International Journal of Research and Review in Applied Science, Humanities, and Technology
Vol 1, Issue 2 December 2024 ISSN: 3048-975X
https://ijrasht.com/

Machine Learning for Earthquake Prediction: Techniques for Seismic Forecasting and Disaster Preparedness

Priyam, Department of Computer Science and Engineering (AI & ML), Panipat Institute of Engineering and Technology (PIET), Samalkha, Panipat, Haryana, India 2822415.cse@pietgroup.co.in

Payal Mor, Department of Computer Science and Engineering (AI & ML), Panipat Institute of Engineering and Technology (PIET), Samalkha, Panipat, Haryana, India 2822892.cse@pietgroup.co.in

Gaurav Tyagi, Department of Computer Science and Engineering (AI & ML), Panipat Institute of Engineering and Technology (PIET), Samalkha, Panipat, Haryana, India 28231324.cse@pietgroup.co.in

Kirti, Department of Computer Science and Engineering (AI & ML), Panipat Institute of Engineering and Technology (PIET), Samalkha, Panipat, Haryana, India 2822408.cse@pietgroup.co.in

Rakhi, Department of Computer Science and Engineering (AI & ML), Panipat Institute of Engineering and Technology (PIET), Samalkha, Panipat, Haryana, India rakhi.cse-aiml@piet.co.in

Abstract— Earthquakes remain among the most difficult natural hazards to forecast, and decades of research have shown that deterministic prediction of the precise time, location, and magnitude of an individual earthquake is not currently achievable with any known method. Machine learning has nonetheless become an increasingly important tool across the broader task of seismic forecasting, supporting probabilistic hazard estimation, rapid earthquake early warning, aftershock sequence forecasting, and automated detection of subtle seismic signals that traditional methods struggle to identify. This paper surveys the application of machine learning to earthquake-related forecasting and disaster preparedness, distinguishing carefully between deterministic prediction, which remains largely unattainable, and probabilistic forecasting and rapid detection tasks, where machine learning has demonstrated measurable practical value. We review deep learning methods for seismic waveform analysis and phase picking, statistical and learned approaches to aftershock forecasting, and applications to earthquake early warning, induced seismicity monitoring, and infrastructure risk assessment. We present a case study illustrating a machine-learning-based earthquake early warning pipeline, discuss benchmarks used to evaluate seismic machine learning systems, and examine the persistent challenges of data imbalance, generalization across tectonic regions, and the risk of conflating improved detection capability with genuine predictive capability. We conclude with future directions connecting seismic machine learning with physics-informed modeling and multi-hazard disaster response systems.

Index Terms—earthquake prediction, seismic forecasting, earthquake early warning, aftershock forecasting, deep learning for seismology, induced seismicity, disaster preparedness

I. INTRODUCTION

Earthquakes cause tens of thousands of deaths and substantial economic damage globally in a typical decade, and the ability to anticipate seismic hazard, whether through long-term probabilistic hazard assessment, short-term forecasting, or rapid detection of an earthquake already in progress, has direct and significant implications for public safety and infrastructure planning. Despite this importance, and despite substantial research investment over more than a century, deterministic earthquake prediction, meaning the ability to specify in advance the precise time, location, and magnitude of a specific future earthquake with actionable accuracy, remains an unsolved and, according to the consensus view of the seismological community, likely unattainable goal given current understanding of the underlying physical processes governing fault rupture initiation.

This distinction between prediction and forecasting is important and is sometimes lost in popular discussion of AI applications to earthquakes. Seismic forecasting, in the sense used throughout this paper, refers to probabilistic statements about the likelihood of seismic activity within a specified spatial region and time window, such as the probability of a damaging aftershock within a given area over the coming week, rather than a deterministic claim about a specific future event. Machine learning has demonstrated genuine and measurable value for this probabilistic forecasting task, as well as for closely related tasks including rapid automated detection and characterization of seismic events, earthquake early warning, and identification of subtle precursory or induced seismic signals that traditional signal processing methods struggle to detect reliably.

