LA-UR-18-28089 (Accepted Manuscript)

Machine Learning in Seismology: Turning Data into Insights

Kong, Qinkai
Trugman, Daniel Taylor
Ross, Zachary E.
Bianco, Michael J.
Meade, Brendan
Gerstoft, Peter

To be published in: Seismological Research Letters

DOI to publisher's version: 10.1785/0220180259

Permalink to record: http://permalink.lanl.gov/object/view?what=info:lanl-repo/lareport/LA-UR-18-28089

Machine Learning in Seismology —Turning Data into Insights

Qingkai Kong, Daniel T. Trugman, Zachary E. Ross, Michael J. Bianco, Brendan Meade, Peter Gerstoft

Abstract:

This paper provides an overview of current applications of machine learning (ML) in seismology. ML techniques are becoming increasingly widespread in seismology, with applications ranging from identifying unseen signals and patterns to extracting features that might improve our physical understanding. The survey of the applications in seismology presented here serves as a catalyst for further use of ML. Five research areas in seismology are surveyed where ML classification, regression, clustering algorithms show promise: earthquake detection and phase picking, earthquake early warning, ground motion prediction, seismic tomography, and earthquake geodesy. We conclude by discussing the need for a hybrid approach combining data-driven ML with traditional physical modeling.

Introduction

In a broad sense, machine learning (ML) is a set of related techniques that extract information directly from data using well-defined optimization rules. ML has recently drawn attention due to its wide-ranging success in various fields (Murphy 2012; Jordan and Mitchell, 2015; Witten et al., 2016). Seismology has been a data-intensive field since its very origin. As the years have progressed and the field has expanded, numerous methods and tools have been developed to detect and characterize earthquakes and to study earth structure. We as a community have already developed a rich set of techniques, but ML can bring a different and complementary set of useful tools.

In seismology, we are currently undergoing rapid changes in the "3 V's" often discussed by the big data community (Sagiroglu and Sinanc, 2013): volume, variety, and velocity. For example, the archive of seismic waveform publicly available from Incorporated Research Institutions for Seismology (IRIS) is increasing in size exponentially (Figure 1). This dramatically increased volume of data (and the secondary products derived from the raw data) makes manual processing difficult. Many ML algorithms are designed with large datasets in mind: typically, more data gives better results. Dataset variety has increased too. Besides seismic data, other types of relevant geophysical datasets (e.g., GPS time series and InSAR images) are readily available from UNAVCO and other resource centers. The use of joint geophysical datasets might provide better resolution in certain problems, and carefully designed ML techniques can help analyze these datasets without introducing unnecessary complexity (Khaleghi et al., 2013). Finally, velocity refers to the speed of data processing and distribution. This is important for real-time earthquake detection and earthquake early warning, which rely on rapid analyses of high-velocity data streams.

An exciting aspect of applying ML to seismology is the potential to find unseen patterns or new and significant features in our datasets. A recent example is using ML to predict the timing of the next slip event in laboratory slip experiments with features extracted from low-amplitude acoustic emissions that were previously considered to be noise (Rouet-Leduc et al., 2017). As seismologists, we have the intuition and logic to analyze data, but ML could work beyond human intuition to facilitate the discovery of unconsidered patterns.

In essence, all ML algorithms learn from data using probability theory, which has been the mainstay of statistical methods for centuries. Most ML algorithms can be grouped into two main categories: supervised learning and unsupervised learning. Depending on whether the data has target labels or not (Figure 2), one category may be preferable. Supervised learning, which comprises predictive modeling and operates on labeled datasets, can be further subdivided into classification and regression algorithms based on whether the target outputs are categorical (classification) or quantitative (regression). Unsupervised learning is subdivided into clustering and dimensionality reduction, depending on whether we are interested in grouping data into categories based on similarity, or simply reducing the input data dimensions. There are other more exotic types of ML algorithms such as semi-supervised learning and reinforcement learning, for which we refer readers to more advanced texts (e.g., Murphy, 2012; Goodfellow et al., 2016).

ML algorithms, while diverse in their implementation, tend to follow a basic workflow that includes the following steps (Figure 3). In Step 1: Data Collection, data is collected and partitioning into training and testing sets. A key aspect of ML is training the model on a random subset of the dataset, and then verifying the model on independent testing data. In Step 2: Preprocessing, data is cleaned and formatted, and missing data are removed or repaired. Feature extraction, which increases the performance of many ML algorithms by transforming the raw data into a more useful state for a given task, may also be performed. In Step 3: Model Training, numerical optimization algorithms are

used to iteratively tune the model parameters based on a cost function specific to the learning task of the problem. In Step 4: Model Evaluation, model performance is evaluated on test data. Finally, in Step 5: Production, the finished ML model is applied in production mode to new data.

ML has received enormous recent interest across a wide range of disciplines (Jordan and Mitchell 2015; LeCun et al., 2015). We hope that this paper will inspire both seismologists to further explore ML theory and techniques and data scientists to apply their latest ML algorithms in seismological fields. While this manuscript provides a high-level overview of potential ML applications to seismology, there are many wonderful textbooks and online courses that provide greater details about individual ML algorithms and their implementations (See Data and Resources). We next discuss recent applications of ML in seismology and their potential for obtaining new geophysical insights.

Applications of ML in Seismology

In the following, we present a detailed survey of five specific applications of ML to earthquake seismology, while acknowledging that there are many other worthy applications that merit discussion.

Earthquake Detection and Phase Picking

Automated detection and picking of earthquakes are long-standing problems in seismology, with the first algorithms developed in the late 1970's (e.g., Allen, 1978).

Today, these subjects are still active research areas, and incorporate technology from computer science, electrical engineering, statistics, and many other fields to extract as much information from the data as possible. Some of the earliest applications of ML learning to seismology were to the problem of discrimination and classification of seismic events (e.g., Dysart et al., 1990; Fedorenko et al., 1999; Musil and Plesinger, 1996, Ursino et al., 2001), and in recent years this research has expanded to include utilizing ML to improve earthquake detection and phase picking capabilities (e.g., Dai & MacBeth 1995; Tiira 1999; Wiszniowski et al., 2014; Zhao & Takano 1999). These efforts have shown considerable promise to date, most notably in the area of deep learning, and suggest that many exciting new developments are coming in the near future. In particular, there is a distinct possibility that these algorithms will surpass the capabilities of human experts for the first time. Here we outline some of the most promising examples of ML applied to the earthquake detection problem.

Over the last decade, there has been an explosion of interest in using the similarity of waveforms between nearby sources to detect previously unidentified earthquakes. This originally began with matched filtering (template matching), which uses waveforms of known events as templates to scan through continuous waveforms for new event detection (Gibbons and Ringdal 2006; Shelly et al. 2007; Peng and Zhao 2009; Kato et al. 2012; Ross et al. 2017; Chamberlain et al. 2018; Beauce et al. 2018). Recently, there has been an interest in applying ML and data mining algorithms for similarity-based event detection. In Perol et al. (2018), a convolutional neural network (CNN) was trained to simultaneously detect and locate earthquakes based on single station

waveform classification. For a given window of data, the goal is to predict which of several spatial regions the event occurred in, with the option for rejecting all of them. Alternatively, the Fingerprinting and Similarity Thresholding (FAST) algorithm (Bergen and Beroza, 2018; Yoon et al., 2015) is a data mining approach that converts an entire continuous waveform dataset into a database of binary fingerprints. These fingerprints are compact representations of short segments of continuous waveform data, and are organized in a special dictionary structure for efficient lookup. A key feature of FAST is that it is essentially unsupervised: earthquakes can be identified without prior knowledge of seismicity because, for highly-similar waveforms, fingerprints are more similar to each other than those of random noise sources. In addition, FAST is computationally more efficient than template matching, which will help to facilitate automated processing of large waveform datasets.

A new category of earthquake detection algorithms that has recently emerged is generalized phase detection (GPD; Ross et al., 2018a). Rather than search for near-identical waveforms, GPD instead trains convolutional networks to learn generalized representations of seismic waves from millions of example seismograms. This knowledge is then used to classify windows of data as P, S, or noise (Figure 4). It has been shown to reliably identify P- and S-waves with excellent temporal sensitivity and performance in low SNR conditions, resulting in typically 5-10 times as many events detected as conventional methods. GPD can simultaneously be used to pick arrival times with high precision. A key advantage of the method is that once trained, the model can be applied to datasets other than just those encompassed by the training set, such

as data recorded in different tectonic regimes, large magnitude earthquakes, and active source explosions. This is advantageous in situations where a seismicity catalog is unavailable to use for template matching, or in seismic monitoring.

In addition to detecting earthquakes, there have been a number of noteworthy developments in algorithms for phase picking with ML. Chen (2018) developed an approach to pick seismic wave arrival times using fuzzy clustering, which is based on the idea that the amplitudes of the seismic data before and after the arrival can be treated as separate, but possibly overlapping, clusters. This enables a decision boundary to be drawn that is taken as the arrival pick. Zhu and Beroza (2018) have found great success in applying fully-convolutional networks to pick P- and S-wave arrival times by training on millions of seismograms picked manually in Northern California. Their method takes complete 3-component seismograms as inputs, and outputs probability time series corresponding to the likelihood of P- and S-wave onsets. They demonstrate state-of-the-art picking performance for both phase types, and their method further provides an important empirical mechanism for estimating the quality of the picks. This includes difficult cases such as clipped seismograms, where even human analysts would have a difficult time. Ross et al. (2018b) trained a CNN to pick P-wave onset times, but instead used the network as a regressor to predict the time index of the phase onset. They also trained a separate convolutional network to pick first-motion polarities of P-waves, which are essential ingredients in calculating focal mechanisms. They demonstrated that the networks can often pick polarities more

accurately than professional seismic analysts, as well as more frequently. This will lead to more detailed and expanded focal mechanism catalogs.

