Chapter 1

Machine learning and fault rupture: a review

Christopher X. Ren[1]*, Claudia Hulbert[2]*, Paul A. Johnson[3], Bertrand Rouet-Leduc[3]**

[1]Intelligence and Space Research Division, Los Alamos National Laboratory, Los Alamos, NM, USA
[2]Laboratoire de Géologie, Département de Géosciences, École Normale Supérieure, Paris, France
[3]Geophysics Group, Los Alamos National Laboratory, Los Alamos, NM, USA

* These two authors contributed equally.
** Corresponding author email: bertrandrl@lanl.gov.

ABSTRACT

Geophysics has historically been a data-driven field, however in recent years the exponential increase of available data has led to increased adoption of machine learning techniques and algorithm for analysis, detection and forecasting applications to faulting. This work reviews recent advances in the application of machine learning in the study of fault rupture ranging from the laboratory to Solid Earth.

KEYWORDS

Machine Learning, Faulting, Earthquakes

1.1 INTRODUCTION

The study of material failure and rupture in geophysics is an extremely broad field [1], involving observation and analysis of geophysical data from failure simulations at laboratory and field scale [2, 3, 4, 5, 6, 7, 8, 9], laboratory experiments [10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26?] and in the solid Earth [1, 27, 28, 29, 30, 31, 32, 33, 34, 35]. Despite the vast apparent scope of this field, one approach to distill its essence is based on the following question: given a seismic signal received by a sensor or a set of sensors, what information can be gleaned concerning the process generating the signal? Often these signals are noisy, numerous, and exist in high-dimensional spaces where it is non-trivial to extract meaningful information from them [36, 37].

It then follows that the study of fault slip, whether in simulated granular fault gouge, or *in situ*, is a rich area of application for machine learning (ML).

Solid Earth geoscience and seismology are historically data-rich areas of research [36, 37], drawing insights from seismic, electromagetic, magenotelluric, gravity datasets among others. Indeed, some applications of what might be considered 'modern' ML techniques originated from the fields of geostatistics and geophysics. For example, Gaussian process regression was initially developed from an effort to estimate the distribution of gold from samples from boreholes [38] in the 1960s. Similarly, least absolute shrinkage and selection operator (LASSO) regression, was initially introduced in the geophysics literature for the linear inversion of band-limited reflection seismograms in 1986 [39], before being independently discovered and further developed in the statistics community in 1996 [40]. Currently these fields are undergoing rapid and substantial changes with respect to the sheer volume of data available to the public: Kong *et al.* noted the volume of seismic waveform data available from Incorporated Research Institutions for Seismology (IRIS) is growing at an exponential rate and that as of September 2018 the total size of the IRIS Data Management Center archive was 496.8 TB [37]. As of May 1, the archive was 644.5 TiB (708 TB), and has been doubling roughly every 3 years for the past 20 years, as shown in Figure 1.1. A similar dramatic increase in acquired data has been taking place in laboratory experiments of faulting. The case of the Rock and Sediments Mechanics Lab at Penn State University is telling: recent advances in experimental setups, related in particular to new high throughput acoustic data acquisition systems have led to a surge in experimental data volume, with a doubling roughly every year for the last 5 years, as shown in Figure 1.2. This explosive increase in available data presents many challenges: increasing volumes can lead to prohibitive access and processing times, however it also provides exciting new opportunities as ML algorithms thrive in problem settings with large volumes of data. The objective of learning algorithms is to extract patterns from training data which *generalize to unseen data* [41], rather than simply fit data. Thus, providing learning algorithms with more high quality data typically enables them to extract relevant patterns more accurately.
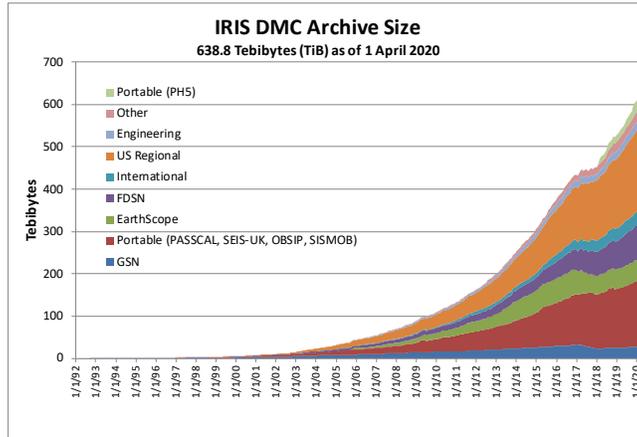
FIGURE 1.1 IRIS Data Management Center archive size since its inception. The archive's size has been doubling every 3 years since the 2000's. Courtesy of IRIS.
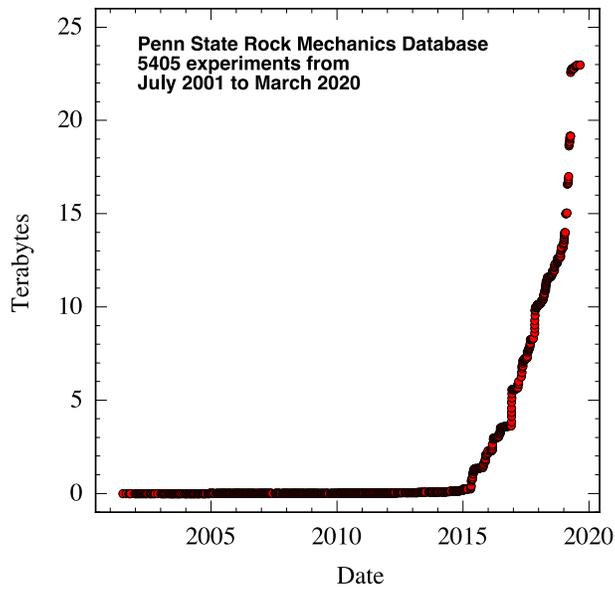
FIGURE 1.2 Experimental data stored at the Rock Mechanics Lab at Penn State. The database's size has been doubling every year since 2015. Courtesy of Chris Marone, Penn State University.

This work will review recent advances in the application of ML in studying seismic signals originating from fault slip, rupture and failure. We will provide an overview of the manners in which ML has been applied to advance research in studying fault slip in laboratory experiments and in the solid Earth.

1.2 MACHINE LEARNING: A SHALLOW DIVE

Machine learning (ML) and deep learning are scientific disciplines in their own right, and readers interested in a deeper dive will find many reviews and books dedicated to the subject [41, 42]. In this section we will provide a brief overview of ML. Previous reviews have already covered many of the types of algorithms used in ML and their properties [37, 43], so we will instead cover overarching definitions and concepts which unify these different algorithms.

1.2.1 Learning Tasks, Performance and Experience

The oft-quoted and rather succinct definition of a learning algorithm by Tom Mitchell [44] goes as follows: "

"A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P , if its performance at tasks in T , as measured by P , improves with experience E ."

Typically, the types of tasks **T** are grouped into two main categories: classification and regression. In a classification task, the objective of an algorithm is to specify which predefined categories a given input belongs to, and outputs a categorical value. For a given regression task, the objective is to predict a numerical value given some input. The main differences between these two types of tasks is the format of the output.

It is interesting to note that often the difference between these two types of learning algorithms is simply the type of output: the 'machinery' behind the two may be the same. Nowhere is this more obvious than with logistic regression, often one of the first 'classification' algorithms taught in introductory ML classes. Logistic regression is itself a regression model which estimates the *probability of class membership* as a multi-linear function of the input features. Only when combined with a decision rule, or a threshold, can we claim to be performing a classification task using logistic regression. A simple example can be seen in Figure 1.3, which shows how both the logistic regression model and a decision rule (set at 0.5 here) can be combined to create a classifier.
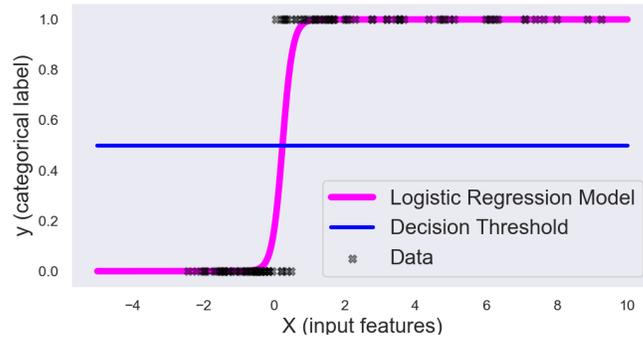
FIGURE 1.3   Logistic regression combined with a decision rule at 0.5. Adapted from [45]

We note that there are many other categories of tasks not covered by this introduction, such as: anomaly detection, data synthesis and inputation, ranking, density estimation, machine translation among others. For a comprehensive overview of these we refer the reader to [41]).

In the context of this review, the task **T** is typically one of detection (e.g. does this continuous seismic signal contain volcanic or nonvolcanic tremor, or an earthquake?), or prediction (e.g. how close is our system to failing, given a portion of some continuous signal?).

The performance measure **P** in this context is a quantitative measure of the ability of the learning algorithm, and is task-specific. The performance metric during the training procedure is measured through the value of the loss function. The act of choosing an appropriate loss is typically quite an involved one, as it can influence model robustness and computational cost (and sometimes even the feasibility of the training). A simple example is shown in Figure 1.4, where we fit two linear models, one with a least-squares loss and one with a Huber loss [46] to a simple linear dataset with outliers. Clearly in this case the effect of the Huber loss is to render the linear model more robust to the outliers. In testing, the key performance measure is the evaluation metric applied to the test-set: since we are most interested in the ability of the learning algorithm to generalize to unseen data, a portion of the available data is reserved for the testing procedure. The trained algorithm is evaluated on this 'unseen' separate dataset, which was not involved in the training procedure. The evaluation metric is usually, but not necessarily, the same as the loss function.

Finally the experience **E** provided to the learning algorithm allows us to draw one of the primary distinctions between types of learning algorithms: **supervised learning** algorithms are provided with a dataset containing input features where each example is associated with a target (the y-axis in Figures 1.3 and 1.4). The term 'supervised' originates from the idea that by providing the learning algorithm with a target we are essentially instructing it to find patterns related to
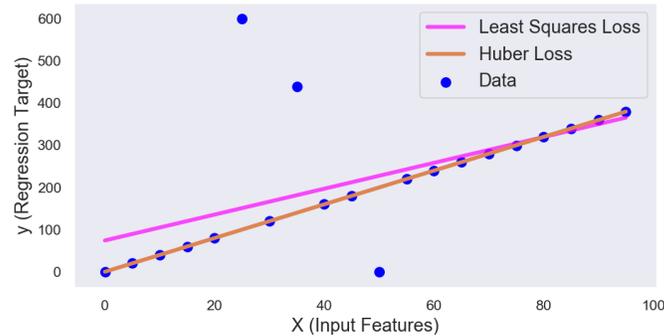
FIGURE 1.4   Linear models fit to data with l2 loss (fuchsia) and Huber loss (orange line).

this target [41].

On the other hand, in an **unsupervised learning** setting the learning algorithm is provided with the input features, but no label or target. The goal of the learning procedures is then to learn useful patterns from the dataset such as dividing the dataset into similar groups (clustering), or learning the generating probability distribution behind the dataset (density estimation, denoising, synthesis) [41].

1.2.2 Learning Capacity

We noted in section 1.2.1 that the central aim of ML is to learn a model which can generalize effectively to unseen data (i.e. hypothesis testing). During the learning process, an ML algorithm has access to some training data, on which it must attempt to reduce some error measure, known as the training error. What separates ML from fitting, is the requirement that not only must this training error be minimized, but the generalization error on an unseen test dataset must also be reduced.

The values of the training and test error and their relationship are linked to two central issues in ML: underfitting, occurring when the model is incapable of minimizing the training error to an satisfactory level and overfitting, when the training error has been reduced but the model cannot reduce the gap between the training error and test error (e.g., the model fits the training dataset but exhibits a poor fit with the training dataset). These issues are governed by the ML model's representational **capacity**, or complexity. Models with low capacity may struggle to capture the full spectrum of patterns in a complex training set, and thus underfit to them. Likewise models with high capacity can memorize properties of the training set which are irrelavant to the model's ability to generalize to the test set, thus overfitting. Figure 1.5 illustrates this principle: we take the simple task of fitting a dataset generated by a quadratic function: the

linear model (green frame) does not have the capacity to fit the data, whereas the high-degree polynomial model (blue frame), with more parameters than training examples, overfits due to the fact is has the representational capacity to pick many potential solutions to the problem, reducing the chances of selecting the correct solution. The model at optimal capacity (red frame), is well matched to the complexity of the task, and will generalize well to unseen data drawn from the same distribution as the training data.
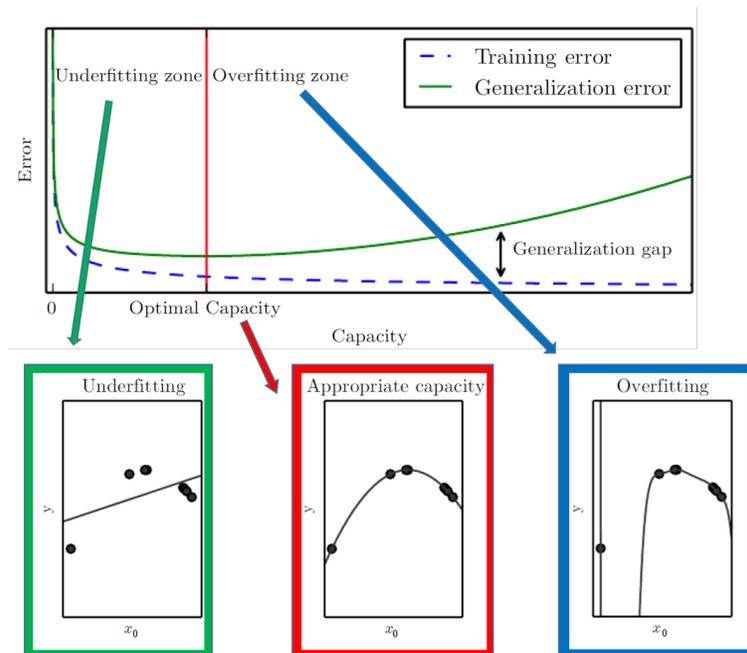
FIGURE 1.5 The relationship between model capacity and error. In the underfitting region shown at left (green frame), both training and test error decrease with increasing capacity. Once the capacity of the model exceeds the complexity of the data, the likelihood of overfitting, and thus increasing the test or generalization error increases (blue frame). The optimal model (red frame), is well matched to the complexity of the task, and will generalize to unseen data drawn from the same distribution as the training data. (Adapted from [41].)

We note that in the case of scientific research, the models and results published are often in the *overfitting* regime. This is likely to be due to several reasons: one aspect is inherent to the scientific publishing process whereby underfit results are less likely to be considered novel or spectacular and thus less likely to be published journal articles. Another is related to data sparsity: though the overall volume of available geophysical data may be increasing, when compared to fields such as computer vision or natural language processing the amount of *labelled* data available for specific tasks is generally still small (examples of labelled data are earthquakes, low frequency earthquakes or tremor that have been identified

and verified by an analyst). Last but not least, the test set will rarely be kept strictly independent from the learning phase, as researchers will go back to tuning their models **after** having seen their results on the test, effectively resulting in 'data snooping' and some overfitting [47]. However, even slightly overfit models are likely to contain useful insights into the physical processes they are applied to if interpreted with caution.