This paper surveys the application of machine learning across this broader landscape of seismic forecasting and disaster preparedness tasks. We are careful throughout to distinguish tasks where machine learning has demonstrated clear practical value, such as automated seismic phase picking and aftershock forecasting, from the substantially more difficult and still largely unattained goal of deterministic earthquake prediction, and we discuss why this distinction matters both scientifically and for responsible public communication about the current capabilities of AI in this domain.

The remainder of this paper is organized as follows. Section II reviews background on the earthquake prediction problem and the historical role of machine learning in seismology. Section III presents a taxonomy of machine learning approaches to seismic forecasting tasks. Section IV describes core methods in technical detail. Section V surveys applications. Section VI summarizes benchmarks and datasets. Section VII offers a comparative discussion. Section VIII presents a case study of an earthquake early warning pipeline. Section IX discusses implementation considerations. Section X considers ethical and communication-related implications. Section XI discusses open challenges, Section XII outlines future directions, and Section XIII concludes.

II. BACKGROUND AND MOTIVATION

A. The Challenge of Deterministic Earthquake Prediction

Earthquakes are generated by the sudden release of accumulated stress along a fault, a process governed by complex, poorly observable subsurface physical conditions, including the precise stress state and frictional properties along a fault at depth, that cannot currently be measured directly with sufficient precision or completeness to support reliable deterministic prediction. Decades of searching for reliable earthquake precursors, observable signals that reliably precede a major earthquake with sufficient lead time and specificity to be operationally useful, have not identified any precursor signal shown to reliably generalize across different faults and tectonic settings, leading most seismologists to conclude that short-term deterministic prediction of individual earthquakes is not achievable with current or foreseeable observational and theoretical understanding.

This scientific reality has important implications for how machine learning applications in this space should be framed and evaluated. A machine learning system that reports high accuracy on a narrowly constructed retrospective prediction task can create a misleading impression of genuine predictive capability if the evaluation does not carefully account for the risk of overfitting to a limited historical dataset, the difficulty of true prospective, out-of-sample evaluation for rare events, and the fundamental scientific uncertainty regarding whether reliable earthquake precursors exist at all for a given fault system.

B. Why Machine Learning Is Nonetheless Valuable in Seismology

Even setting aside deterministic prediction, seismology involves several tasks well suited to modern machine learning methods. Seismic networks generate enormous continuous streams of waveform data from which relatively rare, brief seismic events must be automatically detected and characterized, a pattern recognition task well suited to deep learning architectures that have proven effective at analogous rare-event detection tasks in other domains. Aftershock sequences following a major earthquake follow statistically well-characterized, though imperfectly understood, patterns that machine learning models can learn to forecast probabilistically with demonstrably improved accuracy relative to simpler statistical baselines. Additionally, rapid, low-latency processing of seismic waveform data, a requirement for earthquake early warning systems that must detect and characterize an earthquake within seconds of its onset to provide any meaningful advance warning before damaging shaking arrives at a populated area, benefits directly from the fast inference speed of trained machine learning models relative to some traditional processing pipelines.

Machine learning has also proven valuable for detecting and characterizing seismic signals too subtle or numerous for practical manual analyst review, such as the very large number of small-magnitude earthquakes that occur continuously in most seismically active regions, or slow-slip events and induced seismicity signals associated with specific human activities such as wastewater injection, both of which are discussed further in Section IV, and both of which benefit from automated large-scale signal detection capability that would be impractical to achieve through manual seismogram review alone given the data volumes involved.