There have been several other exciting recent applications of ML to earthquake detection problems. Beyreuther et al. (2012) developed a hidden Markov model (HMM) to classify and detect events for volcanic and geothermal areas. Treating event detections as an object detection problem, Wu et al. (2018) cascaded region-based CNN to capture laboratory slip events of different durations. Finally, Aguiar and Beroza (2014) and Zhang et al. (2014) combined insights from Google's Pagerank and other image-based search engine methods to obtain waveform templates for low frequency earthquakes.

Earthquake Early Warning and Real-time Machine Learning

Earthquake Early Warning (EEW) systems provide seconds to minutes of warning before the strongest shaking by taking advantage of the fact that electronic signals travel much faster than seismic waves, and that the S-wave and surface wave phases that produce the strongest shaking travel slower than the first P-wave arrivals (Allen et al., 2009). There have been several recent efforts in EEW using ML algorithms, either based on hand-selected physical features extracted from seconds of waveforms (Kong et al., 2016a), or using deep learning algorithms to automatically extract features to identify the onset of the earthquakes at a single station (see more examples in the preceding subsection). Li et al. (2018b) trained a Generative Adversarial Network to learn the characteristics of both earthquake P-wave arrivals and background noise, resulting in a discriminator that mitigates false triggering. ML techniques such as

Support Vector Machine regression and neural networks (NN) have also been used to estimate the magnitude, epicentral distance, and other relevant parameters using input features derived from a short time window of waveform data following the P-wave arrival (Böse et al., 2008; Cuéllar et al., 2018; Ochoa et al., 2018). Meier et al. (2015) proposed a method to estimate the magnitude and station-source distance by estimating the posterior probabilities from the observed frequency content to reduce the uncertainties. Böse et al. (2012) developed an image recognition based algorithm to classify the observed ground motion amplitudes into near-source/far-source regions and map a finite fault rupture estimate automatically. Lastly, Cua and Heaton (2007) proposed a unified framework for different components of the EEW, including real-time earthquake source estimation and alert decision-making using a Bayesian approach. Using station locations, previously observed seismicity, and known fault traces as prior information could improve the system performance, especially at regions with low station density (Yin et al., 2018).

Myshake, a recent effort using smartphones to detect nearby earthquakes and provide EEW to the public, has many ML aspects and demonstrates promising results (Kong et al., 2016b). Myshake has two levels of detections: a single-phone supervised approach and a NN-based unsupervised approach. For individual phones, a trained NN is implemented on each phone to distinguish the earthquake signals from everyday human activities. Using a two-second sliding window on a filtered three-component waveform, Myshake extracts three important features from the phone that represent the amplitude and frequency content of the movement. These features are then fed into a

NN algorithm to classify whether the waveform is from an earthquake or from human-activities (Figure 5). When a phone detects an earthquake-like waveform, it sends a trigger message to the cloud server with a timestamp, location, and amplitude to make further confirmation of the earthquake by considering groups of phones within a region at a network level. The triggers from the phones are aggregated to a proper resolution grid to reduce the real-time computation burden, and the DBSCAN (Density-based spatial clustering of applications with noise, Ester et al., 1996) algorithm finds potential clusters of phones that have likely been triggered by an earthquake. Trigger messages in the clusters identified by DBSCAN are used to estimate the earthquake location and magnitude.

Ground Motion Prediction Using Supervised Learning

Ground motion prediction is a crucial aspect of earthquake hazard assessment, and while simple in concept it is challenging to perform in practice. At its core, ground motion prediction answers the question: given a hypothetical earthquake source, how strong is the shaking likely to be? The basic physical factors controlling ground motion are well-established: one can think of the ground motion observed at the surface as a convolution of source, path, and site effects (Boore, 1983). The classical approach to ground motion prediction uses linear regression to model the first-order aspects of these effects (Campbell and Bozorgnia, 2008). In a linear ground motion prediction equation (GMPE), the predicted ground motion Y (in logarithmic units) is a normal-distributed random variable that is a linear function of the input variables, which include the earthquake magnitude M and source-site distance R:

$$\log Y = c_0 + c_1 M + c_2 \log R + \varepsilon; \quad \varepsilon \sim \cdot(0, \sigma^2). \tag{1}$$

Here c_0, c_1, and c_2 are empirical coefficients. The misfit term ε includes both epistemic uncertainty that can be reduced through better observational constraints and more sophisticated modeling approaches, as well as random variability that cannot be reduced (Douglas and Edwards, 2016). Probabilistic seismic hazard assessments are particularly sensitive to epistemic uncertainty in GMPEs (Anderson and Brune, 1999), and its reduction has been a primary focus in developing new GMPEs. For example, several of the most recent linear regression models developed for the Next Generation Attenuation Relationships project (Bozorgnia et al., 2014) include dozens of regression coefficients, which reduces data misfit but at the cost of increased model complexity. There has also been a significant recent effort to develop generic linear GMPEs that are regionally-adjustable and hence exportable to different regions (Yenier and Atkinson, 2015).

Despite these advances, it is challenging to incorporate more complex source, site and path effects within a linear GMPE. A viable alternative is to treat ground motion prediction as a supervised learning problem, with well-defined input and target variables but considerably more flexibility on the model design. Moreover, the central focus on model validation inherent to the ML paradigm, including carefully partitioning of training and testing datasets, would help alleviate the traditional quixotic focus in GMPE model development on reducing the data misfit (Bindi, 2017), and instead allows us shift our attention to improvements in predictive validity.

Some of the earliest ML GMPEs (Alavi and Gandomi, 2011) employed shallow NNs, and this approach is still the most commonly used. Derras et al., (2012) analyzed KiK-net records collected in Japan using a NN with a single hidden layer to predict peak ground acceleration (PGA) as a target variable using five input features: magnitude, epicentral distance, source depth, near-surface shear wave speed, and site resonance frequency. Derras et al. (2014, 2016) generalized this approach to multiple target variables, including PGA but also peak ground velocity and pseudo-spectral accelerations at periods of interest for structural design. NNs are only one of many viable applications of supervised learning techniques to ground motion prediction. Alimoradi and Beck (2015) developed a technique to synthesize realistic strong-motion records by applying Gaussian process regression to a sparse, orthonormal set of basis vectors called eigenquakes, which represent characteristic earthquake records. Thomas et al. (2016) developed a randomized adaptive neuro-fuzzy inference system to analyze records from the Pacific Earthquake Engineering Research database. While these studies differ in technical details, viewed holistically they demonstrate the potential for improved predictive performance over linear GMPEs using similar input and target variables.

The modeling flexibility inherent to supervised learning also allows for the examination of input features that are not traditionally incorporated in linear GMPEs. To this end, Trugman and Shearer (2018) used a generalization of the Random Forest supervised learning algorithm (Breiman, 2001) to quantify the relation between earthquake stress

drop and PGA for earthquakes in the San Francisco Bay Area. While the basic correlation is intuitive - higher stress drop events are enriched in high frequencies and should have systematically larger ground motion amplitudes (Baltay and Hanks, 2014) - there exist few quantitative estimates for the importance of this effect and how it varies with magnitude and source-site distance. Trugman and Shearer (2018) demonstrate that the event residual terms learned by the Random Forest GMPE have a physical basis in the variability in earthquake stress drop (Figure 6), thus highlighting the utility of ML techniques in ground motion modeling.

ML tends to work best in scenarios where high-quality data is plentiful and easily available. This presents a significant challenge in ground motion prediction, where near-source records of large magnitude earthquakes -- which pose the greatest hazard -- are sparse. Future studies may focus on the best ways of integrating limited observational data in this regime with synthetic data from broadband rupture simulations, which will become increasingly prevalent in the coming years (Khoshnevis and Taborda, 2018). Likewise, ML approaches to nonergodic GMPEs, where the predicted ground motions vary spatially (Anderson and Brune, 1999), have yet to be fully explored. While ML itself is not a panacea for the outstanding questions in earthquake hazard analysis, a ML approach to ground motion prediction may prove to be a powerful new tool in the next generation of seismic hazard assessments.

Tomography and Illuminating Geophysical Structure with Machine Learning

ML in seismic tomography has shown great promise for improving our understanding of subsurface geophysical structure. Seismic tomography methods obtain subsurface

models or ‘images’ from sensor array observations of seismic waves, which are generated by anthropogenic sources, earthquakes, or ambient noise processing. Seismic tomography is critical for deducing geophysical structure and characterizing seismic hazard (Rawlinson et al., 2010). However, the demands placed on these methods are great, as tomography models are often estimated from limited and noise-corrupted observations with non-linear forward models. Such ill-posed inverse problems require regularization or assimilation of hypothesized geophysical structure to obtain physically plausible solutions. ML represents a modern paradigm for signal processing, with more sophisticated model priors and latent representations (Murphy, 2012) than classic inverse methods like Tikhonov or total variation regularization (Aster et al., 2011). ML priors include sparsity constraints and latent dictionaries. The non-linear ‘general function approximation’ capability of NNs (Bishop, 2006) permits replacement of seismic data simulation and inversion procedures with NNs. In the following, we introduce seismic tomography and show how it has benefitted from ML theory, including unsupervised and deep learning.

Seismic tomography can be categorized as either travel time tomography or full waveform inversion (FWI; Virieux and Operto, 2009). Travel time tomography calculates slowness (the inverse of seismic wave speed) perturbations to reference models using source-receiver travel time measurements. FWI methods calculate perturbations to a reference model which best predict "full" recorded seismic waveforms. For both methods, the basic optimization problem is

(2)

where is a penalty (e.g. least squares), is the data, denotes the forward model, and are model parameters. is minimized with respect to , for realizations of data. In travel time tomography, are travel time observations, and contains ray information relating the travel times to slownesses . However, solving (2) directly will almost certainly give poor results, because it is ill-posed and ill-conditioned by the non-uniqueness, non-linearity, and sensitivity to noise of the forward operator . ML provides a means of constraining geophysical features in such models, but it is reliant on adequate training data to obtain reasonable performance.

The application of simple ML implementations to the seismic tomography problem is problematic due to a lack of training data, because in regional to global-scale seismic tomography, no geophysical ground truth data exists. This issue has driven development of more advanced ML-based methods in seismic tomography that do not depend on large volumes of training data, or that are trained on simulations. Methods that do not require ground truths are based on adaptive, unsupervised learning frameworks (Elad, 2010; Mairal et al., 2014). In these adaptive approaches, data observations themselves are used for training with unsupervised learning.

