., Michelini, A., Saul, J., and Soto, H. SeisBench—A Toolbox for Machine Learning in Seismology. *Seismological Research Letters*, 93(3):1695–1709, 2022. doi: 10.1785/0220210324.

Yemane, T., Hudson, T. S., Kendall, J. M., Blundy, J., Tadesse, A. Z., Hammond, J. O. S., Ayele, A., Ogubazghi, G., and Lapins, S. Interconnectivity of Magmatic and Hydrothermal Systems of Aluto Volcano in the Main Ethiopian Rift Inferred From Seismicity. *Journal of Geophysical Research: Solid Earth*, 130(6): e2024JB031053, 2025. doi: 10.1029/2024JB031053.

Zhu, W. and Beroza, G. C. PhaseNet: a deep-neural-network-based seismic arrival-time picking method. *Geophysical Journal International*, 216(1):261–273, 2018. doi: 10.1093/gji/ggy423.

Zhu, W., McBrearty, I. W., Mousavi, S. M., Ellsworth, W. L., and Beroza, G. C. Earthquake Phase Association Using a Bayesian Gaussian Mixture Model. *Journal of Geophysical Research: Solid Earth*, 127(5):e2021JB023249, 2022. doi: 10.1029/2021JB023249.

The article *QuakeMigrate: a Python Package for Automatic Earthquake Detection and Location Using Waveform Migration and Stacking* © 2026 by Tom Winder is licensed under CC BY 4.0.