C. Historical Context

Automated seismic event detection has a long history predating modern deep learning, relying on comparatively simple signal processing techniques such as short-term-average/long-term-average triggering algorithms and template matching against known waveform patterns, techniques that remain in operational use at many seismic monitoring networks today but that can struggle with signals that deviate from previously characterized patterns or that occur during periods of elevated background seismic noise. The application of deep learning to seismology accelerated substantially over the past decade, initially focused on automated seismic phase picking, the task of identifying the arrival time of specific seismic wave types within a waveform, a foundational task underlying most downstream seismic monitoring and forecasting applications, before subsequently expanding to more complex forecasting tasks including aftershock sequence forecasting and earthquake early warning, discussed in detail in the following sections. This historical progression mirrors a pattern observed in several other scientific disciplines discussed elsewhere in this survey series, in which machine learning adoption typically begins with comparatively well-defined, objectively verifiable perceptual or detection tasks, such as identifying a seismic wave arrival time against a manually labelled ground truth, before gradually extending to more scientifically ambiguous and harder-to-validate forecasting tasks, where genuine progress is more difficult to establish conclusively and where the risk of over-claiming capability based on limited or improperly validated retrospective evaluation is correspondingly greater, a concern discussed further in Section VI regarding evaluation methodology specific to seismic forecasting research.

III. A TAXONOMY OF MACHINE LEARNING APPROACHES TO SEISMIC FORECASTING

Machine learning applications in seismology can be organized according to the time horizon and the nature of the forecasting or detection task being addressed. Real-time detection and characterization tasks, including seismic phase picking and earthquake early warning, operate on a timescale of seconds and aim to rapidly detect and characterize an earthquake that has already begun rupturing, rather than forecasting a future event. Short-term probabilistic forecasting tasks, including aftershock sequence forecasting, aim to estimate the probability of future seismic activity over a time horizon of hours to weeks following a triggering event such as a mainshock. Long-term hazard assessment tasks aim to estimate the long-term probability of damaging seismic activity within a region over a time horizon of years to decades, informing building codes and infrastructure planning rather than any specific near-term warning.

A separate but related dimension concerns the underlying data source used for forecasting or detection, spanning waveform-based methods that operate directly on raw or minimally processed seismic waveform data, catalogue-based methods that operate on processed earthquake catalogs listing the time, location, and magnitude of previously detected events, and multi-sensor methods that combine seismic data with other observational data sources such as GPS-based ground deformation measurements or satellite-based interferometric synthetic aperture radar imagery of surface deformation. Table I summarizes this taxonomy with representative methods for each combination of task timescale and data source.

TABLE I. TAXONOMY OF MACHINE LEARNING APPROACHES IN SEISMIC FORECASTING

Task Timescale	Representative Task	Typical Data Source	Example Methods
Real-time (seconds)	Phase picking, early warning	Raw seismic waveforms	PhaseNet, EQTransformer, CNN-based EEW
Short-term (hours-weeks)	Aftershock forecasting	Earthquake catalogs, waveforms	Neural network aftershock models, ETAS-ML hybrids
Long-term (years-decades)	Probabilistic hazard assessment	Historical catalogs, geologic data	Statistical seismicity models, ML-enhanced PSHA
Multi-sensor	Slow-slip/induced seismicity detection	Seismic + GPS/InSAR data	Deep learning slow-slip detectors

IV. CORE METHODS

A. Statistical Seismology and Precursor-Based Approaches

Classical statistical seismology models, such as the Epidemic-Type Aftershock Sequence (ETAS) model, describe seismicity as a self-exciting point process in which each earthquake, including small events, triggers a probabilistically elevated rate of subsequent nearby seismicity, providing a well-established statistical baseline against which newer machine learning forecasting methods are typically compared. Historical research into deterministic earthquake precursors, including anomalous radon gas emission, electromagnetic signals, and unusual animal behavior, has not produced a precursor signal demonstrated to generalize reliably across tectonic settings, and modern machine learning research in this space has increasingly shifted focus away from searching for such precursors toward the probabilistic forecasting and rapid detection tasks discussed throughout this paper, reflecting the scientific consensus described in Section II.