Adaptive ML-based seismic tomography methods, inspired by image denoising (Elad, 2010) and medical imaging (Ravishankar and Bresler, 2011; Greenspan et al., 2016), have achieved compelling results. These methods combine sparse modeling with

unsupervised learning. In sparse modeling, signals are represented using few (sparse) atoms from a dictionary of atoms . Such atoms are solved using a least squares objective function with a sparsity inducing prior. For example, a sparsity constraint is added to (2), with with , as

. (3)

Here, are the sparse coefficients, the -norm enforces sparsity by counting the number of non-zero coefficients, and is a tuning parameter controlling sparsity that is analogous to a regularization constant in classical inverse methods. The -norm is non-convex and is typically solved using greedy methods. Under certain conditions, the -norm can be replaced by the the -norm, which is convex (Elad, 2010). The atoms in the dictionary D represent 'elemental' geophysical features, and can be represented by functions such as wavelets.

Sparse representations are appealing because they can model both discontinuous and smooth geophysical features (Loris et al., 2007). In adaptive tomography, the dictionaries are learned directly from signal examples using dictionary learning, a form of unsupervised ML for which many algorithms exist (Mairal et al., 2014). Such learned dictionaries can better represent specific signals than wavelets. Zhu et al. (2017) used this sparse and adaptive framework for FWI by iteratively learning the dictionary of seismic features, while Li and Harris (2018) incorporated a non-local similarity (Mairal et al., 2014) in the dictionary learning procedure. Bianco and Gerstoft (2018) used a

sparse and adaptive framework for 2D (surface wave) travel time tomography, called locally sparse travel time tomography (LST). Assuming dense, straight-ray sampling, LST learns a dictionary of slowness features from patches of a least squares inversion. The learned dictionary is then used to construct a sparse slowness model.

Seismic tomography approaches based on NNs, with early theory developed by Röth and Tarantola (1994), have also achieved compelling results. Moya et al. (2010) apply a NN approach to velocity model inversion. Gupta et al. (2018) address the challenges of limited measurements in travel time tomography using subspace modeling and convolutional NNs. Moseley et al. (2018) present a fast approximate approach for seismic wave propagation and inversion using deep learning, based on deep NN for speech synthesis. Araya-Polo et al., (2018) develop a formulation for FWI that replaces the iterative inversion scheme for velocity features with a deep NN. Lewis and Vig (2017) use a deep NN FWI method to better detect salt domes. Such methods appear to be a future step for generative and inversion architectures.

Earthquake Geodesy and Non-inertial Deformation

While classical seismology has focused on high-frequency inertial deformation of the earth, the full spectrum of earthquake cycle behaviors also includes prolonged non-inertial deformation (Ben-Zion, 2008). These motions include postseismic deformation (durations of years) and interseismic deformation (durations of decades), as well as slow/silent earthquakes (durations of weeks) (Peng and Gomberg, 2010; Ide and Beroza, 2011). Because these motions are non-inertial, they are typically measured using geodetic techniques like GPS and InSAR, estimate time-dependent

displacements at Earth's surface. The precision of these measurements (GPS: ~0.1 mm/yr, InSAR: ~5 mm/yr) limits the resolving power of geodetic data to relatively large earthquakes that occur near Earth's surface. Earthquakes with M<4 are difficult to observe geodetically because they are characterized by relatively small fault slip (<1 cm) and static displacements in the elastic crust fall off as the reciprocal of distance cubed for buried earthquakes. This suggests that there will be orders of magnitude fewer earthquakes observed geodetically than there are seismically. And this bears on ML applications because it implies far fewer earthquakes available for creating labeled datasets from geodetic data. The case is similar for postseismic deformation as well as silent/slow slip events. Currently, the total number of such geodetically observed events may be on the order of 1,000.

Given this relative paucity of classically labeled data, ML applications to the non-inertial part of the earthquake cycle may be somewhat different than those initially applied to seismic waveforms. In particular, seemingly obvious goals like automating the search for slow/silent earthquakes may be challenging due to limited training data. Instead, other opportunities arise. For example, there are numerous non-exclusive mechanisms involved in postseismic deformation including both linear and non-linear versions of afterslip, poroelasticity, and viscoelasticity. Here ML approaches developed to infer the governing partial differential equations directly from observations (Long et al., 2018; Raissi and Karniadakis 2018; Rudy et al., 2017) may play an essential role in resolving the nature and relative contributions of the mechanisms responsible for postseismic deformation. The core idea is that these ML approaches realize the mathematical

structure of the governing physics (both linear and non-linear) directly from observations of surface motions rather than relying on theory-driven concepts that have received traditional focus.

ML approaches also offer the possibility of radically accelerating generative models of earthquake cycle deformation. Numerical rupture and viscoelastic stress transfer models are widely used in earthquake science, but they are not ubiquitous. The primary reason for this is the computational cost of running these simulations and models. In some cases, it may be possible to train deep learning systems to emulate high-performance computing earthquake physics codes so that they are represented in compact mathematical forms as NNs. The central concept here is that we tend to program calculations in terms of mathematical functions that are readily recognizable and comprehensible. However familiar these may be, there may exist far more compact non-linear and non-orthogonal factorizations that enable the solution to be computed quickly, and NNs are free to construct over complete dictionary representation that may be vastly more computationally efficient (DeVries et al., 2017; Moseley et al., 2018).

Other applications and future directions

There are many other exciting ML applications in our field of seismology. For example, the use of probabilistic graphical models and graph theory in seismology has become increasingly prevalent. The deployment of large-N arrays (Karplus and Schmandt, 2018) provides one such opportunity, where weak event signals can be extracted using graph clustering (Riahi and Gerstoft, 2017) or similarity theory (Li et al., 2018a). Separately,

Trugman and Shearer (2017) use graph theory and hierarchical cluster analysis to obtain high-precision earthquake hypocentral estimates using differential travel times from pairs of earthquakes observed at a set of common stations. Telesca and Chelidze (2018) applied a visibility graph method to seismicity near a dam to find anomalous seismic activity.

Additional applications of ML to seismology extend well beyond the realm of graph theory. Araya-Polo et al. (2017) applied a deep NN trained on active seismic data for hydrocarbon exploration to detect subsurface fault structures. Krischer and Fichtner (2017) generate synthetic seismograms using Generative Adversarial Networks (GANs), training the networks using with synthetic seismic data. Using Bayesian networks, Hincks et al. (2018) modeled the joint conditional dependencies between parameters for the Oklahoma seismicity to understand the induced seismicity. Building on the preliminary analyses of Meade et al. (2017), DeVries et al. (2018) trained a deep NN to forecast aftershock locations using as input the static stress change tensor computed from finite fault earthquake rupture models.

The ultimate realization of ML-based methods in seismology would leverage physical models to obtain synergy between the physical theory from domain scientists and the enhanced, data-driven constraints from ML and probability theory. While the application of ML to seismology is becoming increasingly prevalent, ML is often currently applied without physical modeling (Fig. 8, upper left). Geophysical data sets tend to be poorly sampled, noisy, and incomplete, and are often difficult to handle using standard ML

techniques. Thus, often in seismology, we have traditionally deferred to pure physics-based methods (Fig. 8, bottom). It would be transformative if we could develop a hybrid modeling framework that combines data-driven ML methods with explicit physical models (Fig. 8, upper right). This would provide a means of specifying a physical model as a component of the ML algorithm, or conversely, a means of using ML to train better physical parameterizations. Transparency of the learned algorithms would enable human learning and allow validation by testing for physical consistency.

In summary, seismology and machine learning benefit from each other. With its interesting problems and rich datasets, seismology supplies a real-world testbed for various ML algorithms, and even a driving force to compel the development of new algorithms. While ML provides seismology with new tools to extract novel insights directly from the data, combining classical seismology techniques with ML in a hybrid approach might lead to radically new discoveries.

Acknowledgements

We thank Editor-in-Chief Zhigang Peng for the invitation to write this frontier paper and three anonymous reviewers for their thoughtful comments which greatly improved the manuscript. Q. Kong acknowledges support from the Gordon and Betty Moore Foundation through grant GBMF5230 to UC Berkeley. D. Trugman acknowledges institutional support from the Laboratory Directed Research and Development (LDRD) program at Los Alamos National Laboratory. Z. Ross acknowledges support from The Gordon and Betty Moore Foundation and the National Science Foundation. M. Bianco

and P. Gerstoft acknowledge support from the Office of Naval Research (grant no. N00014-18-1-2118).

Data and Resources

For further reading on ML fundamentals, we recommend the following textbooks and online course materials. Bishop (2006) is a more introductory text, whereas Murphy (2012) provides a more in-depth theoretical development. “Deep Learning” (Goodfellow et al., 2016) provides a practical introduction to deep neural networks. There are also many excellent free online courses, such as Ng’s “Machine Learning”, Hinton’s “Neural Networks for Machine Learning”, Tibshirani and Hastie’s “Statistical Learning” , and Li et al.’s “Convolutional Neural Networks for Visual Recognition”. This is not meant to be an exhaustive list of ML resources, but is a good place to get started.

References

Aguiar, A. C., and Beroza, G. C. (2014). PageRank for Earthquakes. Seismological Research Letters, 85(2), 344–350.

Alavi, A. H., and Gandomi, A. H. (2011). Prediction of principal ground-motion parameters using a hybrid method coupling artificial neural networks and simulated annealing. Computers & Structures, 89(23), 2176–2194.

Alimoradi A., and Beck L. J. (2015). Machine-Learning Methods for Earthquake Ground Motion Analysis and Simulation. Journal of Engineering Mechanics, 141(4), 04014147.

Allen, Rex V. (1978). Automatic earthquake recognition and timing from single traces. Bulletin of the Seismological Society of America 68.5 (1978): 1521-1532

Allen, R. M., Gasparini, P., Kamigaichi, O., and Bose, M. (2009). The Status of Earthquake Early Warning around the World: An Introductory Overview. Seismological Research Letters, 80(5), 682–693.