1.2.3 Geophysical Data

An assumption often made in the ML literature is that data samples are i.i.d., independently and identically distributed (i.e. the value from a sample has no influence on the other samples, and every sample is drawn from the same distribution). This i.i.d. assumption is more often than not violated in geophysical data. Time series data, **as well as features derived from it**, are auto-correlated, and the underlying physical process typically evolves over time. Geophysical data are often neither independently distributed, nor identically distributed. Auto-correlation in particular is extremely common in geophysical data analysis, and therefore such data should be handled with care. For example if one is analyzing data from a laboratory experiment: measurements (e.g. shear strain) at a time $t + 1$ are strongly dependent on the state of the experiment at time t, and both measurements are correlated. Similarly, field measurements of geophysical phenomena (e.g. seismicity or GPS displacement) will be strongly correlated depending on their proximity in time. A fundamental corollary to the auto-correlation of geophysical data, is that a contiguous train-test split is required [48] to ensure independent hypothesis testing, e.g. the first n years of data are used to train the model and the following n years are used to test it (e.g. [49]). Moreover, when dealing with geophysical data, strategies to ensure stationarity can be necessary, for example by differentiating the data over time (i.e. trying to relate rates of change of properties to each other instead of the properties themselves [50]).

1.3 LABORATORY STUDIES

Laboratory friction experiments have served as a framework for the study of earthquake physics since the original experiments of stick-slip on a fault analog by Brace and Byerlee [51]. The theory of earthquakes as frictional instabilities and fracture propagation with associated elastic radiation stems from such experiments [1, 51]. The favored phenomenological theory describing fault behavior is rate and state friction, [52, 53, 54] although there are a host of others (e.g., [55]. Rate and state friction has been placed in a more physically-based context by using Arrhennius theory as a basis by Rice and colleagues [56].

Just as in field seismology and geodesy, laboratory experiments are seeing an explosion terms of volume of experimental data recorded, as noted previously. Experimental setups are improving with an increasing number of acquisition

sensors, while these sensors are themselves rapidly improving in acquisition speed, sometimes resulting in terabytes of data for a single experiment [20], making the field ripe for the use of ML. As distributed acoustic sensors [57] are incorporated into laboratory shear experiments, data volumes are expected to grow even further.

There has been a recent focus on using laboratory earthquake experiments to tackle unsolved mysteries of earthquake physics: earthquake precursors[58], earthquake nucleation [59], and the interplay between slow earthquakes and dynamic rupture [59, 60], just to name a few. In this section we will review applications of ML that are helping to answer these fundamental questions in the context of laboratory friction experiments.

1.3.1   Laboratory Geodesy

In this section we will highlight recent advances in applying ML techniques to analysing what we term 'laboratory geodesy' experiments. While point strain sensors have been employed in shear experiments for years using strain gauges or laser measurements, advances in high-resolution strain monitoring enabled by techniques such as digital image correlation [61] are advancing the field of analogue seismotectonic modelling by allowing for the measurement of small-scale deformation on the order of magnitude of displacements related to single earthquakes [24]. We note that digital images including digital image correlation have and are been used for study of nucleation and phenomena such supershear rupture by a number of groups using photoelastic materials, e.g. [12] as well as in photoelastic experiments using 'granular' discs [62, 63].

Seismotectonic scale models, represent a new form of laboratory experiments, possessing several desirable features in terms of rupture and earthquake modelling: co-seismic dynamic weakenening which may occur in nature for reasons such as frictional melting or thermal pressurization can be mimicked as these models feature non-linear frictional material properties [24]. Laboratory seismotectonic models can also feature properly scaled elasticity, allowing for depth-dependent pressurization of the 'faults' in a more realistic manner.

Currently, all of the efforts in applying ML to the analysis of seismotectonic scale models make use of data from a viscoelastic gelatine wedge experiment initially developed by Corbi *et al.* to study subduction megathrust earthquakes [25]. In this setup, a gelatin wedge is underthrust at a constant rate by a dipping, rigid plate. This experiment is meant to be the analogue of a subduction zone, as shown in Figure 1.6 a. The interface between the gelatin wedge and the rigid plate is where the analog megathrust earthquake develops. Two velocity weakening patches (asperities modelled using gelatin on sandpaper) separated by a velocity strengthening patch (a barrier modelled using gelatin on plastic contact) are embedded in this analog megathrust. The viscoelastic seismotectonic model undergoes stick-slip cycles characterized by stress build-up punctuated by the nucleation of frictional instabilities (earthquakes) propagating at the interface

between the gelatin and the subducting plate [64]. The velocity field is extracted
by image cross correlation between consecutive images: this field is discretized
into windows which in this case provide analogues of GPS stations above the
model surface. Corbi *et al.* refer to these discrete windows derived from image
cross correlation as 'synthetic GPS stations' [64].
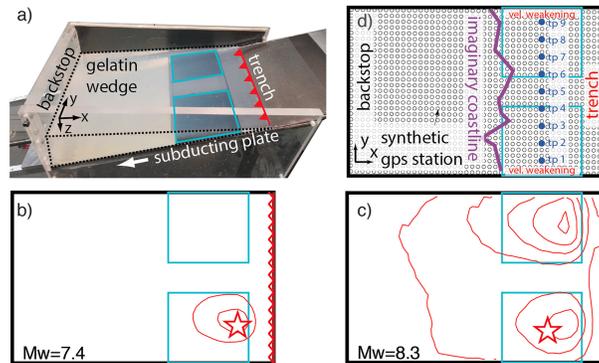
FIGURE 1.6 a) Experimental setup of a viscoelastic model of megathrust earthquakes. Cyan
rectangles delineate the velocity weakening patches, with examples of ruptures involving b) one or
c) two patches characterized in terms of co-seismic displacement shown as the red contour, with the
stars showing the epicenters. d) Monitoring performed using image cross correlation, each circle
represents an discretized area used as a 'synthetic' GPS station. Adapted from [64].

Corbi *et al.* utilised a Gradient Boosted Regression Trees algorithm [65] in
an attempt to predict the time to failure for the next cycle.

The input ML features used are derived from the velocity field measured
at the synthetic GPS stations, and are statistical features such as the variance,
kurtosis and skewness of the velocity field, as well as correlations between
subsequent frames and measures of cumulative displacement (see Supplementary
Information of [64] for more details). The authors use information from previous
cycles as input into the model; this is a hyperparameter that can be optimized:
the authors reported using 10 previous cycles gives the best model performance,
shown in Figure 1.7 a, with a coefficient of correlation R = 0.3 [64]. The authors
then further refined their model by using as label the time to failure at specific
locations, denoted by the nine blue points in Figure 1.6 d, with the results of
these different models shown in Figure 1.7 b. It is stated in [64] that the data
from the points located above the velocity weakening regions allowed for the
learning of a model with $0.7 < R < 0.8$, even when using shorter training
windows of 5 cycles. Analysing the spatio-temporal distribution of the time to
failure predictions allowed for the prediction of not only the timing and location
of upcoming ruptures but also the size by examining how many adjacent points
were expected to rupture simultaneously [64]. The most predictive features were
determined to be measures related to cumulative displacement at specific points,
indicating the learning algorithm is heavily informed by the relative loading

history of these positions.
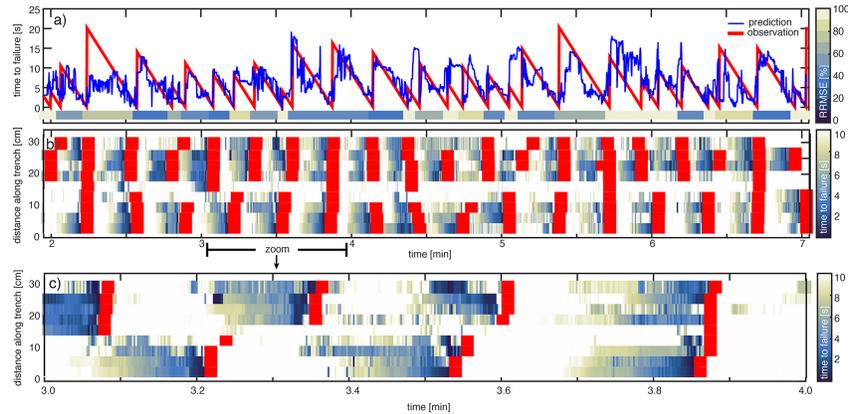
FIGURE 1.7   Results from Corbi et al. ([64]) on a viscoelastic model of a sequence of megathrust earthquakes.   a) Time to failure prediction with the cycle specific performance shown along the color-coded horizontal line, b) Spatio-temporal evolution of the predicted time to failure for the nine points shown in Figure 1.6 d., with the distribution of the earthquakes highlighted by the red squares, c) 1 minute zoom of the period shown in b).

A further study was conducted by Corbi *et al.* on the same dataset in [66]. Rather than framing the problem as a regression task [64], the authors instead chose to train an ML algorithm to determine whether the deformation observed was characteristic of a time window immediately preceding slip onset or not: a binary classification problem [66]. The authors also modified the 'positions' of the synthetic GPS network in the experiment, excluding data originating from the equivalent of the 'offshore' seismogenic zone in an attempt to mimic limitations in geodetic coverage in real subduction zones.

The features used in this study consisted of the displacement measured at a given synthetic GPS station.   The only features used for this study were thus displacements parallel and orthogonal to the 'trench' (or by association, the coastline), as these were determined to be the variables most predictive of slip [64].   The target variable was derived from the same 9 locations as in [64], where the "slip" was identified at times where the displacement at these locations exceeded a threshold. For each event, the category of 'alarm' was assigned to the Δt seconds prior to the onset of this displacement rate. The output thus consists of 9 response vectors, with each element of each vector indicating whether a given time step is part of the 'alarm' category or not.

The authors use an ML algorithm known as Random Undersampling Boosting (RUSBoost) [67], an algorithm which combines random undersampling of the most prominent class with adaptive boosting [68]. This is an effective method in dealing with the imbalanced nature of the dataset: there are simply more 'negative' examples in the dataset than positive 'alarm' examples (depending on the

Δt the alarms represent 3-27% of the total observations).

Since in this case the features are essentially locations in the experimental set up, the authors utilise sequential feature selection to determine informative 'regions' or stations in the set up. Relevant features selected by this process are thus deemed 'informative' to the model, whilst redundant or uninformative features which were discarded by the feature selection process are not. The authors find that regions adjacent to the two velocity weakening zones are the most informative in terms of monitoring the state of the system.

1.3.2 Laboratory Seismology

In this section, we will focus on applications of ML to 'laboratory seismology'. In these experiments, an analogue model (a model of the fault block variety) is used to generate acoustic emissions, which are collected by accelerometers, strain gauges or acoustic sensors.

The first application of ML to the analysis of failure in a laboratory system was performed by Rouet-Leduc *et al.* [15]. This system was the basis of many further ML studies [16, 17, 18, 20, 69], as well as the basis for a Kaggle ML competition that drew more than 4500 teams [70], and as such merits some description in the context of this review. The experimental apparatus consists of a biaxial shearing device: a dual fault-configuration containing fault gouge material is driven at constant normal load and a constant velocity. The driving piston accelerates during slip. Simulated faults fail periodically throughout repetitive loading and failure cycles, with the goal of mimicking loading and failure on a fault patch in Earth. Acoustic emission (AE) produced by the shearing layers is recorded using an accelerometer, and the shear stress resulting from the driving block, shearing rate, gouge layer thickness, friction and applied load are monitored and recorded, as shown in Figure 1.8.

This biaxial shearing system has been thoroughly studied and characterized in the past in many studies [11, 58, 71, 72]. Stick-slip frictional failure causes the shearing block to displace, whilst the gouge material undergoes dilation and strengthening. Characteristics of the 'critical-stress regime' exhibited by the material as it approaches failure, include minute shear failures emitting impulsive bursts of AE, as shown by Johnson *et al.* in [58]. The gradual progression of the system towards instability culminates in a 'laboratory earthquake', or stick-slip failure characterised by rapid displacement of the shearing block, precipitous decreases in friction and shear stress, and compaction of the gouge layers.
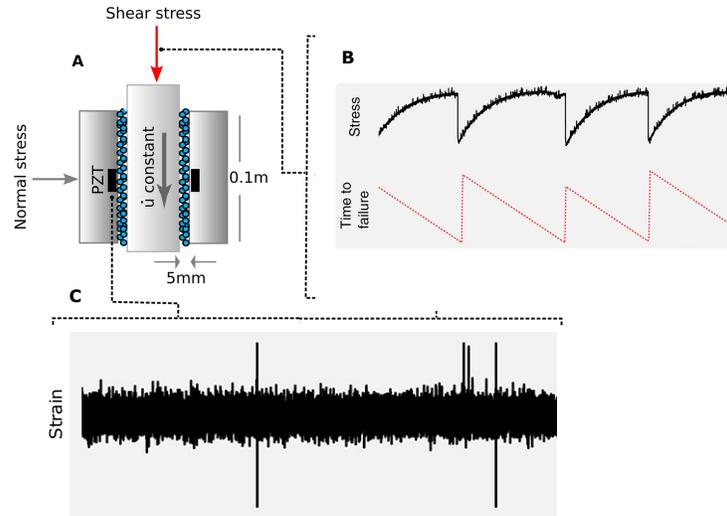
FIGURE 1.8   a) Biaxial shearing apparatus b) recorded shear stress and derived time-to-failure c) recorded strain by the accelerometers. Adapted from [15, 16].

In the initial study by Rouet-Leduc *et al.*, ML was applied in an exploratory manner to identify relevant statistical features derived from the dynamic strain acoustic emission signal. Approximately 100 statistical features (mean, kurtosis, variance etc...) were computed from moving time windows over the signal shown in Figure 1.8 C, which were then used to predict time remaining until the next stick-slip failure (red trace in Figure 1.8 B). In this study, an ML algorithm known as Random Forest (RF) [73], an ensemble method based on decision trees [74], was used to model this relationship. The learning algorithm identifies a relationship between a given time window of the acoustic emissions and the time remaining until the next stick-slip failure, as evidenced by the high R^2 value of the model.