B. Deep Learning for Seismic Waveform Analysis and Phase Picking

Deep learning models, particularly convolutional and, more recently, transformer-based architectures adapted for time-series waveform data, have substantially improved automated seismic phase picking accuracy relative to traditional signal processing methods, particularly for lower-amplitude signals or signals recorded during periods of elevated background noise, where traditional threshold-based triggering algorithms are prone to both missed detections and false triggers. These models are typically trained on large, labelled datasets of historical seismic waveforms with manually annotated phase arrival times, learning to directly map a raw or lightly pre-processed waveform segment to a predicted arrival time and phase type, such as the earlier-arriving, generally lower-amplitude primary wave or the later-arriving, generally higher-amplitude secondary wave.

Beyond phase picking specifically, deep learning models have been applied to broader seismic waveform classification tasks, including discriminating between naturally occurring tectonic earthquakes and other seismic sources such as mining blasts or induced seismic events, and to seismic denoising, in which a learned model separates genuine seismic signal from background noise in a waveform, improving the sensitivity of downstream detection and characterization tasks applied to the denoised waveform.

C. Earthquake Early Warning Systems

Earthquake early warning systems aim to detect an earthquake immediately after rupture initiation, using the earliest-arriving, faster-traveling seismic waves recorded at monitoring stations near the epicentre, and rapidly estimate the earthquake's location and magnitude in order to issue a warning to more distant populated areas before the slower, more damaging shaking waves arrive, providing anywhere from a few seconds to, for sufficiently distant locations, up to a minute or more of advance warning. Machine learning models have been incorporated into several operational and experimental early warning systems to improve the speed and accuracy of this rapid magnitude and location estimation relative to traditional methods, which can be particularly important for reducing warning time and improving accuracy for larger earthquakes whose full rupture extent is not yet apparent from only the initial seconds of recorded waveform data.

A key operational trade-off for earthquake early warning systems, whether based on traditional or machine learning methods, is the balance between warning speed and estimation accuracy: waiting for additional waveform data generally improves the accuracy of the magnitude and location estimate but correspondingly reduces the useful warning time available before damaging shaking arrives, motivating machine learning research specifically focused on improving estimation accuracy from the shortest possible initial waveform window, since even a few additional seconds of warning time can meaningfully improve outcomes for time-critical protective actions such as automated braking of high-speed trains or halting of sensitive industrial processes.

D. Aftershock Sequence Forecasting

Following a major earthquake, the surrounding region typically experiences an elevated rate of aftershock seismicity that decays statistically over time but can itself include damaging events, motivating short-term probabilistic forecasting of aftershock rate and spatial distribution to inform emergency response and public safety guidance in the days and weeks following a mainshock. Deep learning models trained on large historical catalogs of mainshock-aftershock sequences, incorporating spatial stress-change information derived from the mainshock's estimated fault rupture geometry, have demonstrated improved spatial aftershock forecasting accuracy relative to traditional statistical baselines that rely on simpler, more geometrically generic assumptions about aftershock spatial distribution.

These aftershock forecasting models are typically evaluated using retrospective testing against historical mainshock-aftershock sequences withheld from model training, though genuine prospective evaluation, testing a forecasting model's real-time performance on entirely new mainshocks as they occur, provides a stronger test of practical forecasting value and has been increasingly incorporated into operational evaluation frameworks used by government seismic monitoring agencies considering deployment of machine learning aftershock forecasting tools alongside traditional statistical methods.

E. Slow-Slip and Induced Seismicity Detection

Slow-slip events, in which fault slip occurs gradually over days to months rather than in a sudden earthquake rupture, produce subtle, low-amplitude seismic and geodetic signals that are difficult to detect using traditional methods but that may carry information relevant to understanding subsequent seismic hazard along the same fault segment. Machine learning models trained on combined seismic and GPS ground deformation data have demonstrated improved sensitivity for detecting and characterizing these subtle slow-slip signals relative to traditional signal processing approaches, contributing to improved scientific understanding of the relationship between slow and fast fault slip processes.

Induced seismicity, seismic activity caused or triggered by human activities such as wastewater injection associated with oil and gas extraction, or reservoir impoundment behind large dams, presents a related detection and forecasting challenge of direct regulatory and public safety relevance, motivating machine learning approaches specifically trained to detect and, where possible, forecast induced seismic activity associated with specific industrial operations, supporting operational decisions such as adjusting injection rates in response to detected or forecast elevated induced seismic risk.