Anderson, J. G., and Brune, J. N. (1999). Probabilistic Seismic Hazard Analysis without the Ergodic Assumption. Seismological Research Letters, 70(1), 19–28.

Araya-Polo, M., Dahlke, T., Frogner, C., Zhang, C., Poggio, T., and Hohl, D. (2017). Automated fault detection without seismic processing. Leading Edge, 36(3), 208–214.

Araya-Polo, M., Jennings, J., Adler, A., and Dahlke, T. (2018). Deep-learning tomography. Leading Edge, 37(1), 58–66.

Aster, R. C., Borchers, B., and Thurber, C. H. (2011). Parameter Estimation and Inverse Problems. Academic Press.

Baltay, A. S., and Hanks, T. C. (2014). Understanding the magnitude dependence of PGA and PGV in NGA West 2 data. Bulletin of the Seismological Society of America, 104 (6): 2851–2865.

Beaucé, E., Frank, W. B., and Romanenko, A. (2018). Fast Matched Filter (FMF): An Efficient Seismic Matched Filter Search for Both CPU and GPU Architectures. Seismological Research Letters, 89(1), 165–172.

Ben-Zion, Y. (2008). Collective behavior of earthquakes and faults: Continuum-discrete transitions, progressive evolutionary changes, and different dynamic regimes. Reviews of Geophysics , 46(4).

Bergen, K. J., and Beroza, G. C. (2018). Detecting earthquakes over a seismic network using single-station similarity measures. Geophysical Journal International, 213(3), 1984–1998.

Beroza, G. C., and Ide, S. (2011). Slow Earthquakes and Nonvolcanic Tremor. Annual Review of Earth and Planetary Sciences, 39(1), 271–296.

Beyreuther, M., Hammer, C., Wassermann, J., Ohrnberger, M., and Megies, T. (2012). Constructing a Hidden Markov Model based earthquake detector: application to induced seismicity. Geophysical Journal International, 189(1), 602–610.

Bianco, M. J., and Gerstoft, P. (2018). Travel time tomography with adaptive dictionaries. IEEE Transactions on Computational Imaging, doi:10.1109/TCI.2018.2862644.

Bindi, D. (2017). The predictive power of ground-motion prediction equations. Bulletin of the Seismological Society of America, 107(2), 1005–1011.

Bishop, C. M. (2006). Pattern Recognition and Machine Learning. Springer Verlag.

Boore, D. M. (1983). Stochastic simulation of high-frequency ground motions based on seismological models of the radiated spectra. Bulletin of the Seismological Society of America, 73(6A), 1865–1894.

Böse, M., Wenzel, F., and Erdik, M. (2008). PreSEIS: A neural network-based approach to earthquake early warning for finite faults. Bulletin of the Seismological Society of America, , 98 (1). pp. 366-382.

Böse, M., Heaton, T. H., and Hauksson, E. (2012). Real-time Finite Fault Rupture Detector (FinDer) for large earthquakes. Geophysical Journal International, 191(2), 803–812.

Bozorgnia, Y., Abrahamson, N. A., Atik, L. A., Ancheta, T. D., Atkinson, G. M., Baker, J. W., et al. (2014). NGA-West2 research project. Earthquake Spectra, 30(3), 973–987.

Breiman, L. (2001). Random Forests. Machine Learning, 45(1), 5–32.

Campbell, K. W., and Bozorgnia, Y. (2008). NGA ground motion model for the geometric mean horizontal component of PGA, PGV, PGD and 5% damped linear elastic response spectra for periods ranging from 0.01 to 10 s. Earthquake Spectra, 24(1), 139–171.

Chamberlain, C. J., Hopp, C. J., Boese, C. M., Warren Smith, E., Chambers, D., Chu, S. X., et al. (2018). EQcorrscan: Repeating and Near Repeating Earthquake Detection and Analysis in Python. Seismological Research Letters, 89(1), 173–181.

Chen, Y. (2018). Automatic microseismic event picking via unsupervised machine learning. Geophysical Journal International, 212(1), 88–102.

Cua, G., and Heaton, T. (2007). The Virtual Seismologist (VS) Method: a Bayesian Approach to Earthquake Early Warning. In P. Gasparini, G. Manfredi, and J. Zschau (Eds.), Earthquake Early Warning Systems (pp. 97–132). Berlin, Heidelberg: Springer Berlin Heidelberg.

Cuéllar, A., Suárez, G., and Espinosa-Aranda, J. M. (2018). A Fast Earthquake Early Warning Algorithm Based on the First 3 s of the P-Wave Coda. Bulletin of the Seismological Society of America, 108(4), 2068–2079.

Dai, H., and MacBeth, C. (1995). Automatic picking of seismic arrivals in local earthquake data using an artificial neural network. Geophysical Journal International, 120(3), 758–774.

Derras, B., Bard, P.-Y., Cotton, F., and Bekkouche, A. (2012). Adapting the neural network approach to PGA prediction: An example based on the KiK-net data. Bulletin of the Seismological Society of America, 102(4), 1446–1461.

Derras, B., Bard, P. Y., and Cotton, F. (2014). Towards fully data driven ground-motion prediction models for Europe. Bulletin of Earthquake Engineering, 12(1), 495–516.

Derras, B., Bard, P.-Y., and Cotton, F. (2016). Site-condition proxies, ground motion variability, and data-driven GMPEs: Insights from the NGA-West2 and RESORCE data sets. Earthquake Spectra, 32(4), 2027–2056.

DeVries, P. M. R., Thompson, T. B., and Meade, B. J. (2017). Enabling large-scale viscoelastic calculations via neural network acceleration. Geophysical Research Letters, 44(6), 2662–2669.

DeVries, P. M. R., Viégas, F., Wattenberg, M., and Meade, B. J. (2018). Deep learning of aftershock patterns following large earthquakes. Nature, 560(7720), 632–634.

Douglas, J., and Edwards, B. (2016). Recent and future developments in earthquake ground motion estimation. Earth-Science Reviews, 160, 203–219.

Dysart, P. S., and Pulli, J. J. (1990). Regional seismic event classification at the NORESS array: Seismological measurements and the use of trained neural networks. Bulletin of the Seismological Society of America, 80(6B), 1910–1933.

Elad, M. (2010). From Exact to Approximate Solutions. In M. Elad (Ed.), Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing (pp. 79–109). New York, NY: Springer New York.

Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In Kdd (Vol. 96, pp.

226–231).

Fedorenko, Y. V., Husebye, E. S., and Ruud, B. O. (1999). Explosion site recognition; neural net discriminator using single three-component stations. Physics of the Earth and Planetary Interiors, 113(1), 131–142.

Gibbons, S. J., and Ringdal, F. (2006). The detection of low magnitude seismic events using array-based waveform correlation. Geophysical Journal International, 165(1), 149–166.

Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. (2016). Deep learning (Vol. 1). MIT press Cambridge.

Greenspan, H., van Ginneken, B., and Summers, R. M. (2016). Guest Editorial Deep Learning in Medical Imaging: Overview and Future Promise of an Exciting New Technique. IEEE Transactions on Medical Imaging, 35(5), 1153–1159.

Gupta, S., Kothari, K., de Hoop, M. V., and Dokmanić, I. (2018). Deep Mesh Projectors for Inverse Problems. arXiv preprint arXiv:1805.11718v1

Hincks, T., Aspinall, W., Cooke, R., and Gernon, T. (2018). Oklahoma's induced seismicity strongly linked to wastewater injection depth. Science, 359(6381), 1251–1255.

Jordan, M. I., and Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. Science, 349(6245), 255–260.

Karplus, M., and Schmandt, B. (n.d.). Preface to the Focus Section on Geophone Array Seismology. Seismological Research Letters. https://doi.org/10.1785/0220180212

Kato, A., Obara, K., Igarashi, T., Tsuruoka, H., Nakagawa, S., and Hirata, N. (2012). Propagation of slow slip leading up to the 2011 M(w) 9.0 Tohoku-Oki earthquake.

Science, 335(6069), 705–708.

Khaleghi, B., Khamis, A., Karray, F. O., and Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. An International Journal on Information Fusion, 14(1), 28–44.

Khoshnevis, N., and Taborda, R. (2018). Prioritizing Ground Motion Validation Metrics Using Semisupervised and Supervised Learning. Bulletin of the Seismological Society of America, 108 (4): 2248–2264. doi: https://doi.org/10.1785/0120180056

Kong, Q., Allen, R. M., Schreier, L., and Kwon, Y.-W. (2016). MyShake: A smartphone seismic network for earthquake early warning and beyond. Science Advances, 2(2), e1501055.

Kong, Q., Allen, R. M., and Schreier, L. (2016). MyShake: Initial observations from a global smartphone seismic network. Geophysical Research Letters, 43(18), 9588–9594.

Krischer, L. and Fichtner, A. (2017). Generating Seismograms with Deep Neural Networks, AGU Fall Meeting Abstracts.

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444.

Lewis, W., and Vigh, D. (2017). Deep learning prior models from seismic images for full-waveform inversion. In SEG Technical Program Expanded Abstracts 2017 (pp. 1512–1517). Society of Exploration Geophysicists.

Li, D., and Harris, J. M. (2018). Full waveform inversion with nonlocal similarity and model-derivative domain adaptive sparsity-promoting regularization. Geophysical Journal International, https://doi.org/10.1093/gji/ggy380.

Li, Z., Peng, Z., Hollis, D., Zhu, L., and McClellan, J. (2018). High-resolution seismic event detection using local similarity for Large-N arrays. Scientific Reports, 8(1), 1646.

Li, Z., Meier, M.-A., Hauksson, E., Zhan, Z., and Andrews, J. (2018). Machine Learning Seismic Wave Discrimination: Application to Earthquake Early Warning. Geophysical Research Letters, 45(10), 4773–4779.

Loris, I., Nolet, G., Daubechies, I., & Dahlen, F. A. (2007), Tomographic inversion using ℓ1-norm regularization of wavelet coefficients, Geophysical Journal International, 170(1), 359-370.