The key takeaway from this experiment is that the time to failure in this system can be derived with fairly good accuracy even in the initial stages of the slip cycle, indicating a continuous progression towards failure. It is also important to note that even large fluctuations in slip cycle periodicity, such as the drastically shorter cycles shown in Figure 1.9 e, were accurately modeled by the RF algorithm. The system is clearly deterministic; however, in the critical state approaching failure it may exhibit some stochastic behavior. This is a topic that needs exploring.
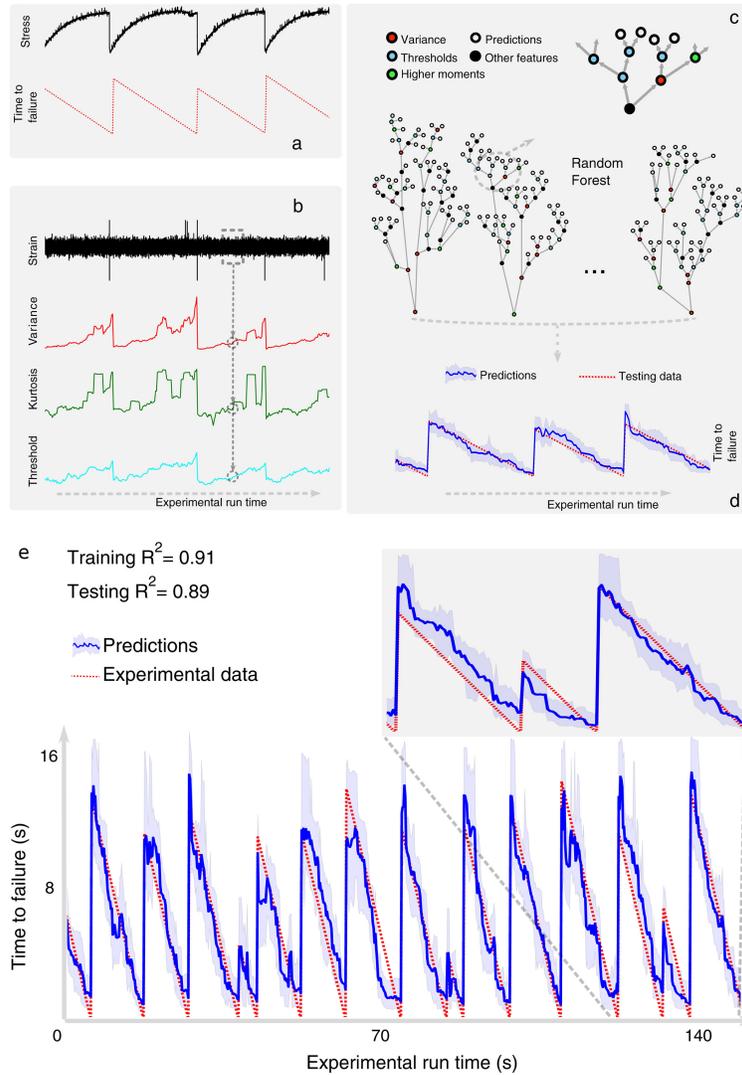
FIGURE 1.9 Biaxial experimental results. a) Shear stress and time to failure (labels), b) acoustic emission and derived features c) RF model schematic: the model averages the predictions of multiple decision trees, where each tree makes a prediction (white node) given an example based on a series of decisions made on the features associated with that example (colored nodes) d) and e) test-set performance of the trained RF model for the time to failure prediction task. Adapted from [15].

This initial success indicating that the state of the laboratory fault can be derived from the time windowed features prompted a further study with the aim of explicitly predicting the frictional state, or shear stress of the fault using a

similar approach [16]. Feature importance derived from the model revealed that the variance of the acoustic signal in a given time window was the most predictive feature in determining the instantaneous friction of the fault system. Even under varying normal loads, the biaxial apparatus demonstrated a predictable relationship between frictional state (shear stress) and the seismic power (variance of the signal), with the observation that scaling the seismic power by the cube of the normal load revealed a single, empirical 'friction law' for all normal loads.

Lubbers *et al.* demonstrated a different methodology employing the same data used in [15] but deriving input features from event catalogs built from acoustic emission rather than directly from the continuous signal [18]. In this study the features were chosen to capture cumulative statistics of event counts and amplitudes for a given time window. The authors demonstrated the ability of a RF model trained on the features to predict shear stress to a high degree of accuracy, as well as time *since* failure, with a lower performance for time to failure. The results suggest that models using cataloged-based features can reach a performance similar to models built on features from the continuous signal, but only in the presence of an extremely complete catalog, something currently not possible in Earth. The performance rapidly deteriorates when the smallest events are progressively removed [18].

Hulbert *et al.* [75] used a similar approach to analyse an experiment with aperiodic slip cycles consisting of both short and long slip durations, showing good performance in predicting not only the instantaneous shear stress of the system, but also fault displacement, fault gouge layer thickness, slip duration, and event magnitude. Signals originating from both slow and fast earthquakes behave in the same way when attempting to deduce the state of the fault. The level of acoustic power early on in the slip cycle is found to be correlated to the duration and the stress drop of the impending slip event.

Bolton et al. [20] further analyzed the data used in [15, 18, 75], using unsupervised learning to study the evolution of features of the seismic data throughout the laboratory earthquake cycle, with a focus on identifying precursors from their anomalous distribution.

Jasperson et al. [69] analyzed the same data set. They applied a Conscience Self-Organizing Map to perform topologically ordered vector quantization based on seismic waveform properties. The resulting map was used to interactively cluster catalogued events according to 'damage mechanism' of the granular fault gouge material, a term borrowed from the domain of nondestructive evaluation. They applied an event-based Long Short Term Memory (LSTM) network to test the predictive power of each cluster. By tracking cumulative waveform features over the seismic cycle, the network forecast the time-to-failure of the fault.

1.4 FIELD STUDIES

This section will review recent progress in Earth at the interface between ML and geophysics for the observation of the earthquake cycle, with a focus on novel

applications of ML to better detect and characterize earthquakes, precursors to earthquakes (foreshocks, slow slip), and aftershocks.

Most of the recent advances at this interface between ML and geophysics have been made in seismology, with the goal of improving the detection of earthquakes of small magnitude, with broad-ranging applications: (i) for the study of earthquake dynamics – clusters in time, e.g. earthquake swarms, clusters in space, e.g. finer fault structures, (ii) for the study of foreshocks – elevated seismicity prior to main earthquakes that may indicate the nucleation of the main shock, (iii) for the study of aftershocks. These advances have relied for the most part on deep learning techniques, which perform exceedingly well on complex problems containing unstructured data such as seismic data. Another avenue of research at the interface between ML and geophysics that has seen a number of studies in recent years has been the search for predictability in earthquake catalogs. Finally, we will review the early works emerging from a field of study that is poised for dramatic advances: the interface between ML and geodesy, interferometric synthetic aperture radar (InSAR) in particular.

The stakes are high in terms of the detection of precursory signals, as these may lead to an improved understanding of earthquake nucleation, one of the most challenging and elusive phenomena in geoscience. This challenge is characterized most notably by the discrepancy between laboratory studies, where clear precursors are generally observed (foreshocks [76], aseismic slip [59, 77]) as discussed in section 1.3, and field earthquakes where precursors are only sometimes observed [78, 78, 79] for reasons that may related to earthquake catalog fidelity, or to the fact that not all earthquakes exhibit foreshocks.

In a broad sense, two end-members of earthquake nucleation exist in the literature [80, 81]. The cascade model of earthquake nucleation posits that a spontaneous rupture triggers an avalanche of foreshocks that leads to the main shock, with no prior activity. The preslip model of earthquake nucleation posits on the contrary that earthquakes are triggered by previous aseismic slip, of which foreshocks are a byproduct (see Fig. 1.10).
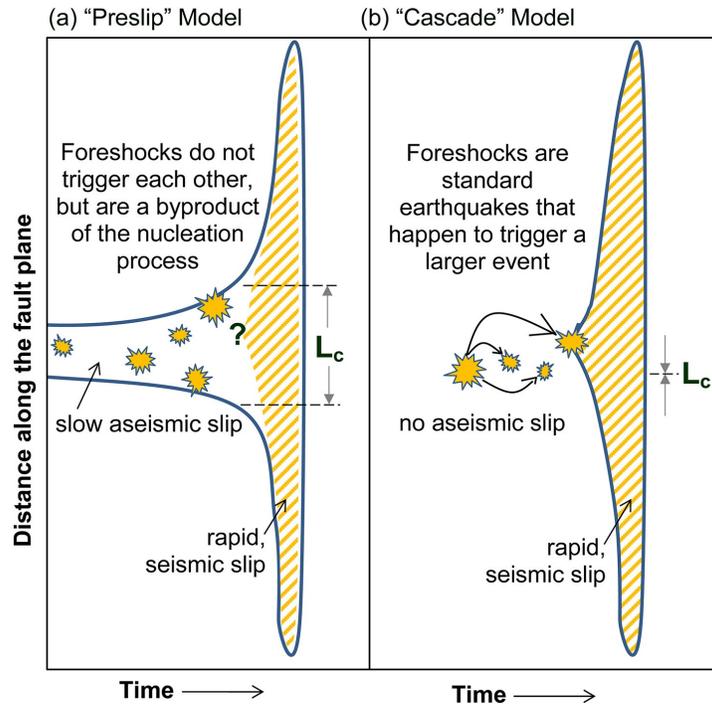
FIGURE 1.10 **End-members of earthquake nucleation models.** Figure from [81].

These two models have dramatic implications for the observation of earthquake nucleation, and open very different avenues for their detection using ML. The cascade model only leaves the improvement of earthquake catalogs, the current focus of the majority of ML research in geophysics [37], as an avenue for improving earthquake nucleation detection. The preslip model, on the other hand, leaves hope for a variety of possible improvements in the observation of earthquake nucleation through ML. The preslip model's hypothesis of foreshock and main shock triggering by aseismic slip adds slow slip to the list of potential observable earthquake dynamics. As a result, in addition to improved foreshock detection from improved earthquakes catalogs, improvements in a variety of geodetic observations of slow slip (GRACE, GPS, InSAR, etc.) stemming from ML may improve the detection of earthquake nucleation along with analysis of the continuous seismic signals emitted by the fault.

Recent direct [82, 83] and indirect [78, 84] observations of aseismic slip preceding large earthquakes hold the promise that with ML-enhanced detection of aseismic slip from novel seismic and geodetic data analysis, will come an enhanced detectability of earthquake nucleation. Likewise, recent improvements in the detection threshold of earthquake catalogs have shown that foreshocks may

be more prevalent before earthquakes than previously thought [31], indicating that improvements in precursor detection may also aid in detecting the elusive nucleation of earthquakes.

In the following sections of the review, in a first sub-section we will review recent ML-driven technical advances in earthquake detection and characterization, with a focus on (i) building higher-fidelity catalogs, a requisite in particular in the quest for precursory foreshocks, and (ii) improving the detection of slow earthquakes from novel seismology and geodesy processing, a promising avenue for the detection of earthquake nucleation.

In a second sub-section, we will review few developments at the intersection between rupture physics in the earth and ML that require specific developments in addition to event detection and characterization. This work will cover induced seismicity, a topic of intense focus in recent years, including the development of ML-based techniques to better relate it to injection or extraction practices. We will also provide an overview of the specific developments made for earthquake early warning with a focus on faster and more reliable detection as well as better early-magnitude estimation.

Lastly we will briefly cover earthquake forecasting, for which progress from ML techniques has arguably been lacking.

1.4.1 Techniques

1.4.1.1 Earthquake catalog building

Building catalogs with lower magnitudes of completeness has long been a major effort in seismology. More complete catalogs are useful for many applications such as the study of fault rupture, the identification of faulting networks, the estimation of seismic hazard and analyses of induced seismicity among others. They are also of major interest to identify foreshocks preceding large earthquakes.

The creation of catalogs has historically been conducted by hand in a long and tedious process involving the identification an event and its seismic phase arrivals on seismograms, associating these phases to a particular earthquake, and pinpointing it to a location. The manual nature of this task is a hindrance in terms of scalability. With the explosion in the volume of seismic data recorded over the last decade (see Figure 1.1), the need has arisen for automation of catalog building. Methods developed to combine the automatic detection of events and further verification by an analyst are now commonly used; however, because they still require human intervention, these approaches nonethless still suffer scalability issues.

Given that they remain unchecked by analysts, fully automated methodologies for building catalogs require sufficient robustness to ensure a minimal false detection rate. In the scope of fully automated event detection, using robust confidence metrics (and ideally measuring uncertainty within a Bayesian framework) is therefore of paramount importance.

ML algorithms are particularly adapted to this problem, as they provide an opportunity for fast and automatic detection. While standard confidence metrics for the existence of an event is often based upon simple statistics such as a correlation threshold, ML algorithms provide an assessment of performance that can be evaluated (in the case of supervised learning), and in some cases can enable the estimation of uncertainty of event detection. As a result, catalog building based upon ML techniques has been actively pursued and was introduced early in the literature, starting in the early 1990s [85, 86]. In the following, we introduce the original work in this field and describe more recent developments for event detection, picking, association, and location. We also conduct a critical analysis of the shortcomings of ML algorithms with respect to catalog creation.

Earthquake detection
The first applications of ML for earthquake recognition focused on discriminating the spectra of seismic waveforms due to earthquakes from the those due to explosions [85, 86]. While not directly related to construction of earthquake catalogs, these analyses pioneered the detection of characteristic seismic signatures of earthquakes with neural networks. They were soon followed by studies aimed at identifying earthquakes from seismic signal and discriminating them from noise. In 1995, Wang and Teng [87] found that neural networks trained with data from the Landers earthquake aftershocks outperformed the results of a standard STA/LTA standard threshold classification approach. In 1999, Tiira [88] reached similar conclusions when analyzing teleseism data in Finland, reporting that a neural network approach could improve the training catalog by 25%.

All the studies mentioned above relied on shallow, feed-forward neural networks (with the exception of [88], which utilized slightly more complex architectures in the form of simple recurrent networks). They relied on hand-built features (either waveform spectra, STA/LTA functions, or both), and were trained on a small number of examples (from a few dozen to several hundred).

Most of the following studies have relied on the same tools, and were also based on relatively small datasets, with examples in the 100-1000 range [89, 90, 91, 92]. However the majority of these studies also incorporated ML approaches other than feed-forward networks, such as SVMs [89, 90], Random Forests [91], or Logistic Regressions [91], which were found to systematically outperform the simple perceptron [93] models. A parallel algorithmic avenue explored to detect earthquakes has relied on probabilistic graphical models, using either Hidden Markov Models [94] or Dynamic Bayesian Networks [95]. These approaches provide further benefits in associating a true Bayesian uncertainty measure to each detection, which can be used to help filter out false positives. In the two studies, the authors comment that their algorithms produced a number of false detections, but the uncertainties associated to those false detections were much higher than those associated to real earthquakes. Unsupervised algorithms based upon regrouping event waveforms by similarity have also been explored

[96, 97, 98].

Recently, researchers have begun to train models using larger datasets [99, 100, 101, 102], enabling the use of more complex neural network architectures and automatic feature extraction. In particular, Mousavi et al. [101] used a database of 250000 earthquakes and 250000 noise samples in California to train a residual network based upon convolutional and recurrent units, and reported that their algorithm outperforms other approaches. Interestingly, when attempting to generalize their model from California to Arkansas, the model produced a higher rate of false positives, suggesting that complex models built upon large databases can have difficulties generalizing to other areas.

Overall, the analysis of the existing literature shows that earthquake detection with ML algorithms clearly outperform a standard Short Time Average over Long Time Average (STA/LTA) approach, in particular due to a lower rate of false positives as shown in Figure 1.11. The comparison with template matching algorithms is less clear-cut. Some definitive advantages of ML earthquake detection over template matching are i) reduced computation time, and ii) the fact that events not included in the template list can be detected. However, ML algorithms may generate more false positives than template matching. Further developments of Bayesian deep networks trained on large datasets to associate clear uncertainties to each detection, such as that proposed by [103], would be a straight-forward improvement of the state-of-the-art.
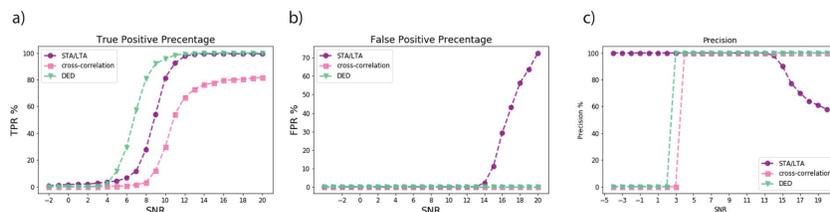
FIGURE 1.11 Comparison between neural network detection (DED), STA/LTA, and cross-correlation on semi-synthetic data, at different signal-to-noise ratios. Figure from [104].