Long, Z., Lu, Y., Ma, X., and Dong, B. (2018). PDE-Net: Learning PDEs from Data. Proceedings of the 35th International Conference on Machine Learning, 3208-3216.

Mairal, J., Bach, F., and Ponce, J. (2014). Sparse Modeling for Image and Vision Processing. Foundations and Trends® in Computer Graphics and Vision, 8(2-3), 85–283.

Meade, B. J., DeVries, P. M. R., Faller, J., Viegas, F., and Wattenberg, M. (2017). What Is Better Than Coulomb Failure Stress? A Ranking of Scalar Static Stress Triggering Mechanisms from 105 Mainshock-Aftershock Pairs. Geophysical Research Letters, 44(22).

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

Moseley, B., Markham, A., and Nissen-Meyer, T. (2018). Fast approximate simulation of seismic waves with deep learning. arXiv preprint arXiv:1807.06873v1.

Moya, A., and Irikura, K. (2010). Inversion of a velocity model using artificial neural networks. Computers & Geosciences, 36(12), 1474–1483.

Murphy, K. P. (2012). Machine Learning: A Probabilistic Perspective. MIT Press.

Musil, M., and Plešinger, A. (1996). Discrimination between local microearthquakes and quarry blasts by multi-layer perceptrons and Kohonen maps. Bulletin of the Seismological Society of America, 86(4), 1077–1090.

Ochoa, L. H., Niño, L. F., and Vargas, C. A. (2018). Fast magnitude determination using a single seismological station record implementing machine learning techniques. Geodesy and Geodynamics, 9(1), 34–41.

Peng, Z., and Zhao, P. (2009). Migration of early aftershocks following the 2004 Parkfield earthquake. Nature Geoscience, 2, 877.

Peng, Z., and Gomberg, J. (2010). An integrated perspective of the continuum between earthquakes and slow-slip phenomena. Nature Geoscience, 3, 599.

Perol, T., Gharbi, M., and Denolle, M. (2018). Convolutional neural network for earthquake detection and location. Science Advances, 4(2), e1700578.

Raissi, M., and Karniadakis, G. E. (2018). Hidden physics models: Machine learning of nonlinear partial differential equations. Journal of Computational Physics, 357, 125–141.

Ravishankar, S., and Bresler, Y. (2011). MR image reconstruction from highly undersampled k-space data by dictionary learning. IEEE Transactions on Medical Imaging, 30(5), 1028–1041.

Rawlinson, N., Pozgay, S., and Fishwick, S. (2010). Seismic tomography: A window into deep Earth. Physics of the Earth and Planetary Interiors, 178(3), 101–135.

Riahi, N., and Gerstoft, P. (2017). Using graph clustering to locate sources within a dense sensor array. Signal Processing, 132, 110–120.

Ross, Z. E., Rollins, C., Cochran, E. S., Hauksson, E., Avouac, J.-P., and Ben-Zion, Y. (2017). Aftershocks driven by afterslip and fluid pressure sweeping through a fault-fracture mesh. Geophysical Research Letters, 44(16), 8260–8267.

Ross, Z. E., Meier, M.-A., Hauksson, E., and Heaton, T. H. (2018a). Generalized Seismic Phase Detection with Deep Learning. Bulletin of the Seismological Society of America, doi: https://doi.org/10.1785/0120180080

Ross, Z. E., Meier, M.-A., and Hauksson, E. (2018b). P-wave arrival picking and first-motion polarity determination with deep learning. Journal of Geophysical Research: Solid Earth, 123, 5120–5129.

Röth, G., and Tarantola, A. (1994). Neural networks and inversion of seismic data. Journal of Geophysical Research, 99(B4), 6753.

Rouet-Leduc, B., Hulbert, C., Lubbers, N., Barros, K., Humphreys, C. J., and Johnson, P. A. (2017). Machine Learning Predicts Laboratory Earthquakes. Geophysical Research Letters, 44(18), 9276–9282.

Rudy, S. H., Brunton, S. L., Proctor, J. L., and Kutz, J. N. (2017). Data-driven discovery of partial differential equations. Science Advances, 3(4), e1602614.

Sagiroglu, S., and Sinanc, D. (2013). Big data: A review. In 2013 International Conference on Collaboration Technologies and Systems (CTS) (pp. 42–47).

Shelly, D. R., Beroza, G. C., and Ide, S. (2007). Non-volcanic tremor and low-frequency earthquake swarms. Nature, 446(7133), 305–307.

Telesca, L., and Chelidze, T. (2018). Visibility Graph Analysis of Seismicity around

Enguri High Arch Dam, Caucasus. Bulletin of the Seismological Society of America, doi: https://doi.org/10.1785/0120170370

Thomas, S., Pillai, G. N., Pal, K., and Jagtap, P. (2016). Prediction of ground motion parameters using randomized ANFIS (RANFIS). Applied Soft Computing, 40, 624–634.

Tiira, T. (1999). Detecting teleseismic events using artificial neural networks. Computers & Geosciences, 25(8), 929–938.

Trugman, D. T., and Shearer, P. M. (2017). GrowClust: A Hierarchical Clustering Algorithm for Relative Earthquake Relocation, with Application to the Spanish Springs and Sheldon, Nevada, Earthquake Sequences. Seismological Research Letters, 88(2A), 379–391.

Trugman, D. T., and Shearer, P. M. (2018). Strong Correlation between Stress Drop and Peak Ground Acceleration for Recent M 1–4 Earthquakes in the San Francisco Bay Area. Bulletin of the Seismological Society of America, 108 (2): 929–945.

Ursino, A., Langer, H., Scarfì, L., Di Grazia, G., and Gresta, S. (2001). Discrimination of quarry blasts from tectonic microearthquakes in the Hyblean Plateau (Southeastern Sicily). Annals of Geophysics, 44(4). https://doi.org/10.4401/ag-3569

Virieux, J., and Operto, S. (2009). An overview of full-waveform inversion in exploration geophysics. Geophysics, 74(6), WCC1–WCC26.

Wang, J., and Teng, T.-L. (1995). Artificial neural network-based seismic detector. Bulletin of the Seismological Society of America, 85(1), 308–319.

Wiszniowski, J., Plesiewicz, B. M., and Trojanowski, J. (2014). Application of real time recurrent neural network for detection of small natural earthquakes in Poland. Acta

Geophysica, 62(3), 469–485.

Witten, I. H., Frank, E., Hall, M. A., and Pal, C. J. (2016). Data Mining: Practical Machine Learning Tools and Techniques. Morgan Kaufmann.

Wu, Y., Lin, Y., Zhou, Z., Bolton, D. C., Liu, J., and Johnson, P. (2018). DeepDetect: A Cascaded Region-Based Densely Connected Network for Seismic Event Detection. IEEE Transactions on Geoscience and Remote Sensing: A Publication of the IEEE Geoscience and Remote Sensing Society, 1–14.

Yenier, E., and Atkinson, G. M. (2015). Regionally adjustable generic ground motion prediction equation based on equivalent point source simulations: Application to central and eastern North America. Bulletin of the Seismological Society of America, 105 (4): 1989–2009.

Yin, L., Andrews, J., and Heaton, T. (2018). Rapid Earthquake Discrimination for Earthquake Early Warning: A Bayesian Probabilistic Approach Using Three Component Single Station Waveforms and Seismicity Forecast. Bulletin of the Seismological Society of America, 108(4), 2054–2067.

Yoon, C. E., O'Reilly, O., Bergen, K. J., and Beroza, G. C. (2015). Earthquake detection through computationally efficient similarity search. Science Advances, 1(11), e1501057.

Zhang, J., Zhang, H., Chen, E., Zheng, Y., Kuang, W., and Zhang, X. (2014). Real-time earthquake monitoring using a search engine method. Nature Communications, 5, 5664.

Zhu, L., Liu, E., and McClellan, J. (2017). Sparse-promoting full-waveform inversion based on online orthonormal dictionary learning. Geophysics, 82(2), R87–R107.

Zhu, W., and Beroza, G. C. (2018). PhaseNet: A Deep-Neural-Network-Based Seismic Arrival Time Picking Method. arXiv preprint arXiv:1803.03211v1.

Figure 1. The IRIS DMC (Data Management Center) archive growth (modified from IRIS). The growth of the seismic waveform data at the IRIS DMC from the time it was established until the 1st Sep 2018. (http://ds.iris.edu/files/stats/data/archive/Archive_Growth.jpg)

Figure 2. Types of machine learning (ML) algorithms. Supervised ML operates on labeled datasets with the objective to develop models that predict either categorical or quantitative target variables. Unsupervised ML operates on unlabeled datasets with the objective to group data by similarity or reduce the dimensionality of the input datasets. Some common ML algorithms are listed at the bottom for each category.

Figure 3. A generic machine learning workflow that guides many applications: (1) Data Collection, (2) Preprocessing, (3) Model Training, (4) Model Evaluation, and (5) Production.

Figure 4. Example of Generalized Phase Detection. a) Cartoon schematic of a CNN for GPD. A convolutional feature extraction system is combined with a fully connected neural network to produce class probabilities for P-waves, S-waves, and noise (Ross et al. 2018a). b) Application of GPD to the 2016 Borrego Springs, CA sequence. Red and

blue colors indicate P- and S-waves, respectively. Vertical bars indicate automated picks.

Figure 5. The neural network (NN) used in MyShake earthquake early warning phone application. (a) The workflow of the NN algorithm on the phone, including extraction of features from recorded phone motion and implementation of a NN classifier to distinguish between motions from humans and earthquakes. (b) The interquartile range and maximum zero-crossing rate are two important features for distinguishing between earthquake and non-earthquake motions. (c) Example application of Myshake at the network level to a M4.4 earthquake that occured in January of 2018. NN triggers from individual users are compared against theoretical P and S arrivals.

Figure 6. Random Forest Ground Motion Prediction Equation (GMPE; Trugman and Shearer 2018) and earthquake stress drop versus peak ground acceleration (PGA). (a) Schematic workflow for training Random Forest GMPE. (b) Peak ground acceleration vs. hypocentral distance for seismicity in the San Francisco Bay Area. Each point represents a site-corrected PGA measurement from an earthquake at a single station. Also shown is the median value in equally-spaced magnitude-distance bins (large markers) and predicted values from the Random Forest GMPE (dashed lines). (c) Event PGA residuals learned from Random Forest GMPE versus earthquake stress drop. The least-squares linear fit and correlation coefficient are marked for reference.

Figure 7. Locally sparse travel time tomography (LST) of checkerboard slowness. (a)

Synthetic checkerboard slowness patterns with 100x100 pixel grid (km) are sampled by (b) 2016 straight rays from 64 seismic stations. (c) Conventional inversion using damping and smoothing regularization, and (d) LST. Profiles from the 2D inversion are shown with true and estimated slownesses. The root-mean-squared error (ms/km) estimated relative to the true slowness is printed on the 2D estimates. (e) Dictionary learned from LST contains checkerboard-like atom (100 atoms shown). Each atom (patch) is 10x10 pixels.

Figure 8. Geophysical insight will be maximized by leveraging the strengths of both physical and machine learning (ML)-based, data driven models. Analytic physical models (lower left) give basic insights about physical systems. More sophisticated models, reliant on computational methods (lower right), can model more complex phenomena. Whereas physical models are reliant on rules, which are updated by physical evidence (data), ML is purely data-driven (upper left). By augmenting ML methods with physical models to obtain hybrid models (upper right), a synergy can be obtained that balance the complementary strengths of physical intuition with data-driven insights.

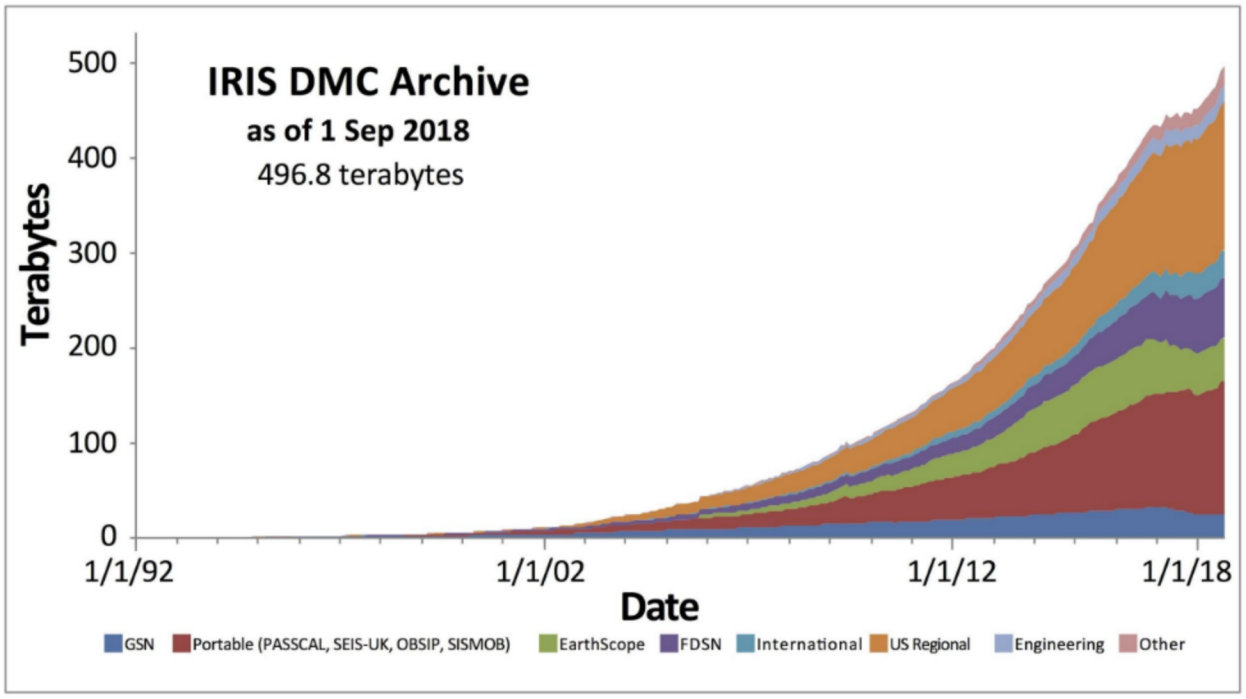

Figure 1. The IRIS DMC (Data Management Center) archive growth (modified from IRIS). The growth of the seismic waveform data at the IRIS DMC from the time it was established until the 1st Sep 2018. (http://ds.iris.edu/files/stats/data/archive/Archive_Growth.jpg)

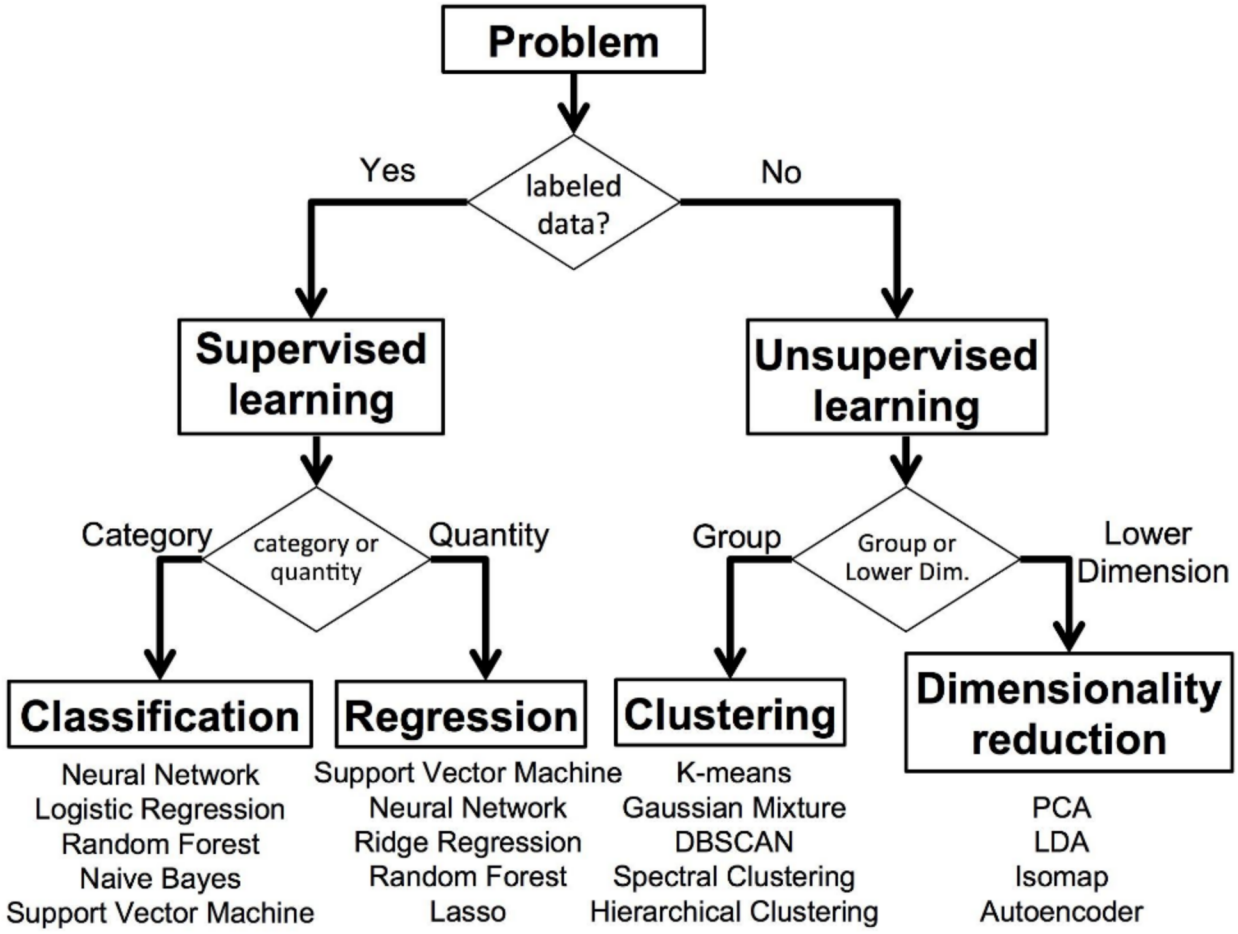

Figure 2. Types of machine learning (ML) algorithms. Supervised ML operates on labeled datasets with the objective to develop models that predict either categorical or quantitative target variables. Unsupervised ML operates on unlabeled datasets with the objective to group data by similarity or reduce the dimensionality of the input datasets. Some common ML algorithms are listed at the bottom for each category.

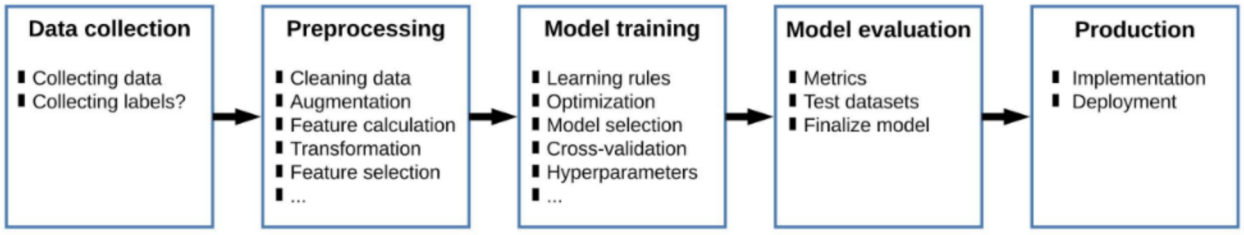

Figure 3. A generic machine learning workflow that guides many applications: (1) Data Collection, (2) Preprocessing, (3) Model Training, (4) Model Evaluation, and (5) Production.