The comparison between template matching and neural network detection using convolutional networks is of particular interest. Because convolutional layers perform a cross-correlation operation, template matching corresponds effectively to a single 1D convolutional layer using one kernel of known weights for each template, and as many channels as the number of stations. In contrast, within each convolutional layer a neural network updates multiple smaller kernels, thereby learning multiple templates at different resolution size of the data.

Another useful comparison relates to the generalization of earthquake detection algorithms. STA/LTA algorithms are obviously straight-forward to generalize to other areas and/or other seismic stations, but tend to generate a high rate of false positives. On the contrary, template matching algorithms do not

transfer well, and can only detect events for which templates have been identified by analysts. ML algorithms are promising in terms of building robust models that generalize due to their ability to learn templates and identify events different from those included in training. However, the current state-of-the-art does not seem to include models generalizing well to regions outside the training area. This is problematic as it precludes the use of such models for the analysis of regions of interest that are poorly instrumented, or where only a small number of events were cataloged. A likely future area of research may be focused on improving the generalization of these models.

Phase picking and polarity determination

Once an earthquake is detected, the arrival times of P- and S-waves have to be picked on seismograms in order to allow for the determination the hypocenter; the polarity of the P-wave is required to determine an event's focal mechanism. The identification of phase picks, as well as polarity, have been actively pursued with ML recently.

The first applications of ML for seismic picking was focused on recovering first arrivals from seismic reflection data, in the early 1990s [105, 106]. ML estimation of earthquake P- and S-wave arrivals was further investigated in 1995 by Dai and MacBeth [107], where a shallow feed-forward neural network (in combination with a discriminant function) was tasked with recovering arrivals from the vector modulus of the three components of seismic waveforms. A similar analysis by the same authors [108] extended the approach to arrival detections from a single station component. In both studies, a model trained on a very small number of examples (9 pairs of seismic arrivals and noise) was able to reach a relatively high picking accuracy on hundreds of earthquakes. Shallow perceptron architectures were also applied for the same purpose by [109, 110]. Gentili and Michelini [111] found that picking from a neural tree was robust to errors from noise data and outperforms a standard STA/LTA algorithm. Chen [112] also found that unsupervised algorithms outperform STA/LTA picking for small events.

In the late 2010's models were developed using much larger datasets which, combined with increases in computational power, enabled the use of more complex neural network architectures and automated feature extraction for phase picking. Zhu et al. [114] developed a U-Net type network for picking estimations, trained on over 600000 earthquakes in Northern California. The network takes as input the three components of seismic waveforms, and outputs either a P arrival, S arrival, or noise class. Wang et al. [115] trained two separate convolutional networks to pick P and S arrival times, on a database of over 700000 picks in Japan. Zhu et al. [113] also utlized on a CNN to study the aftershock sequences of the 2008 M W 7.9 Wenchuan Earthquake. For all studies, results in testing were found to outperfom existing methods. An online version of their algorithm, identifying picks on continuous waveforms, was also proposed. A sketch of the associated classification pipeline is shown in Figure 1.12.
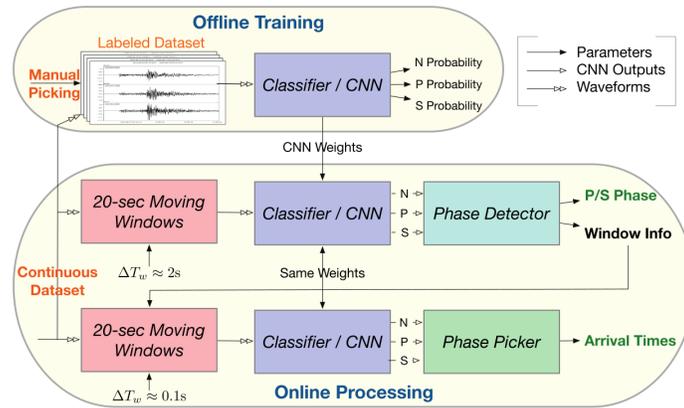
FIGURE 1.12 Schematic of an offline and online picking classifier. N stands for noise, and P and S represent the pick arrivals of the P and S waves. Figure from [113].

Besides picking arrival times, a few studies attempt to estimate P-wave first motion polarity. Ross [116] trained two different CNNs on 18.2 million manually picked seismograms in Southern California: one to identify P-wave arrival time, and a second to determine the polarity of the first motion. Hara [117] similarly relied on a CNN classifier to determine first-motion polarity from earthquakes in Western Japan.

Interestingly, models for phase picking appear to generalize better than models trained to detect earthquakes. Two recent studies report that their algorithms generalize well to regions outside of the training area, either without finetuning [115], or with minimal finetuning [113]. While this improved generalization could be attributed to larger training sets or model specificities, it might be the case that earthquake detection is intrinsically a more challenging problem in terms of generalization due to models indirectly learning specific Green's functions.

Event association

Once waveform picks are identified, they have to be associated to a given earthquake; the timing of arrivals appearing on multiple stations is analyzed to identify patterns characteristic of an earthquake. Standard methods mostly rely on travel time information for determining associations, which can be challenging in the presence of false arrivals, or when the earthquake rate is very high as shown in Figure 1.13. ML algorithms have only very recently been applied to this problem.
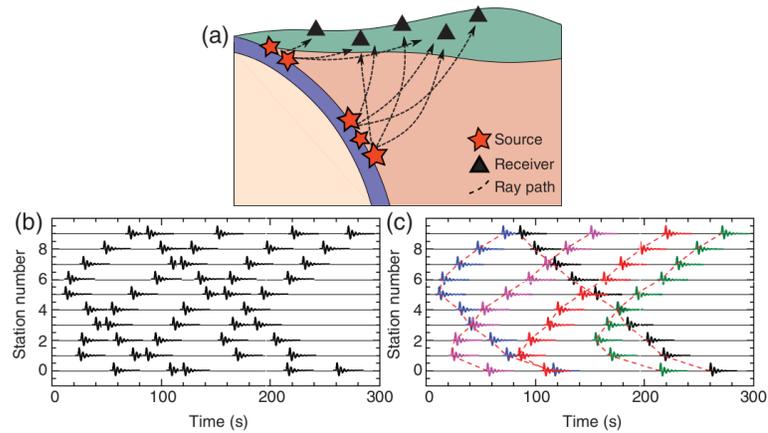
FIGURE 1.13   Illustration of the earthquake association problem. In particular, in the presence of dense seismicity, associating the earthquake arrivals recorded at different sensors to a given earthquake (subplots (a) and (b)) can be challenging. (c) shows a schematic of the associated arrivals (each color corresponds to an earthquake). Figure from [118].

The main approach employed in the literature is to associate several picks to an event by classification [49, 119]. McBrearty *et al.* [49] captained the use of convolutional neural networks to recognize whether or not a pair of seismic waveforms originated from the same earthquake. In Ross *et al.* [119], phases, travel times and station locations are fed as input to a recurrent network; the network then estimates whether picks within a moving window belong to the same event.

Another approach has focused on the use of clustering instead of classification for associating picks. In [118] a backprojection algorithm is tasked with identifying candidate sources. A spectral clustering algorithm, combined with an optimization routine, is then used to associate the arrivals to the smallest number of sources possible that still match the data. The application of the algorithm in Northern Chile allows the authors outperform existing catalogs by nearly an order of magnitude.

Event location

Automatic earthquake location with ML has been actively explored in recent years. Once detected, earthquakes need to be located, which usually involves the back-propagation of seismic waves towards a source, or their association to the location of a given template. ML algorithms provide a promising avenue to

refine and accelerate the location process. Two approaches have been explored in the literature: i) multi-station methods, that are likely to favor event move-out patterns for event location similar to existing seismology approaches; ii) single-station location, which employs the ability of algorithms to recognize the path of the wave's propagation encoded in seismic waveforms.

Among multi-station location techniques, past attempts to locate seismic sources have relied upon clustering algorithms and neural networks. Riahi and Gerstoft [120] use a graph-based clustering approach to regroup sensors affected by a common source; the source's location can then be inferred from the area spanned by the clusters. Trugman and Shearer [121] use hierarchical clustering for earthquake relocation. The clustering algorithm takes as input differential travel times, cross-correlation values, and starting locations; events grouped within similar clusters are then simultaneously relocated. This approach is also able to return estimates of location errors. Kriegerowski *et al.* [122] trained a neural network to locate earthquakes from multiple stations during an earthquake swarm in West Bohemia. The network takes as input the orthogonal components of all stations, and returns the depth, east and north source coordinates. The model is trained with events from a double-difference relocated catalog. Zhang et al. [123] rely on a fully convolutional network to estimate a 3D image of the earthquake location, from data recorded at multiple stations. While location errors are small for earthquakes greater than magnitude 2, smaller events were reported to be more challenging to locate using this approach and associated with larger errors.

Studies of earthquake location from single stations all rely on neural networks, either in the classification or regression setting. Perol *et al.* [100] and Lomax *et al.* [102] trained neural networks to output classes that correspond to geographical areas associated to source locations. In [100], a convolutional neural network was tasked with regrouping earthquakes within 6 geographic clusters. However, while the accuracy of the model's locations is reasonably good (74.5%), this metric drops drastically with increasing number of classes – potentially due to the small number of training samples. Lomax *et al.* [102] employed a similar idea, but with a larger number of classes; however the model returns high error rates. Mousavi and Beroza [124] used an innovative approach for earthquake location: two Bayesian neural networks were trained separately, tasked with estimating respectively the epicentral distance and the P travel time associated to an event. A parallel neural network was then used to estimate the back-azimuth and its uncertainty. By regrouping the results of the three models, an event location is proposed along with an associated uncertainty. Epicenter, origin time, and depth errors are small in this approach, and because the models were trained on globally distributed events, this approach is found to generalize to different areas.

Location algorithms based on a single station are of particular interest in this field. First, the fact that models are able to locate events from single stations shows that they can reconstruct the wave propagation path from seismic data,

thereby learning from the physics of the medium. Second, such models are portable, in the sense that they can be applied to areas outside of the training region. A model trained on several stations from a specific network will not be useful on data from other stations, which makes single-station models potentially more generalizable. Single-station models can also be useful for locating earthquakes that are recorded on a very few number of stations.

1.4.1.2 Seismic waveform denoising and enhancing

A very promising, recent use of ML in seismology is related to the ability of several algorithms to process or denoise data, thereby helping identify signals of interest, such as the the detection of small events in seismology. Denoised data can afterwards be combined with standard seismology analysis or other ML algorithms for event detection and location.

The ability of many algorithms to reduce the dimension of the data and build a sparse representation from original signals makes ML tools particularly suited for the task of denoising seismic signals. In the literature, two main approaches have been used for such analysis of seismic waveforms: dictionary learning and auto-encoders. Dictionary learning has been applied in particular to the denoising of seismic signals. In [79, 125], the authors show that dictionary learning algorithms outperform several other approaches for this task (such as wavelets, curvelets, etc,), and result in higher signal-to-noise ratios.

In one of the first application of auto-encoders in seismology Valentine et al. [126] show that seismic waveforms can be successfully encoded and decoded, therefore learning an effective low-dimension representation of seismic signals. This procedure is applied to the identification of seismograms of good quality versus those of bad quality, with the goal of building a selection criteria of retaining only waveforms that have high signal-to-noise ratios, leading to more robust analyses. A recent study published by Zhu *et al.* [113] similarly employs auto-encoders with the objective of separating seismic signals from noise. In this work, an auto-encoder trained on a large number of earthquake waveforms with high signal-to-noise ratios that are overlayed with noise waveforms was tasked with separating the signal from the noise. Once trained, the auto-encoder was able to produce masks corresponding respectively to the signal and to the noise. These masks were used to clean new waveform data; on a few examples, retaining only the masked signal is shown to drastically improve the STA/LTA functions associated to earthquakes (Figure 1.14), a demonstration of the potential of combining ML based waveform denoising with standard seismology tools. A different approach for the same problem has been developed by Rouet-Leduc *et al.* [127], who showed that neural network interpretation techniques can be used to denoise tremor waveforms and produce much cleaner recovered signals of interest.

A different and interesting application aimed at recovering seismic signal hidden in noise was proposed by Sun [128]. The authors rely on a convolu-
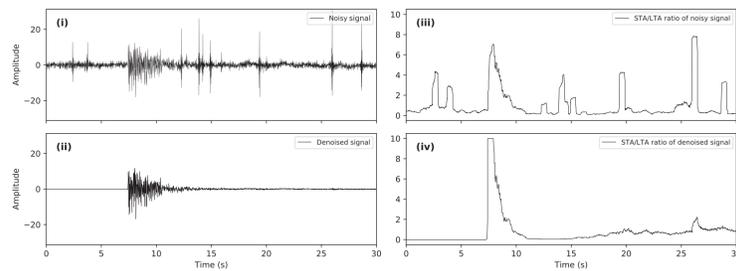
FIGURE 1.14   Denoising of seismic waveforms with an auto-encoder. Figure from [113]. Improvement of STA/LTA characteristic function after denoising. (i) Example of a noisy signals. (ii) Denoised signals. (iii) and (iv) Corresponding STAT/LTA characteristic functions. (a) and (b) two examples of earthquake detection.

tional neural network to extrapolate missing low frequencies in synthetic seismic records. Inputs to the model were bandlimited recordings of seismograms, with low frequencies removed. The model is tasked with recovering the missing low frequencies from these inputs. The model was able to estimate the unobserved low-frequencies, both in terms of phase and in amplitude, reasonably well.

Further developments of ML algorithms to denoise or enhance the quality of seismic data may be of considerable interest in seismology. These approaches are particularly promising for the identification of small events hidden in the noise. They can be used to leverage standard seismology tools while limiting the risks of false detection, and might help to increase waveform cross-correlation levels.

1.4.1.3   Tectonic tremor detection

Tectonic tremor, along with GPS measures of deformation, is the most direct evidence for aseismic slip [129]. As such, progress driven by ML in the detection of tectonic tremor is poised to advance the understanding of slow earthquakes and their interplay with regular earthquakes.

Recent advances in tectonic tremor detection using ML fall into two categories: developments of models to estimate tectonic tremor intensity and slow slip displacement rate from features of continuous seismic data, and developments of models to directly detect tectonic tremor.

In terms of the first category mentioned above: Rouet-Leduc *et al.* demonstrated in [50] that a ML model can be trained to accurately track slow slip rate on Vancouver Island, using features of the continuous seismic noise that are characteristic of tectonic tremor. Hulbert *et al.* showed in [130] that slow earthquakes under Vancouver Island are likely preceded by a months-long nucleation phase, during which seismic features characteristic of tremor rise exponentially.

In terms of the second category: we note here that direct detection of tectonic tremor is a task perfectly suited for deep neural networks, which deal particularly

well with rich unstructured data (e.g. text, images, videos, sound). Indeed, this has been demonstrated by Nakano *et al.*, who showed showed in [131] that a CNN could reliably discriminate between earthquakes and tectonic tremor using a catalog of tremors recorded near the Nankai trough in Japan. Rouet-Leduc *et al.* showed in [34] that a CNN trained to distinguish tectonic tremor from background noise could be used to extract the tremor signals to improve event detection, using seismic records from Vancouver Island. In [130], Hulbert *et al.* showed that this deep learning-based extraction of tectonic tremor signals enables the location of many more tremor events in Cascadia.

1.4.1.4 Fault slip inversion

Fault slip inversion relies on seismic and/or geodetic observations to determine earthquake source properties. Determining source properties is of importance for a variety of problems, in particular for early warning applications (see Section 1.4.2.1). Slip distribution also has implications for the distribution of subsequent aftershocks (e.g. [132]), as well as the general understanding of slip processes on faults (e.g. [133]).

Early applications of neural networks for slip inversion were developed by Kaufl et al. [134]. The authors trained a number of small neural networks to invert for slip properties from synthetic geodetic data, and successfully applied their models to real data from the 2010 M_w 10 El Mayor Cupah earthquake, using the predictions from their ensemble of neural networks as a proxy for an a posteriori distribution of the source parameters.

Slip inversion is a particularly difficult problem for ML algorithms, for two reasons: (i) inversion is a generally non-unique problem, and therefore the estimation of fault slip properties requires probabilistic approaches, which inhibits the available tools; (ii) slip inversion often relies on very heterogeneous sources of data (e.g. seismic and GPS data), where deep learning does not perform well in general. Recent advances in probabilistic deep learning [135] may solve this apparent incompatibility.

1.4.1.5 Automatic detection of geodetic deformation

Interferometric Synthetic Aperture Radar (InSAR) is a powerful geodetic technique developed in the early 90's [136, 137] for measuring ground surface displacements. InSAR has been successfully applied to monitor large displacements due to earthquakes [136, 138, 139], as well as smaller displacements related to interseismic deformation [140, 141, 142, 143], slow moving landslides [144], and slow slip events [145, 146]. However, in general these measurements have been successful only through painstaking manual exploration of the deformation data by InSAR experts.

In the past low satellite revisit rates typically resulted in large time delays between measurements. The launch of the Sentinel-1 satellite constellation in 2014 has been revolutionary in that it provides systematic radar mapping of
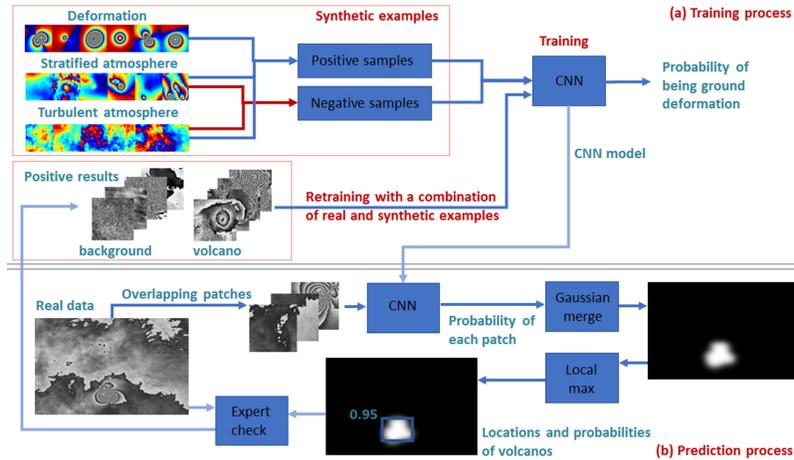
FIGURE 1.15 Deformation detection in a wrapped InSAR interferogram. Figure from [147]. Anantrasirichai *et al.* finetuned the AlexNet model on synthetic data and on real data to classify whether or not a portion of wrapped InSAR interferograms contains ground deformation.

all actively deforming regions in the world with a 6-day return period. This wealth of data represents an opportunity as well as a challenge. InSAR data processing is not fully automatic, with unwrapping errors in particular often requiring manual correction, although rapid progress is being made towards full InSAR unwrapping automation [148]. Furthermore, the analysis of InSAR deformation data requires expert interpretation, and usually requires a priori knowledge on the analyzed area. The launch of additional InSAR satellites with even shorter return period (NASA's NISAR mission in 2022) will prove even more challenging in terms of data volumes and computational and time costs. Automation of InSAR processing and analysis has been identified as one of the main challenges of the field of InSAR geodesy (R. Lohman, AGU Fall meeting 2019), and is poised to become a rapidly growing avenue of research, that may enable the automatic detection of small or slow deformation on faults.

Anantrasirichai *et al.* [147, 149, 150] began developing deep learning tools for InSAR analysis by finetuning the AlexNet model [151] on manually labeled wrapped InSAR data as well as synthetic data. Their models flag volcano deformation (see Fig. 1.15), demonstrating a 3 cm detection capability.

In [152], Rouet-Leduc et al. built a deep auto-encoder to output cumulative ground deformation from unwrapped InSAR time series, with a focus on fault deformation. Trained on synthetic data, the authors demonstrate the recovery of millimeter scale deformation on real data from creeping faults as well as small earthquakes.

1.4.2 Applications

1.4.2.1 Early warning

Real-time alerts of destructive earthquakes while they are in progress are based on Earthquake Early Warning (EEW) systems. The very proposition of earthquake early warning (EEW) – a rapid assessment of immediate seismic hazard – demands the use of algorithmic solutions. As a result, EEW has seen considerable use of ML methods as early as the 2000's. These systems rely on the fact that i) damaging S-waves follow identifiable P-waves and ii) seismic waves propagate at a few km/s, much slower than speed-of-light telecommunication messages (seismic stations closer to the earthquake source can therefore be used to alert locales further away).

Böse *et al.* [153] designed a series of small, fully connected neural networks on features of seismic recordings following the onset of P-wave arrivals to estimate location and magnitude of earthquakes, on simulated data. Ochoa *et al.* [154] used a similar approach on real data, feeding features of the waveform that follows the P arrival to a support vector machine to estimate earthquake magnitude.

Ongoing efforts by Kong *et al.* [155] to build a very large array of EEW detectors using a smartphone application (ShakeAltert), make use of a neural network applied to features of the phone's accelerometer data to estimate if the phone is recording an earthquake or some other motion. With the same goal of reducing false alarm rates from false detections, Li *et al.* [156] trained a generative adversarial network (GAN), consisting a generative network and a discriminator network, [157, 158], and use the discriminator network as a feature extractor for a random forest classifier tasked with telling earthquakes from impulsive noise on short waveforms from ShakeAlert data.

Meier *et al.* [159] developed a Bayesian algorithm that starts from single station analysis for the very early onset of the detection on nearby stations, and evolves into a more accurate multi-station analysis. Trained on a dataset of 60000+ waveforms, their algorithm can provide a timely assessment of earthquake magnitude for areas close to the epicenter, where damage is greater. Their ML approach of training and testing an empirical model also demonstrates that information about earthquake magnitude saturates at magnitudes that depend on seismic record length, but are in general about magnitude 7-7.5.

1.4.2.2 Induced seismicity

A major effort in earthquake detection and location with ML algorithms has been focused on induced seismicity, either in geothermal fields [160], in mines [92], or in gas [103] and oil [91, 100, 101, 113, 123] extraction fields (in particular in Oklahoma). These studies aim at building automatic and more complete catalogs, and were described in Section 1.

In this section we will therefore focus on another application of ML algorithms: the analysis of existing earthquake catalogs to better understand the

connection between injected fluids and induced seismicity. Two studies are of particular interest, both relying on probabilistic graphical models to extract patterns from catalog data.

Holtzman *et al.* [161] analyzed the waveforms of 46000 cataloged earthquakes in the Geysers geothermal field, in order to identify subtle changes in seismic characteristics of the signal. By combining Non-negative Matrix Factorization and Hidden Markov Models, the authors were able to extract low-dimension representations of the earthquake seismic waveforms, which were then fed to a clustering algorithm. The resulting clusters did not reflect the geographical proximity of earthquakes, and showed instead clear temporal patterns that appear to be connected to injection rates. Interestingly, not all clusters correspond to higher seismic rates at times of high water injection; in some cases, clusters are correlated with the minimum injected rates. An interpretation of the results is that the model is able to identify changes in faulting mechanisms, and that the spectral-temporal patterns identified mmay correspond to changes in earthquake source physics, possibly due to changes in the thermomechanical conditions of the reservoir. As a result these observations may have applications for operational monitoring,aid in mitigating the risk of dangerous earthquakes, and enable the optimizion of production efficiency from the geothermal field. Hincks *et al.* [162] also relied on probabilistic graphical models to analyze the connections between geological characteristics, well operations, and induced seismicity in Oklahoma. They provided well characteristics (injection volume, pressure, etc.), geological features, as well as seismic moment release (calculated from cataloged earthquakes) to a Bayesian network. The model was then tasked with modeling the joint conditional dependencies between these different variables. Interestingly, the injection depth with respect to the crystalline basement was found to correlate with seismic moment release much more strongly than the total injected volume. Therefore, lithologic and fault network characteristics were found to play a more fundamental role with respect to induced seismicity than simple injection volumes: deeper fluid injection may favor migration into the basement, and favor the reactivation of faults as others have observed, e.g., [163]. This analysis has important implications for operation management and regulation; in particular, it suggests that injecting water at lower depths (further away from the crystalline basement) may significantly reduce annual seismic moment release.

Both studies show that ML algoritms are promising tools for identifying patterns between known earthquakes and injection data, potentially leading to a better understanding of the physics of injection systems. Such findings may be useful for the operational management of geothermal reservoirs and extraction fields, by helping mitigate the risks of potentially damaging seismicity, and/or help to increase production. The combination of the aforementioned creation of denser catalogs, analysis of the continuous seismic signal as in the lab and Cascadia, and similar analyses that link seismicity to injection operations may help to better characterize injection systems in the future.

1.4.2.3   Earthquake catalog forecasting

A. Mignan *et. al* wrote a detailed analysis [164] of the existing neural network-based earthquake forecasting literature (77 articles from 1994 to 2019). They concluded that no new insights on earthquake predictability stemmed from these works, and demonstrated that simpler models based on classical empirical laws from statistical seismology offer similar or superior predictive power, perhaps indicating that prediction from earthquake catalogs is still a distant goal. In particular, the authors discuss the lack of baseline in the earthquake prediction literature, and show for example that on synthetic time series of earthquakes (ETAS model [165]) a neural network model does not outperform predicting mainshock magnitude in a window using the Gutenberg-Richter law (see Fig. 1.16).
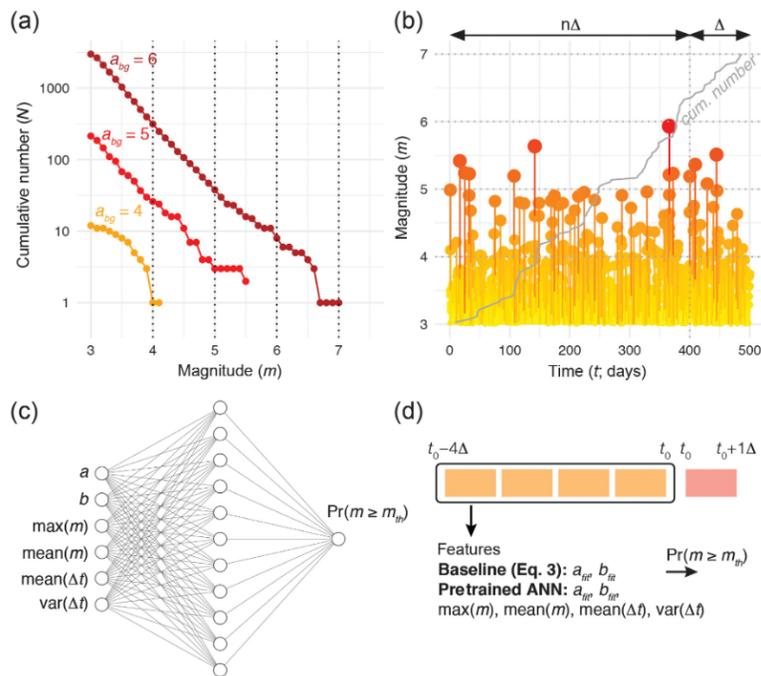
FIGURE 1.16   Neural network forecasting of earthquakes: no improvement over classical statistical seismology so far. Figure from [164] showing the approach of the authors to compare neural network-based earthquake predictions with classical statistical seismology laws as baseline. (a) The Gutenberg-Richter law as a baseline, (b) earthquake time series from the ETAS model, (c) a neural network model for earthquake prediction, (d) earthquake prediction in a window of time based on previous history.

CONCLUSION

In this review, we have attempted to present the current state of the art at the interface between fault physics and ML. From our analysis, two trends emerge. On the one hand, the field of seismology embraced ML very early on: seismology has historically been a highly data intensive discipline, and on occasions even preceded the field of ML in the development of certain algorithms. On the other hand, the field of fault physics (the why and how of rupture nucleation and propagation) has historically been much less data intensive. Earlier studies focused on the understanding of a few isolated slip events, whether in the laboratory or in the field. The shift towards ML and data science is only a very recent development, most notably for laboratory studies, for which large scale acquisition of seismic data has only recently become a reality.

As the amount of geophysical data continues increasing at ever faster rates, it also increases in complexity. Laboratory studies are recording increasing amounts of data from a variety of mechanical and seismic data from increasingly dense arrays of sensors, challenging models (such as neural networks) that can struggle with structured, heterogeneous data. The coverage of seismic networks in the field is also improving, challenging the generalization ability of models trained on particular regions, while also opening up the possibility of detecting events below the noise level of any single station. Distributed Acoustic Sensing (DAS) has recently started to be used for the study of earthquakes [166], and is poised to enable seismology at high frequency, large distances, and very fine spatial sampling (on the order of meters) – resolutions unattainable with classical seismic sensors. As a counterpart, DAS are also likely to generate staggering amounts of data. The sheer volume of currently available InSAR data already provides a challenge for existing processing methods and algorithms which sufffer from scalability issues and can require manual intervention. Next generation InSAR satellites will soon launch and dramatically increase the amount of acquisitions, making the automation of InSAR data analysis an even more pressing challenge.

As the stage is set for geophysical data to become ever harder to analyze using traditional methods, ML and data science techniques may well become fundamental tools of the trade in the geosciences. Notably, as field and laboratory studies of earthquakes and earthquake nucleation processes incorporate ML in the search for ever finer observations and slip characterization, new information may be revealed concerning fault behavior. This in turn will inform of us of the accuracy of earthquake nucleation models and analogs. Further, as ML is applied to progressively more seismic data from faults, we will gain a far better understanding regarding the deterministic versus the stochastic behavior of fault slip.

BIBLIOGRAPHY

[1] C. H. Scholz, The Mechanics of Earthquakes and Faulting, 3rd Edition, Cambridge University Press, 2019. doi:10.1017/9781316681473.

[2] K. Richards-Dinger, J. Dieterich, Rsqsim earthquake simulator, Seismological Research Letters 83 (2012) 983–990. doi:10.1785/0220120105.