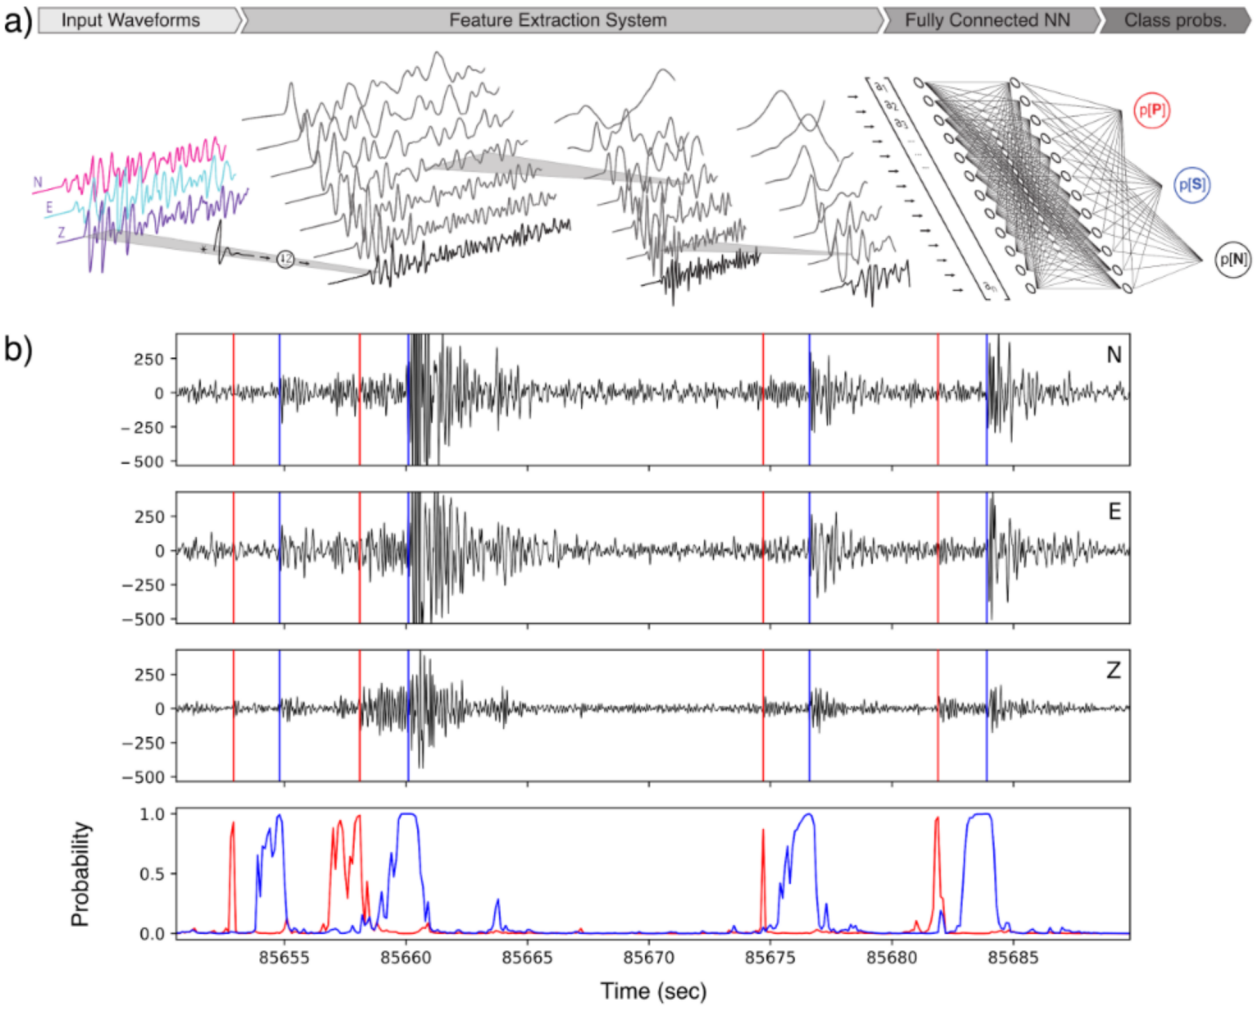

Figure 4. Example of Generalized Phase Detection. a) Cartoon schematic of a CNN for GPD. A convolutional feature extraction system is combined with a fully connected neural network to produce class probabilities for P-waves, S-waves, and noise (Ross et al. 2018a). b) Application of GPD to the 2016 Borrego Springs, CA sequence. Red and

blue colors indicate P- and S-waves, respectively. Vertical bars indicate automated picks.

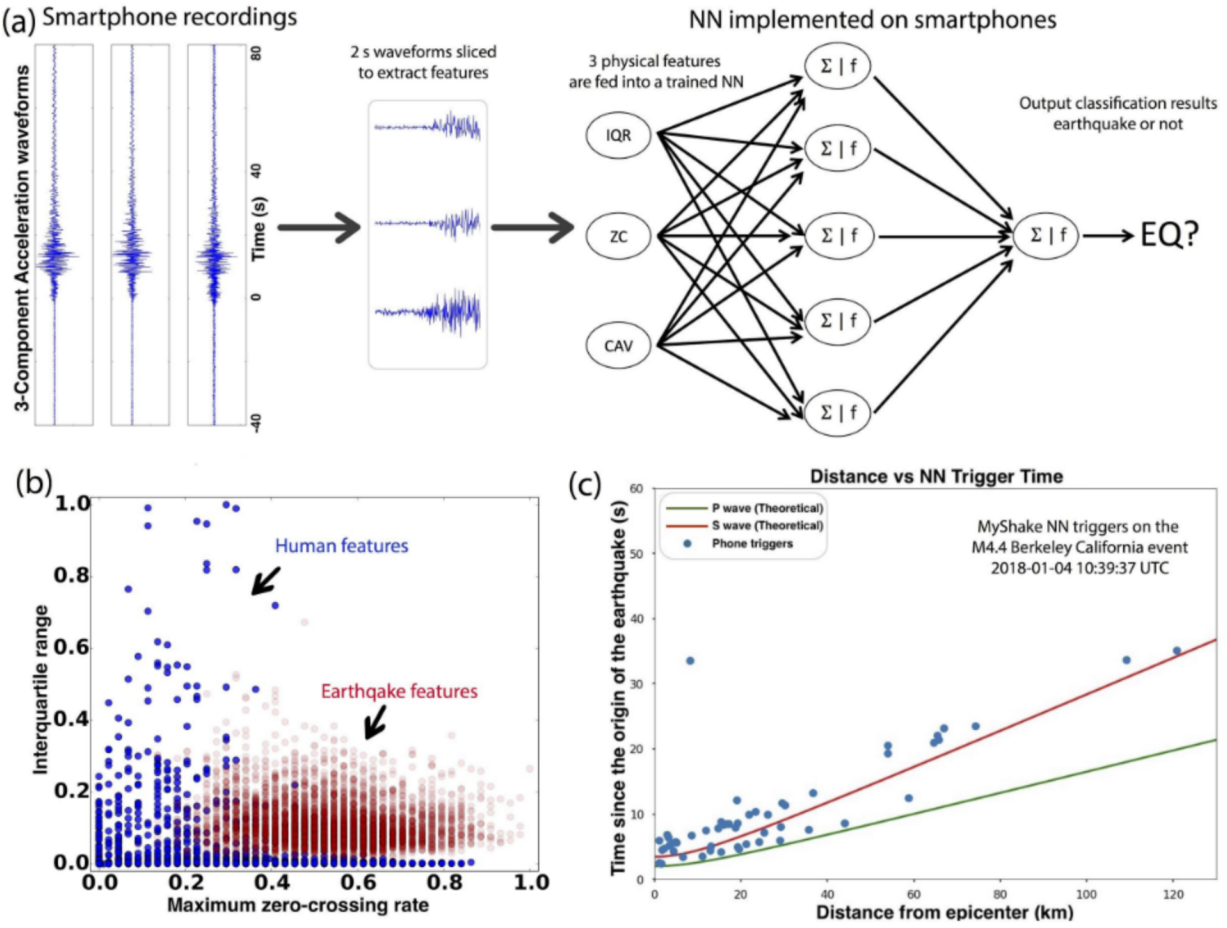

Figure 5. The neural network (NN) used in MyShake earthquake early warning phone application. (a) The workflow of the NN algorithm on the phone, including extraction of features from recorded phone motion and implementation of a NN classifier to distinguish between motions from humans and earthquakes. (b) The interquartile range and maximum zero-crossing rate are two important features for distinguishing between earthquake and non-earthquake motions. (c) Example application of Myshake at the

network level to a M4.4 earthquake that occured in January of 2018. NN triggers from individual users are compared against theoretical P and S arrivals.