[3] O. Dorostkar, R. A. Guyer, P. A. Johnson, C. Marone, J. Carmeliet, On the role of fluids in stick-slip dynamics of saturated granular fault gouge using a coupled computational fluid dynamics-discrete element approach, Journal of Geophysical Research: Solid Earthdoi:10.1002/2017JB014099.

[4] O. Dorostkar, J. Carmeliet, Potential Energy as Metric for Understanding Stick–Slip Dynamics in Sheared Granular Fault Gouge: A Coupled CFD–DEM Study, Rock Mechanics and Rock Engineering 51 (10) (2018) 3281–3294. doi:10.1007/s00603-018-1457-6.
URL http://link.springer.com/10.1007/s00603-018-1457-6

[5] O. Dorostkar, R. A. Guyer, P. A. Johnson, C. Marone, J. Carmeliet, Cohesion-induced stabilization in stick-slip dynamics of weakly wet, sheared granular fault gouge, Journal of Geophysical Research: Solid Earthdoi:10.1002/2017JB015171.

[6] M. Thomas, H. Bhat, Dynamic evolution of off-fault medium during an earthquake: A micromechanics based model, Geophysical Journal International 214 (2018) 1267–1280. doi:10.1093/GJI/GGY129.

[7] K. Gao, B. J. Euser, E. Rougier, R. A. Guyer, Z. Lei, E. E. Knight, J. Carmeliet, P. A. Johnson, Modeling of Stick-Slip Behavior in Sheared Granular Fault Gouge Using the Combined Finite-Discrete Element Method, Journal of Geophysical Research: Solid Earth 123 (7) (2018) 5774–5792, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2018JB015668. doi:10.1029/2018JB015668.
URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018JB015668

[8] K. Gao, R. Guyer, E. Rougier, C. X. Ren, P. A. Johnson, From Stress Chains to Acoustic Emission, Physical Review Letters 123 (4) (2019) 048003, publisher: American Physical Society. doi:10.1103/PhysRevLett.123.048003.
URL https://link.aps.org/doi/10.1103/PhysRevLett.123.048003

[9] C. X. Ren, O. Dorostkar, B. Rouet-Leduc, C. Hulbert, D. Strebel, R. A. Guyer, P. A. Johnson, J. Carmeliet, Machine learning reveals the state of intermittent frictional dynamics in a sheared granular fault, Geophysical Research Letters 46 (13) (2019) 7395–7403.

[10] J. Logan, Friction in rocks, Reviews of Geophysics and Space Physics 13 (1975) 358–361. doi:10.1029/RG013i003p00358.

[11] C. Marone, Laboratory-derived friction laws and their application to seismic faulting, Annual Review of Earth and Planetary Sciences 26 (1) (1998) 643–696.

[12] K. Xia, A. Rosakis, H. Kanamori, Supershear and subrayleigh to supershear transition observed in laboratory earthquake experiments, Experimental Techniques 29 (2005) 63–66.

[13] N. Beeler, Review of the physical basis of laboratory-derived relations for brittle failure and their implications for earthquake occurrence and earthquake nucleation, Pure and Applied Geophysics 161 (2004) 1853–1876. doi:10.1007/s00024-004-2536-z.

[14] N. Beeler, Inferring earthquake source properties from laboratory observations and the scope of lab contributions to source physics, Washington DC American Geophysical Union Geophysical Monograph Series 170. doi:10.1029/170GM12.

[15] B. Rouet-Leduc, C. Hulbert, N. Lubbers, K. Barros, C. J. Humphreys, P. A. Johnson, Machine Learning Predicts Laboratory Earthquakes, Geophysical Research Lettersdoi:10.1002/2017GL074677.

[16] B. Rouet-Leduc, C. Hulbert, D. C. Bolton, C. X. Ren, J. Riviere, C. Marone, R. A. Guyer, P. A. Johnson, Estimating Fault Friction From Seismic Signals in the Laboratory, Geophysical Research Letters 45 (3) (2018) 1321–1329, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2017GL076708. doi:10.1002/2017GL076708.
URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2017GL076708

[17] C. Hulbert, B. Rouet-Leduc, P. A. Johnson, C. X. Ren, J. Rivière, D. C. Bolton, C. Marone, Similarity of fast and slow earthquakes illuminated by machine learning, Nature Geoscience 12 (1) (2019) 69–74, number: 1 Publisher: Nature Publishing Group. doi:10.1038/s41561-018-0272-8.
URL https://www.nature.com/articles/s41561-018-0272-8

[18] N. Lubbers, D. C. Bolton, J. Mohd-Yusof, C. Marone, K. Barros, P. A. Johnson, Earthquake Catalog-Based Machine Learning Identification of Laboratory Fault States and the Effects of Magnitude of Completeness, Geophysical Research Letters 45 (24) (2018) 13,269–13,276. doi:10.1029/2018GL079712.
URL https://onlinelibrary.wiley.com/doi/abs/10.1029/2018GL079712

[19] A. R. Kenigsberg, J. Rivière, C. Marone, D. M. Saffer, The Effects of Shear Strain, Fabric, and Porosity Evolution on Elastic and Mechanical Properties of Clay-Rich Fault Gouge, Journal of Geophysical Research: Solid Earth 124 (11) (2019) 10968–10982, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2019JB017944. doi:10.1029/2019JB017944.
URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019JB017944

[20] D. C. Bolton, P. Shokouhi, B. Rouet-Leduc, C. Hulbert, J. Rivière, C. Marone, P. A. Johnson, Characterizing Acoustic Signals and Searching for Precursors during the Laboratory Seismic Cycle Using Unsupervised Machine Learning, Seismological Research Letters 90 (3) (2019) 1088–1098, publisher: GeoScienceWorld. doi:10.1785/0220180367.

[21] K. L. Ryan, J. Rivière, C. Marone, The Role of Shear Stress in Fault Healing and Frictional Aging, Journal of Geophysical Research: Solid Earth 123 (12) (2018) 10,479–10,495, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2018JB016296. doi:10.1029/2018JB016296.
URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2018JB016296

[22] P. A. Johnson, J. Carmeliet, H. M. Savage, M. Scuderi, B. M. Carpenter, R. A. Guyer, E. G. Daub, C. Marone, Dynamically triggered slip leading to sustained fault gouge weakening under laboratory shear conditions, Geophysical Research Letters 43 (4) (2016) 1559–1565, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1002/2015GL067056. doi:10.1002/2015GL067056.
URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/2015GL067056

[23] J. R. Leeman, D. M. Saffer, M. M. Scuderi, C. Marone, Laboratory observations of slow earthquakes and the spectrum of tectonic fault slip modes, Nature Communications 7 (1) (2016) 1–6, number: 1 Publisher: Nature Publishing Group. doi:10.1038/ncomms11104.
URL https://www.nature.com/articles/ncomms11104

[24] M. Rosenau, F. Corbi, S. Dominguez, Analogue earthquakes and seismic cycles: experimental modelling across timescales.

[25] F. Corbi, F. Funiciello, M. Moroni, Y. Van Dinther, P. M. Mai, L. Dalguer, C. Faccenna, The seismic cycle at subduction thrusts: 1. insights from laboratory models, Journal of Geophysical Research: Solid Earth 118 (4) (2013) 1483–1501.

[26] G. C. McLaskey, Earthquake initiation from laboratory observations and implications for foreshocks, Journal of Geophysical Research: Solid Earth 124 (12) (2019) 12882–12904.

[27] C. Collettini, A. Niemeijer, C. Viti, C. Marone, Fault zone fabric and fault weakness, Nature 462 (2009) 907–10. doi:10.1038/nature08585.

[28] Z. E. Ross, D. T. Trugman, E. Hauksson, P. M. Shearer, Searching for hidden earthquakes in Southern California, Science 364 (6442) (2019) 767–771, publisher: American Association for the Advancement of Science Section: Report. doi:10.1126/science.aaw6888.
URL https://science.sciencemag.org/content/364/6442/767

[29] Z. E. Ross, M.-A. Meier, E. Hauksson, P Wave Arrival Picking and First-Motion Polarity Determination With Deep Learning, Journal of Geophysical Research: Solid Earth 123 (6) (2018) 5120–5129, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2017JB015251. doi:10.1029/2017JB015251.
URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2017JB015251

[30] Z. E. Ross, M.-A. Meier, E. Hauksson, T. H. Heaton, Generalized Seismic Phase Detection with Deep Learning, Bulletin of the Seismological Society of America 108 (5A) (2018) 2894–2901, publisher: GeoScienceWorld. doi:10.1785/0120180080.

[31] D. T. Trugman, Z. E. Ross, Pervasive Foreshock Activity Across Southern California, Geophysical Research Letters 46 (15) (2019) 8772–8781, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2019GL083725. doi:10.1029/2019GL083725.
URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019GL083725

[32] F. Trippetta, P. Petricca, A. Billi, C. Collettini, M. Cuffaro, A. Lombardi, D. Scrocca, G. Ventura, A. Morgante, C. Doglioni, From mapped faults to fault-length earthquake magnitude (flem): a test on italy with methodological implications, Solid Earth 10 (2019) 1555–1579. doi:10.5194/se-10-1555-2019.

[33] D. T. Trugman, Z. E. Ross, P. A. Johnson, Imaging stress and faulting complexity through earthquake waveform similarity, Geophysical Research Letters 47 (1) (2020) e2019GL085888. doi:10.1029/2019GL085888.
URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019GL085888

[34] B. Rouet-Leduc, C. Hulbert, I. W. McBrearty, P. A. Johnson, Probing slow earthquakes with deep learning, Geophysical Research Letters 47 (4) (2020) e2019GL085870, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2019GL085870. doi:10.1029/2019GL085870.
URL https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019GL085870

[35] E. E. Brodsky, J. J. Mori, L. Anderson, F. M. Chester, M. Conin, E. M. Dunham, N. Eguchi, P. M. Fulton, R. Hino, T. Hirose, M. J. Ikari, T. Ishikawa, T. Jeppson, Y. Kano, J. Kirkpatrick, S. Kodaira, W. Lin, Y. Nakamura, H. S. Rabinowitz, C. Regalla, F. Remitti, C. Rowe, D. M. Saffer, S. Saito, J. Sample, Y. Sanada, H. M. Savage, T. Sun, S. Toczko, K. Ujiie, M. Wolfson-Schwehr, T. Yang, The state of stress on the fault before, during, and after a major earthquake, Annual Review of Earth and Planetary Sciences 48 (1) (2020) null. arXiv:https://doi.org/10.1146/annurev-earth-053018-060507, doi:10.1146/annurev-earth-053018-060507.
URL https://doi.org/10.1146/annurev-earth-053018-060507

[36] K. J. Bergen, P. A. Johnson, M. V. de Hoop, G. C. Beroza, Machine learning for data-driven discovery in solid Earth geoscience, Science 363 (6433) (2019) eaau0323. doi:10.1126/science.aau0323.
URL http://www.sciencemag.org/lookup/doi/10.1126/science.aau0323

[37] Q. Kong, D. T. Trugman, Z. E. Ross, M. J. Bianco, B. J. Meade, P. Gerstoft, Machine Learning in Seismology: Turning Data into Insights, Seismological Research Letters 90 (1)

(2019) 3–14, publisher: GeoScienceWorld. doi:10.1785/0220180259.

[38] G. Matheron, Traité de géostatistique appliquée. 1 (1962), Vol. 1, Editions Technip, 1962.

[39] F. Santosa, W. W. Symes, Linear inversion of band-limited reflection seismograms, SIAM Journal on Scientific and Statistical Computing 7 (4) (1986) 1307–1330.

[40] R. Tibshirani, Regression shrinkage and selection via the lasso, Journal of the Royal Statistical Society: Series B (Methodological) 58 (1) (1996) 267–288.

[41] I. Goodfellow, Y. Bengio, A. Courville, Deep learning, MIT press, 2016.

[42] Y. S. Abu-Mostafa, M. Magdon-Ismail, H.-T. Lin, Learning from data, Vol. 4, AMLBook New York, NY, USA:, 2012.

[43] K. J. Bergen, P. A. Johnson, M. V. d. Hoop, G. C. Beroza, Machine learning for data-driven discovery in solid Earth geoscience, Science 363 (6433), publisher: American Association for the Advancement of Science Section: Review. doi:10.1126/science.aau0323.
URL https://science.sciencemag.org/content/363/6433/eaau0323

[44] T. M. Mitchell, et al., Machine learning (1997).

[45] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: Machine learning in Python, Journal of Machine Learning Research 12 (2011) 2825–2830.

[46] P. J. Huber, Robust estimation of a location parameter: Annals mathematics statistics, 35.

[47] Y. S. Abu-Mostafa, M. Magdon-Ismail, H. T. Lin, Learning from Data: A Short Course, 2012. doi:10.1108/17538271011063889.

[48] C. Bergmeir, J. M. Benítez, On the use of cross-validation for time series predictor evaluation, Information Sciences 191 (2012) 192–213. doi:10.1016/j.ins.2011.12.028.

[49] I. W. McBrearty, A. A. Delorey, P. A. Johnson, Pairwise Association of Seismic Arrivals with Convolutional Neural Networks 7.

[50] B. Rouet-Leduc, C. Hulbert, P. A. Johnson, Continuous chatter of the Cascadia subduction zone revealed by machine learning , Nature Geoscience 12 (1) (2019) 908–1752. doi:10.1038/s41561-018-0274-6.
URL http://www.nature.com/articles/s41561-018-0274-6

[51] W. F. Brace, J. D. Byerlee, Stick-Slip as a Mechanism for Earthquakes, Science 153 (3739) (1966) 990–992. doi:10.1126/science.153.3739.990.

[52] J. H. Dieterich, Modeling of rock friction: 1. Experimental results and constitutive equations, Journal of Geophysical Research: Solid Earth 84 (B5) (1979) 2161–2168.

[53] A. Ruina, Slip instability and state variable friction laws, Journal of Geophysical Research: Solid Earth 88 (B12) (1983) 10359–10370. doi:10.1029/JB088iB12p10359.
URL http://doi.wiley.com/10.1029/JB088iB12p10359

[54] C. Marone, Laboratory-derived friction laws and their application to seismic faulting, Annual Review of Earth and Planetary Sciences 26 (1) (1998) 643–696.

[55] E. Daub, D. Shelly, R. Guyer, P. Johnson, Brittle and ductile friction and the physics of tectonic tremor, Geophysical Research Letters 38. doi:10.1029/2011GL046866.

[56] J. Rice, N. Lapusta, R. Kunnath, Rate and state dependent friction and the stability of sliding between elastically deformable solids, Journal of the Mechanics and Physics of Solids 49 (2001) 1865–1898. doi:10.1016/S0022-5096(01)00042-4.

[57] T. Parker, S. Shatalin, M. Farhadiroushan, Distributed acoustic sensing - a new tool for seismic applications, First Break 32. doi:10.3997/1365-2397.2013034.

[58] P. Johnson, B. Ferdowsi, B. Kaproth, M. Scuderi, M. Griffa, J. Carmeliet, R. Guyer, P.-Y. Le Bas, D. Trugman, C. Marone, Acoustic emission and microslip precursors to stick-slip failure in sheared granular material, Geophysical Research Letters 40 (21) (2013) 5627–5631.

[59] S. Latour, A. Schubnel, S. Nielsen, R. Madariaga, S. Vinciguerra, Characterization of nucleation during laboratory earthquakes, Geophysical Research Letters 40 (19) (2013) 5064–5069. doi:10.1002/grl.50974.
URL https://onlinelibrary.wiley.com/doi/abs/10.1002/grl.50974

[60] J. R. Leeman, D. M. Saffer, M. M. Scuderi, C. Marone, Laboratory observations of slow earthquakes and the spectrum of tectonic fault slip modes, Nature communications 7.

[61] J. Adam, J. Urai, B. Wieneke, O. Oncken, K. Pfeiffer, N. Kukowski, J. Lohrmann, S. Hoth, W. Van Der Zee, J. Schmatz, Shear localisation and strain distribution during tectonic faulting—new insights from granular-flow experiments and high-resolution optical image correlation techniques, Journal of Structural Geology 27 (2) (2005) 283–301.

[62] T. Majmudar, R. Behringer, Majmudar, t. s. and behringer, r. p. contact force measurements and stress-induced anisotropy in granular materials. nature, 435, 1079-1082, Nature 435 (2005) 1079–82. doi:10.1038/nature03805.

[63] K. Daniels, J. Kollmer, J. Puckett, Photoelastic force measurements in granular materials, Review of Scientific Instruments 88. doi:10.1063/1.4983049.

[64] F. Corbi, L. Sandri, J. Bedford, F. Funiciello, S. Brizzi, M. Rosenau, S. Lallemand, Machine learning can predict the timing and size of analog earthquakes, Geophysical Research Letters 46 (3) (2019) 1303–1311.

[65] J. H. Friedman, Greedy function approximation: a gradient boosting machine, Annals of statistics (2001) 1189–1232.

[66] F. Corbi, J. Bedford, L. Sandri, F. Funiciello, A. Gualandi, M. Rosenau, Predicting imminence of analog megathrust earthquakes with machine learning: Implications for monitoring subduction zones, Geophysical Research Letters (2020) e2019GL086615.

[67] C. Seiffert, T. M. Khoshgoftaar, J. Van Hulse, A. Napolitano, Rusboost: Improving classification performance when training data is skewed, in: 2008 19th International Conference on Pattern Recognition, IEEE, 2008, pp. 1–4.

[68] Y. Freund, R. Schapire, N. Abe, A short introduction to boosting, Journal-Japanese Society For Artificial Intelligence 14 (771-780) (1999) 1612.

[69] H. Jasperson, C. Bolton, P. Johnson, C. Marone, M. Hoop, Unsupervised classification of acoustic emissions from catalogs and fault time-to-failure prediction (12 2019).

[70] LANL Earthquake Prediction | Kaggle.
URL https://www.kaggle.com/c/LANL-Earthquake-Prediction

[71] A. Niemeijer, C. Marone, D. Elsworth, Frictional strength and strain weakening in simulated fault gouge: Competition between geometrical weakening and chemical strengthening, Journal of Geophysical Research: Solid Earth 115 (B10).

[72] M. M. Scuderi, C. Marone, E. Tinti, G. D. Stefano, C. Collettini, Precursory changes in seismic velocity for the spectrum of earthquake failure modes, Nature Geoscience 9 (9) (2016) 695–700, number: 9 Publisher: Nature Publishing Group. doi:10.1038/ngeo2775.
URL https://www.nature.com/articles/ngeo2775

[73] L. Breiman, Random forests, Machine learning 45 (1) (2001) 5–32.

[74] L. Breiman, Classification and Regression Trees., Routledge, New York:, 1984.
URL https://doi.org/10.1201/9781315139470

[75] C. Hulbert, B. Rouet-Leduc, P. A. Johnson, C. X. Ren, J. Rivière, D. C. Bolton, C. Marone, Similarity of fast and slow earthquakes illuminated by machine learning, Nature Geosciencedoi:10.1038/s41561-018-0272-8.

[76] P. A. Johnson, B. Ferdowsi, B. M. Kaproth, M. Scuderi, M. Griffa, J. Carmeliet, R. A. Guyer, P.-Y. Le Bas, D. T. Trugman, C. Marone, Acoustic emission and microslip precursors to stick-slip failure in sheared granular material, Geophysical Research Letters 40 (21) (2013)

5627–5631.

[77] J. H. Dieterich, Preseismic fault slip and earthquake prediction, Journal of Geophysical Research 83 (B8) (1978) 3940. doi:10.1029/jb083ib08p03940.

[78] M. Bouchon, V. Durand, D. Marsan, H. Karabulut, J. Schmittbuhl, The long precursory phase of most large interplate earthquakes, Nature Geoscience 6 (4) (2013) 299–302. doi:10.1038/ngeo1770.
URL http://www.nature.com/articles/ngeo1770

[79] S. Beckouche, J. Ma, Simultaneous dictionary learning and denoising for seismic data, Geophysics 79 (3) (2014) A27–A31. doi:10.1190/GEO2013-0382.1.

[80] G. C. Beroza, W. L. Ellsworth, Properties of the seismic nucleation phase, Tectonophysics 261 (1-3 SPEC. ISS.) (1996) 209–227. doi:10.1016/0040-1951(96)00067-4.

[81] G. C. McLaskey, Earthquake Initiation From Laboratory Observations and Implications for Foreshocks, Journal of Geophysical Research: Solid Earth 124 (12) (2019) 12882–12904. doi:10.1029/2019JB018363.
URL https://onlinelibrary.wiley.com/doi/abs/10.1029/2019JB018363

[82] Y. Ito, R. Hino, M. Kido, H. Fujimoto, Y. Osada, D. Inazu, Y. Ohta, T. Iinuma, M. Ohzono, S. Miura, Others, Episodic slow slip events in the Japan subduction zone before the 2011 Tohoku-Oki earthquake, Tectonophysics 600 (2013) 14–26.

[83] S. Ruiz, M. Metois, A. Fuenzalida, J. Ruiz, F. Leyton, R. Grandin, C. Vigny, R. Madariaga, J. Campos, Intense foreshocks and a slow slip event preceded the 2014 Iquique Mw 8.1 earthquake, Science 345 (6201) (2014) 1165–1169. doi:10.1126/science.1256074.
URL http://www.ncbi.nlm.nih.gov/pubmed/25061132

[84] A. Kato, K. Obara, T. Igarashi, H. Tsuruoka, S. Nakagawa, N. Hirata, Propagation of Slow Slip Leading Up to the 2011 Mw 9.0 Tohoku-Oki Earthquake, Science 335 (6069) (2012) 705–708. doi:10.1126/science.1215141.

[85] F. U. Dowla, S. R. Taylor, R. W. Anderson, Seismic discrimination with artificial neural networks: preliminary results with regional spectral data, Bulletin - Seismological Society of America 80 (5) (1990) 1346–1373.

[86] P. S. Dysart, J. J. Pulli, Regional seismic event classification at the NORESS array: seismological measurements and the use of trained neural networks, Bulletin - Seismological Society of America 80 (6 B) (1991) 1910–1933.

[87] Jin Wang, Ta-Liang Teng, Artificial neural network-based seismic detector, Bulletin - Seismological Society of America 85 (1) (1995) 308–319. doi:10.1016/0148-9062(96)86904-x.

[88] T. Tiira, Detecting teleseismic events using artificial neural networks, Computers and Geosciences 25 (8) (1999) 929–938. doi:10.1016/S0098-3004(99)00056-4.

[89] G. Madureira, A. E. Ruano, A neural network seismic detector, in: IFAC Proceedings Volumes (IFAC-PapersOnline), Vol. 2, Elsevier, 2009, pp. 304–309. doi:10.3182/20090921-3-TR-3005.00054.

[90] A. E. Ruano, G. Madureira, O. Barros, H. R. Khosravani, M. G. Ruano, P. M. Ferreira, Seismic detection using support vector machines, Neurocomputing 135 (2014) 273–283. doi:10.1016/j.neucom.2013.12.020.

[91] A. Reynen, P. Audet, Supervised machine learning on a network scale: application to seismic event classification and detection, Geophysical Journal International 210 (3) (2017) 1394–1409. doi:10.1093/gji/ggx238.
URL https://doi.org/10.1093/gji/ggx238

[92] S. M. Mousavi, S. P. Horton, C. A. Langston, B. Samei, Seismic features and automatic discrimination of deep and shallow induced-microearthquakes using neural network and logistic regression, Geophysical Journal International 207 (1) (2016) 29–46. doi:10.1093/gji/ggw258.

URL `https://doi.org/10.1093/gji/ggw258`

[93] Y. Freund, R. Schapire, Large margin classification using the perceptron algorithm, Machine Learning 37. doi:10.1023/A:1007662407062.

[94] M. Beyreuther, J. Wassermann, Continuous earthquake detection and classification using discrete hidden markov models, Geophysical Journal International 175 (3) (2008) 1055–1066. doi:10.1111/j.1365-246X.2008.03921.x.
URL `https://academic.oup.com/gji/article-lookup/doi/10.1111/j.1365-246X.2008.03921.x`

[95] C. Riggelsen, M. Ohrnberger, A Machine Learning Approach for Improving the Detection Capabilities at 3C Seismic Stations, Pure and Applied Geophysics 171 (3-5) (2014) 395–411. doi:10.1007/s00024-012-0592-3.

[96] A. C. Aguiar, G. C. Beroza, PageRank for earthquakes, Seismological Research Letters 85 (2) (2014) 344–350. doi:10.1785/0220130162.

[97] C. E. Yoon, O. O'Reilly, K. J. Bergen, G. C. Beroza, O. O'Reilly, K. J. Bergen, G. C. Beroza, Earthquake detection through computationally efficient similarity search, Science Advances 1 (11) (2015) e1501057. doi:10.1126/sciadv.1501057.
URL `https://advances.sciencemag.org/content/1/11/e1501057`

[98] B. Sick, M. Guggenmos, M. Joswig, Chances and limits of single-station seismic event clustering by unsupervised pattern recognition, Geophysical Journal International 201 (3) (2015) 1801–1813. doi:10.1093/gji/ggv126.

[99] J. Wiszniowski, B. M. Plesiewicz, J. Trojanowski, Application of real time recurrent neural network for detection of small natural earthquakes in Poland, Acta Geophysica 62 (3) (2014) 469–485. doi:10.2478/s11600-013-0140-2.

[100] T. Perol, M. Gharbi, M. Denolle, Convolutional neural network for earthquake detection and location, Science Advances 4 (2) (2018) e1700578. arXiv:1702.02073, doi:10.1126/sciadv.1700578.

[101] S. M. Mousavi, W. Zhu, Y. Sheng, G. C. Beroza, CRED: A Deep Residual Network of Convolutional and Recurrent Units for Earthquake Signal Detection, Scientific Reports 9 (1) (2019) 1–14. arXiv:1810.01965, doi:10.1038/s41598-019-45748-1.

[102] A. Lomax, A. Michelini, D. Jozinović, An Investigation of Rapid Earthquake Characterization Using Single-Station Waveforms and a Convolutional Neural Network, Seismological Research Letters 90 (2 A) (2019) 517–529. doi:10.1785/0220180311.

[103] C. Gu, Y. M. Marzouk, M. N. Toksoz, Bayesian deep learning and uncertainty quantification applied to induced seismicity locations in the Groningen gas field in the Netherlands: What do we need for safe AI? (2019).
URL `https://doi.org/`

[104] S. M. Mousavi, W. Zhu, Y. Sheng, G. C. Beroza, CRED: A Deep Residual Network of Convolutional and Recurrent Units for Earthquake Signal Detection, Scientific Reports 9 (1) (2019) 1–14, number: 1 Publisher: Nature Publishing Group. doi:10.1038/s41598-019-45748-1.
URL `https://www.nature.com/articles/s41598-019-45748-1`

[105] M. E. MURAT, A. J. RUDMAN, AUTOMATED FIRST ARRIVAL PICKING: A NEURAL NETWORK APPROACH, Geophysical Prospecting 40 (6) (1992) 587–604. doi:10.1111/j.1365-2478.1992.tb00543.x.
URL `http://doi.wiley.com/10.1111/j.1365-2478.1992.tb00543.x`

[106] M. D. McCormack, D. E. Zaucha, D. W. Dushek, First-break refraction event picking and seismic data trace editing using neural networks, Geophysics 58 (1) (1993) 67–78. doi:10.1190/1.1443352.

[107] H. Dai, C. MacBeth, Automatic picking of seismic arrivals in local earthquake data using

an artificial neural network, Geophysical Journal International 120 (3) (1995) 758–774. doi:10.1111/j.1365-246X.1995.tb01851.x.
URL https://academic.oup.com/gji/article-lookup/doi/10.1111/j.1365-246X.1995.tb01851.x

[108] H. Dai, C. MacBeth, The application of back-propagation neural network to automatic picking seismic arrivals from single-component recordings, Journal of Geophysical Research: Solid Earth 102 (B7) (1997) 15105–15113. doi:10.1029/97jb00625.

[109] J. Wang, T.-l. Teng, Identification and picking of S phase using an artificial neural network, Bulletin of the Seismological Society of America 87 (5) (1997) 1140–1149.

[110] Y. Zhao, K. Takano, An artificial neural network approach for broadband seismic phase picking, Bulletin of the Seismological Society of America 89 (3) (1999) 670–680.

[111] S. Gentili, A. Michelini, Automatic picking of P and S phases using a neural tree, Journal of Seismology 10 (1) (2006) 39–63. doi:10.1007/s10950-006-2296-6.

[112] Y. Chen, Automatic microseismic event picking via unsupervised machine learning, Geophysical Journal International 212 (1) (2017) 88–102. doi:10.1093/gji/ggx420.
URL https://doi.org/10.1093/gji/ggx420

[113] L. Zhu, Z. Peng, J. McClellan, C. Li, D. Yao, Z. Li, L. Fang, Deep learning for seismic phase detection and picking in the aftershock zone of 2008 Mw7.9 Wenchuan Earthquake, Physics of the Earth and Planetary Interiors 293 (2019) 106261. arXiv:1901.06396, doi:10.1016/j.pepi.2019.05.004.

[114] W. Zhu, G. C. Beroza, {PhaseNet}: a deep-neural-network-based seismic arrival-time picking method, Geophysical Journal International 216 (1) (2019) 261–273. doi:10.1093/gji/ggy423.
URL https://academic.oup.com/gji/article/216/1/261/5129142 https://doi.org/10.1093/gji/ggy423

[115] K. Wang, Y. Li, Q. Kemao, J. Di, J. Zhao, One-step robust deep learning phase unwrapping, Optics Express 27 (10) (2019) 15100. doi:10.1364/OE.27.015100.
URL https://www.osapublishing.org/abstract.cfm?URI=oe-27-10-15100

[116] Z. E. Ross, M. A. Meier, E. Hauksson, T. H. Heaton, Generalized seismic phase detection with deep learning, Bulletin of the Seismological Society of America 108 (5) (2018) 2894–2901. arXiv:1805.01075, doi:10.1785/0120180080.

[117] S. Hara, Y. Fukahata, Y. Iio, P-wave first-motion polarity determination of waveform data in western Japan using deep learning, Earth, Planets and Space 71 (1) (2019) 1–11. doi:10.1186/s40623-019-1111-x.