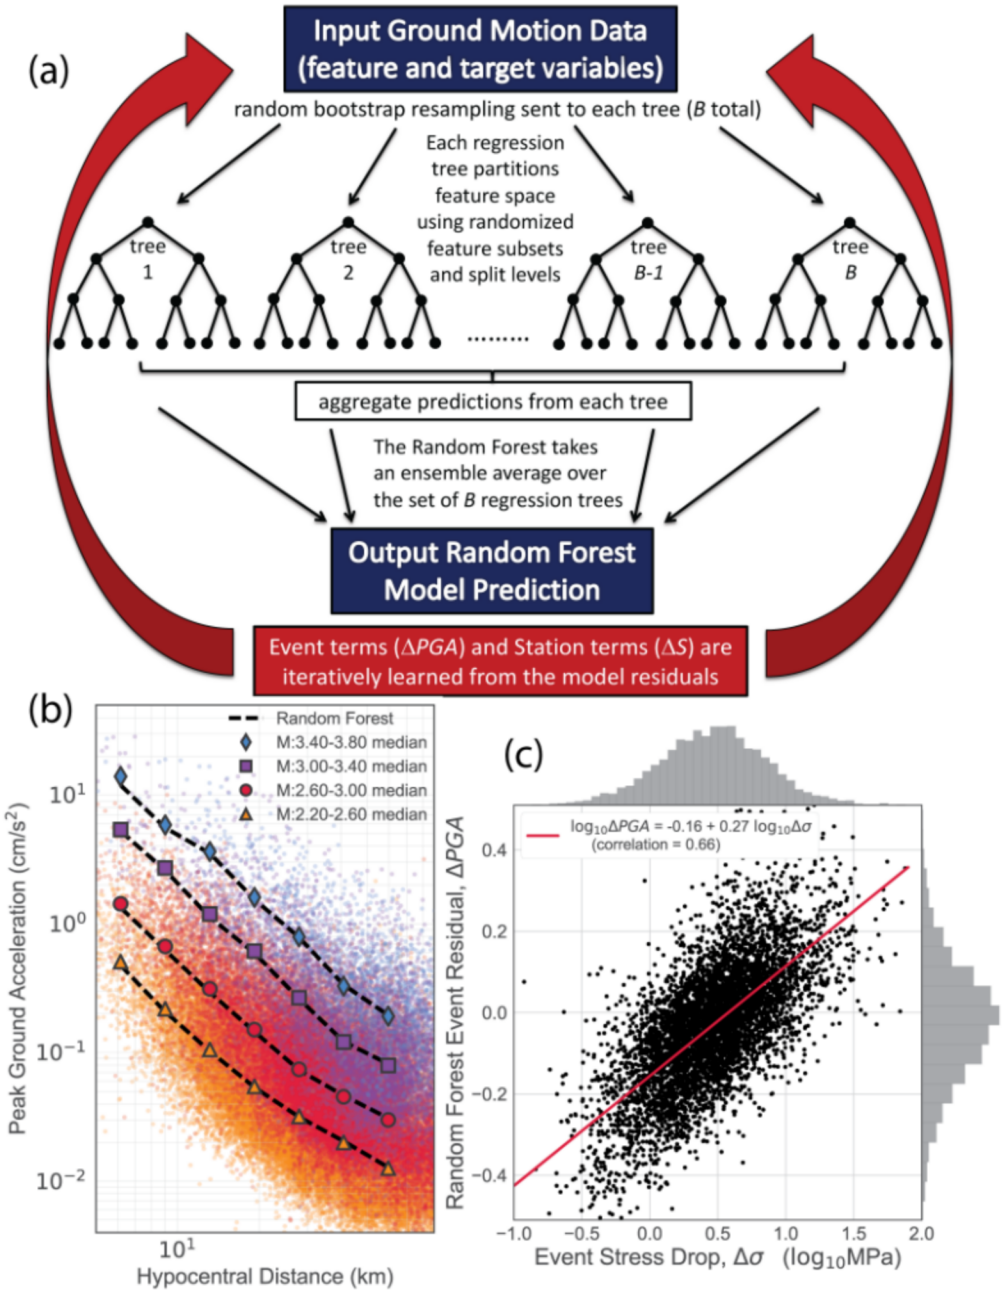

Figure 6. Random Forest Ground Motion Prediction Equation (GMPE; Trugman and Shearer 2018) and earthquake stress drop versus peak ground acceleration (PGA). (a)

Schematic workflow for training Random Forest GMPE. (b) Peak ground acceleration vs. hypocentral distance for seismicity in the San Francisco Bay Area. Each point represents a site-corrected PGA measurement from an earthquake at a single station. Also shown is the median value in equally-spaced magnitude-distance bins (large markers) and predicted values from the Random Forest GMPE (dashed lines). (c) Event PGA residuals learned from Random Forest GMPE versus earthquake stress drop. The least-squares linear fit and correlation coefficient are marked for reference.

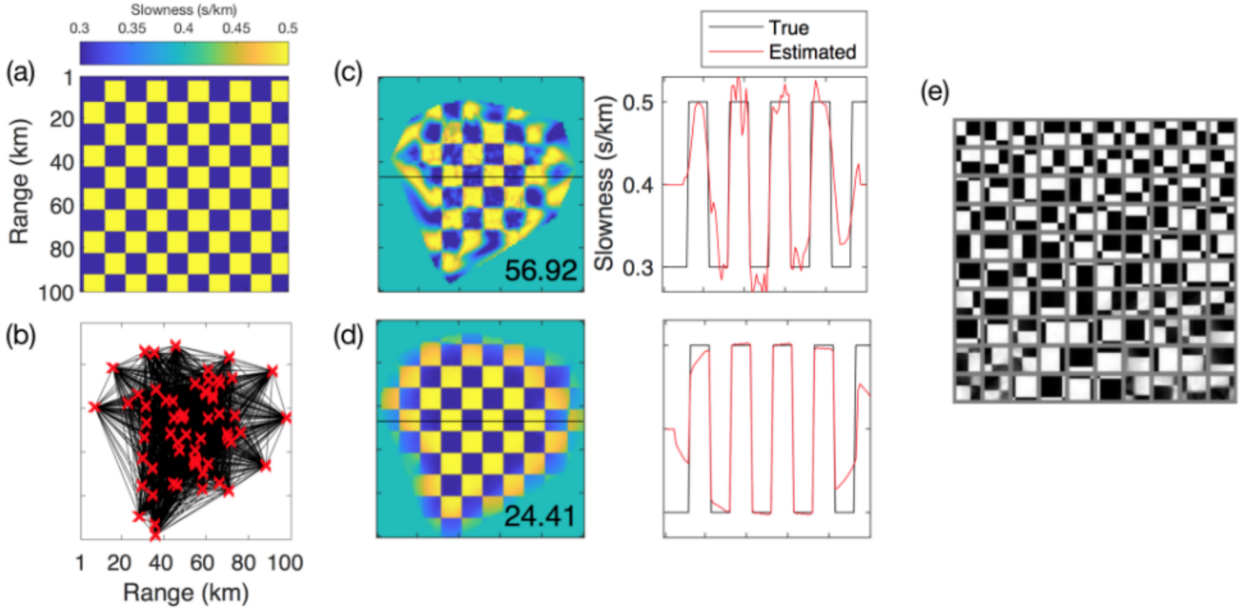

Figure 7. Locally sparse travel time tomography (LST) of checkerboard slowness. (a) Synthetic checkerboard slowness patterns with 100x100 pixel grid (km) are sampled by (b) 2016 straight rays from 64 seismic stations. (c) Conventional inversion using damping and smoothing regularization, and (d) LST. Profiles from the 2D inversion are shown with true and estimated slownesses. The root-mean-squared error (ms/km) estimated relative to the true slowness is printed on the 2D estimates. (e) Dictionary

learned from LST contains checkerboard-like atom (100 atoms shown). Each atom (patch) is 10x10 pixels.

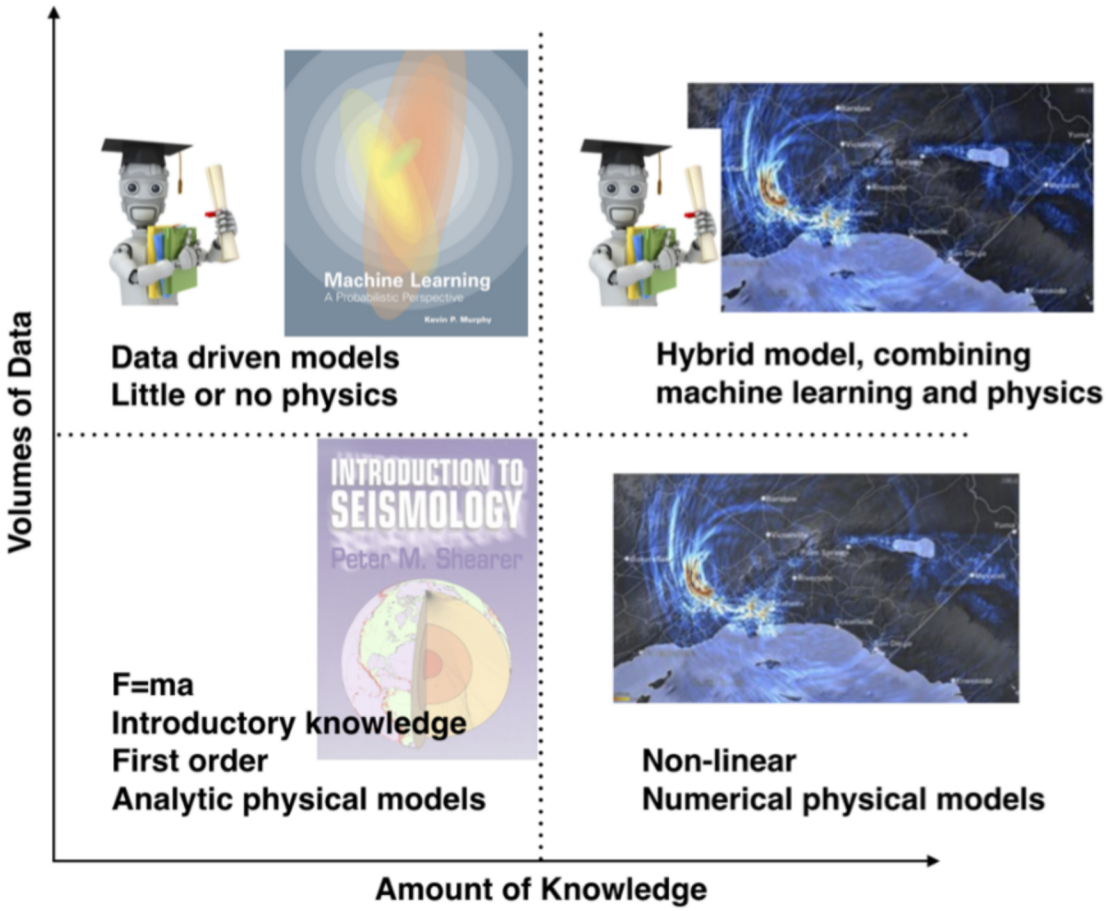

Figure 8. Geophysical insight will be maximized by leveraging the strengths of both physical and machine learning (ML)-based, data driven models. Analytic physical models (lower left) give basic insights about physical systems. More sophisticated models, reliant on computational methods (lower right), can model more complex phenomena. Whereas physical models are reliant on rules, which are updated by physical evidence (data), ML is purely data-driven (upper left). By augmenting ML methods with physical models to obtain hybrid models (upper right), a synergy can be

obtained that balance the complementary strengths of physical intuition with data-driven insights.