[118] I. W. McBrearty, J. Gomberg, A. A. Delorey, P. A. Johnson, Earthquake Arrival Association with Backprojection and Graph TheoryEarthquake Arrival Association with Backprojection and Graph Theory, Bulletin of the Seismological Society of America 109 (6) (2019) 2510–2531, publisher: GeoScienceWorld. doi:10.1785/0120190081.

[119] Z. E. Ross, Y. Yue, M. A. Meier, E. Hauksson, T. H. Heaton, PhaseLink: A Deep Learning Approach to Seismic Phase Association, Journal of Geophysical Research: Solid Earth 124 (1) (2019) 856–869. arXiv:1809.02880, doi:10.1029/2018JB016674.
URL https://onlinelibrary.wiley.com/doi/abs/10.1029/2018JB016674

[120] N. Riahi, P. Gerstoft, Using graph clustering to locate sources within a dense sensor array, Signal Processing 132 (2017) 110–120. doi:10.1016/j.sigpro.2016.10.001.

[121] D. T. Trugman, P. M. Shearer, GrowClust: A Hierarchical clustering algorithm for relative earthquake relocation, with application to the Spanish Springs and Sheldon, Nevada, earthquake sequences, Seismological Research Letters 88 (2) (2017) 379–391. doi:10.1785/0220160188.

[122] M. Kriegerowski, G. M. Petersen, H. Vasyura-Bathke, M. Ohrnberger, A deep convolutional neural network for localization of clustered earthquakes based on multistation full waveforms,

Seismological Research Letters 90 (2 A) (2019) 510–516. doi:10.1785/0220180320.

[123] X. Zhang, J. Zhang, C. Yuan, S. Liu, Z. Chen, W. Li, Locating induced earthquakes with a network of seismic stations in Oklahoma via a deep learning method, Scientific Reports 10 (1) (2020) 1–12. doi:10.1038/s41598-020-58908-5.

[124] S. M. Mousavi, G. C. Beroza, Bayesian-Deep-Learning Estimation of Earthquake Location from Single-Station ObservationsarXiv:1912.01144.
URL http://arxiv.org/abs/1912.01144

[125] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, ACM, 2016, pp. 785–794.

[126] A. P. Valentine, J. Trampert, Data space reduction, quality assessment and searching of seismograms: Autoencoder networks for waveform data, Geophysical Journal International 189 (2) (2012) 1183–1202. doi:10.1111/j.1365-246X.2012.05429.x.
URL https://academic.oup.com/gji/article-lookup/doi/10.1111/j.1365-246X.2012.05429.x

[127] B. Rouet-Leduc, C. Hulbert, I. W. McBrearty, P. A. Johnson, Probing Slow Earthquakes With Deep Learning, Geophysical Research Letters 47 (4). doi:10.1029/2019GL085870.
URL https://onlinelibrary.wiley.com/doi/abs/10.1029/2019GL085870

[128] H. Sun, L. Demanet, Extrapolated full waveform inversion with deep learning, GEOPHYSICS (2020) 1–71arXiv:1909.11536, doi:10.1190/geo2019-0195.1.

[129] G. Rogers, H. Dragert, Episodic tremor and slip on the Cascadia subduction zone: The chatter of silent slip, Science 300 (5627) (2003) 1942–1943.

[130] C. Hulbert, B. Rouet-Leduc, P. A. Johnson, A Silent Build-up in Seismic Energy Precedes Slow Slip Failure in the Cascadia Subduction ZonearXiv:1909.06787.
URL http://arxiv.org/abs/1909.06787

[131] M. Nakano, D. Sugiyama, T. Hori, T. Kuwatani, S. Tsuboi, Discrimination of Seismic Signals from Earthquakes and Tectonic Tremor by Applying a Convolutional Neural Network to Running Spectral Images, Seismological Research Letters 90 (2 A) (2019) 530–538. doi:10.1785/0220180279.
URL https://pubs.geoscienceworld.org/ssa/srl/article/90/2A/530/567979/Discrimination-of-Seismic-Signals-fr

[132] A. Rietbrock, I. Ryder, G. Hayes, C. Haberland, D. Comte, S. Roecker, H. Lyon-Caen, Aftershock seismicity of the 2010 Maule Mw=8.8, Chile, earthquake: Correlation between co-seismic slip models and aftershock distribution?, Geophysical Research Letters 39 (8) (2012) n/a–n/a. doi:10.1029/2012GL051308.
URL http://doi.wiley.com/10.1029/2012GL051308

[133] B. Rousset, R. Jolivet, M. Simons, C. Lasserre, B. Riel, P. Milillo, Z. Çakir, F. Renard, An aseismic slip transient on the North Anatolian Fault, Geophysical Research Letters 43 (7) (2016) 3254–3262. doi:10.1002/2016GL068250.
URL   https://onlinelibrary.wiley.com/doi/abs/10.1002/2016GL068250
http://doi.wiley.com/10.1002/2016GL068250

[134] P. Käufl, A. P. Valentine, T. B. O'Toole, J. Trampert, A framework for fast probabilistic centroid-moment-tensor determination-inversion of regional static displacement measurements, Geophysical Journal International 196 (3) (2014) 1676–1693. doi:10.1093/gji/ggt473.

[135] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, N. D. Goodman, Pyro: Deep Universal Probabilistic Programming, J. Mach. Learn. Res. 20 (1) (2019) 973–978.

[136] D. Massonnet, M. Rossi, C. Carmona, F. Adragna, G. Peltzer, K. Feigl, T. Rabaute, The displacement field of the Landers earthquake mapped by radar interferometry, Nature 364 (6433) (1993) 138–142. doi:10.1038/364138a0.

URL http://www.nature.com/articles/364138a0

[137] D. Massonnet, P. Briole, A. Arnaud, Deflation of Mount Etna monitored by spaceborne radar interferometry, Nature 375 (6532) (1995) 567–570. doi:10.1038/375567a0.
URL http://www.nature.com/articles/375567a0

[138] G. Peltzer, F. Crampé, G. King, Evidence of nonlinear elasticity of the crust from the Mw7.6 Manyi (Tibet) earthquake, Science 286 (5438) (1999) 272–276. doi:10.1126/science.286.5438.272.
URL http://www.ncbi.nlm.nih.gov/pubmed/10514367

[139] M. Simons, Y. Fialko, L. Rivera, Coseismic deformation from the 1999 Mw 7.1 Hector Mine, California, earthquake as inferred from InSAR and GPS observations, Bulletin of the Seismological Society of America 92 (4) (2002) 1390–1402. doi:10.1785/0120000933.
URL https://pubs.geoscienceworld.org/bssa/article/92/4/1390-1402/120788

[140] G. Peltzer, F. Crampé, S. Hensley, P. Rosen, Transient strain accumulation and fault interaction in the Eastern California shear zone, Geology 29 (11) (2001) 975–978. doi:10.1130/0091-7613(2001)029<0975:TSAAFI>2.0.CO;2.
URL https://pubs.geoscienceworld.org/geology/article/29/11/975-978/197897

[141] T. J. Wright, B. Parsons, P. C. England, E. J. Fielding, InSAR observations of low slip rates on the major faults of western Tibet, Science 305 (5681) (2004) 236–239. doi:10.1126/science.1096388.
URL http://www.ncbi.nlm.nih.gov/pubmed/15247476

[142] R. Jolivet, C. Lasserre, M. P. Doin, S. Guillaso, G. Peltzer, R. Dailu, J. Sun, Z. K. Shen, X. Xu, Shallow creep on the Haiyuan fault (Gansu, China) revealed by SAR interferometry, Journal of Geophysical Research: Solid Earth 117 (6) (2012) n/a–n/a. doi:10.1029/2011JB008732.
URL http://doi.wiley.com/10.1029/2011JB008732

[143] R. Jolivet, T. Candela, C. Lasserre, F. Renard, Y. Klinger, M. P. Doin, The burst-like behavior of aseismic slip on a rough fault: The creeping section of the Haiyuan fault, China, Bulletin of the Seismological Society of America 105 (1) (2015) 480–488. doi:10.1785/0120140237.
URL https://pubs.geoscienceworld.org/bssa/article/105/1/480-488/323642

[144] G. E. Hilley, R. Bürgmann, A. Ferretti, F. Novali, F. Rocca, Dynamics of slow-moving landslides from permanent scatterer analysis, Science 304 (5679) (2004) 1952–1955. doi:10.1126/science.1098821.
URL http://www.ncbi.nlm.nih.gov/pubmed/15218146

[145] O. Cavalié, E. Pathier, M. Radiguet, M. Vergnolle, N. Cotte, A. Walpersdorf, V. Kostoglodov, F. Cotton, Slow slip event in the Mexican subduction zone: Evidence of shallower slip in the Guerrero seismic gap for the 2006 event revealed by the joint inversion of InSAR and GPS data, Earth and Planetary Science Letters 367 (2013) 52–60. doi:10.1016/j.epsl.2013.02.020.
URL https://www.sciencedirect.com/science/article/pii/S0012821X13000873

[146] R. Jolivet, C. Lasserre, M. P. Doin, G. Peltzer, J. P. Avouac, J. Sun, R. Dailu, Spatio-temporal evolution of aseismic slip along the Haiyuan fault, China: Implications for fault frictional properties, Earth and Planetary Science Letters 377-378 (2013) 23–33. doi:10.1016/j.epsl.2013.07.020.
URL https://www.sciencedirect.com/science/article/pii/S0012821X13003956

[147] N. Anantrasirichai, J. Biggs, F. Albino, D. Bull, A deep learning approach to detecting volcano deformation from satellite imagery using synthetic datasets, Remote Sensing of Environment 230 (2019) 111179. arXiv:1905.07286, doi:10.1016/J.RSE.2019.04.032.
URL https://www.sciencedirect.com/science/article/pii/S003442571930183X?via%3Dihub

[148] A. Benoit, B. Pinel-Puysségur, R. Jolivet, C. Lasserre, CorPhU: an algorithm based on phase closure for the correction of unwrapping errors in SAR interferometry, Geophysical Journal

International 221 (3) (2020) 1959–1970. doi:10.1093/gji/ggaa120.
URL https://doi.org/10.1093/gji/ggaa120

[149] N. Anantrasirichai, F. Albino, P. Hill, D. Bull, J. Biggs, Detecting Volcano Deformation in InSAR using Deep learningarXiv:1803.00380.
URL http://arxiv.org/abs/1803.00380

[150] N. Anantrasirichai, J. Biggs, F. Albino, D. Bull, The Application of Convolutional Neural Networks to Detect Slow, Sustained Deformation in InSAR Time Series, Geophysical Research Letters 46 (21) (2019) 11850–11858. doi:10.1029/2019GL084993.
URL https://onlinelibrary.wiley.com/doi/abs/10.1029/2019GL084993

[151] A. Krizhevsky, I. Sutskever, G. E. Hinton, ImageNet Classification with Deep Convolutional Neural Networks (jan 2012).
URL http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networ http://papers.nips.cc/paper/4824-imagenet-classification-w papers3://publication/uuid/1ECF396A-CEDA-45CD-9A9F-03344449DA2A

[152] B. Rouet-Leduc, M. Dalaison, C. Hulbert, P. Johnson, R. Jolivet, A deep learning approach for detecting transient deformation in insar.

[153] M. Böse, F. Wenzel, M. Erdik, PreSEIS: A neural network-based approach to earthquake early warning for finite faults, Bulletin of the Seismological Society of America 98 (1) (2008) 366–382. doi:10.1785/0120070002.

[154] L. H. Ochoa, L. F. Niño, C. A. Vargas, Fast magnitude determination using a single seismological station record implementing machine learning techniques, Geodesy and Geodynamics 9 (1) (2018) 34–41. doi:10.1016/j.geog.2017.03.010.

[155] Q. Kong, R. M. Allen, L. Schreier, Y. W. Kwon, Earth Sciences: MyShake: A smartphone seismic network for earthquake early warning and beyond, Science Advances 2 (2) (2016) e1501055. doi:10.1126/sciadv.1501055.

[156] Z. Li, M. A. Meier, E. Hauksson, Z. Zhan, J. Andrews, Machine Learning Seismic Wave Discrimination: Application to Earthquake Early Warning, Geophysical Research Letters 45 (10) (2018) 4773–4779. doi:10.1029/2018GL077870.
URL http://doi.wiley.com/10.1029/2018GL077870

[157] J. Schmidhuber, Learning Factorial Codes by Predictability Minimization, Neural Computation 4 (6) (1992) 863–879. doi:10.1162/neco.1992.4.6.863.

[158] S. D. Goodfellow, R. P. Young, A laboratory acoustic emission experiment under in situ conditions, Geophysical Research Letters 41 (10) (2014) 3422–3430. doi:10.1002/2014GL059965.
URL http://doi.wiley.com/10.1002/2014GL059965

[159] M. A. Meier, T. Heaton, J. Clinton, The Gutenberg algorithm: Evolutionary Bayesian magnitude estimates for earthquake early warning with a filter bank, Bulletin of the Seismological Society of America 105 (5) (2015) 2774–2786. doi:10.1785/0120150098.

[160] M. Beyreuther, C. Hammer, J. Wassermann, M. Ohrnberger, T. Megies, Constructing a Hidden Markov Model based earthquake detector: application to induced seismicity, Geophysical Journal International 189 (1) (2012) 602–610. doi:10.1111/j.1365-246X.2012.05361.x.
URL https://doi.org/10.1111/j.1365-246X.2012.05361.x

[161] B. K. Holtzman, A. Paté, J. Paisley, F. Waldhauser, D. Repetto, Machine learning reveals cyclic changes in seismic source spectra in Geysers geothermal field, Science Advances 4 (5) (2018) eaao2929. doi:10.1126/sciadv.aao2929.
URL http://advances.sciencemag.org/lookup/doi/10.1126/sciadv.aao2929

[162] T. Hincks, W. Aspinall, R. Cooke, T. Gernon, Oklahoma's induced seismicity strongly linked to wastewater injection depth, Science 359 (6381) (2018) 1251–1255. doi:10.1126/science.aap7911.

[163] F. R. Walsh, M. D. Zoback, Oklahoma's recent earthquakes and saltwater disposal, Science Advances 1 (5). arXiv:https://advances.sciencemag.org/content/1/5/e1500195.full.pdf, doi:10.1126/sciadv.1500195.
URL https://advances.sciencemag.org/content/1/5/e1500195

[164] A. Mignan, M. Broccardo, Neural Network Applications in Earthquake Prediction (1994–2019): Meta-Analytic and Statistical Insights on Their Limitations, Seismological Research Lettersdoi:10.1785/0220200021.
URL https://pubs.geoscienceworld.org/ssa/srl/article/586519/Neural-Network-Applications-in-Earthquake

[165] Y. Ogata, Statistical models for earthquake occurrences and residual analysis for point processes, Journal of the American Statistical Association 83 (401) (1988) 9–27. doi:10.1080/01621459.1988.10478560.

[166] N. J. Lindsey, E. R. Martin, D. S. Dreger, B. Freifeld, S. Cole, S. R. James, B. L. Biondi, J. B. Ajo-Franklin, Fiber-Optic Network Observations of Earthquake Wavefields, Geophysical Research Letters 44 (23) (2017) 11,792–11,799. doi:10.1002/2017GL075722.
URL http://doi.wiley.com/10.1002/2017GL075722