F. Physics-Informed and Hybrid Approaches

Physics-informed approaches to seismic machine learning incorporate known physical constraints, such as the geometric relationship between fault rupture parameters and expected stress changes in the surrounding region, directly into model architecture or training objective, aiming to combine the flexibility of learned models with the generalization benefits of respecting known physical relationships, an approach directly analogous to the physics-informed methods discussed in broader AI-for-science research and applied here specifically to the seismic forecasting domain. Hybrid statistical-machine-learning approaches similarly combine well-established statistical seismicity models, such as the ETAS model discussed in Section IV-A, with learned components that capture additional structure not represented in the simpler statistical baseline, aiming to retain the statistical model's interpretability and demonstrated baseline reliability while improving forecasting accuracy where sufficient training data is available to support the additional learned component.

V. APPLICATIONS

A. Earthquake Early Warning Deployment

Several national and regional earthquake early warning systems have incorporated or experimented with machine learning components to improve rapid magnitude and location estimation, providing automated public alerts and, in some deployments, automated triggering of protective actions such as halting high-speed trains or opening elevator doors at the nearest floor, reducing risk of injury during the shaking that follows a warning.

Deployment of machine learning components within operational early warning systems has generally proceeded cautiously, reflecting the high reliability requirements of a public safety system where both missed detections and false alarms carry meaningful costs, with machine learning components typically deployed initially as a complement to, rather than a wholesale replacement for, existing traditional early warning algorithms, allowing operational performance to be compared directly before any decision to rely more heavily on the machine learning component.

B. Aftershock Sequence Forecasting in Practice

Following major earthquakes, government seismic agencies increasingly incorporate machine-learning-informed aftershock forecasts into public communication and emergency management guidance, providing probabilistic statements regarding the likelihood of further damaging aftershocks over specified time windows to inform decisions such as building re-occupancy timing and emergency shelter planning in the aftermath of a significant mainshock.

Operational deployment of aftershock forecasting typically involves presenting forecast information alongside, rather than as a replacement for, traditional statistical forecasts such as those derived from the ETAS model, allowing emergency managers and the public to see a degree of convergence or divergence between methods, which itself provides useful qualitative information about the reliability of the current forecast given the inherent scientific uncertainty involved, an approach that reflects a broader, appropriately cautious norm across operational seismic forecasting practice of presenting multiple complementary model outputs rather than a single unqualified point estimate.

C. Induced Seismicity Monitoring

Regulatory agencies overseeing wastewater injection and other industrial activities associated with induced seismicity risk increasingly use machine-learning-based seismic monitoring to detect elevated induced seismic activity in near-real time, supporting traffic-light regulatory frameworks in which injection operations are automatically flagged for review, reduced rate, or suspension if monitored seismic activity exceeds specified thresholds.

Machine learning models applied to induced seismicity monitoring must be particularly attentive to reliably distinguishing induced events from naturally occurring background seismicity in the same region, since misclassification in either direction carries meaningful regulatory and economic consequences, whether through unnecessary operational restrictions imposed in response to seismicity that was not actually induced by the monitored activity, or through failure to detect genuinely induced seismicity that poses an elevated hazard requiring operational response.

D. Building and Infrastructure Risk Assessment

Machine learning models trained on historical earthquake damage data, building characteristics, and estimated ground shaking intensity are used to support rapid post-earthquake damage assessment, prioritizing structural inspection resources toward buildings estimated to be at greatest risk of serious damage, and to support longer-term seismic risk assessment informing building code development and infrastructure retrofit prioritization decisions.

E. Integration with Tsunami Warning Systems

For subduction zone earthquakes with tsunami-generating potential, machine learning models supporting rapid earthquake magnitude and rupture characterization feed directly into downstream tsunami warning systems, which must similarly balance warning speed against estimation accuracy, given that coastal tsunami arrival can occur within minutes of a nearby offshore earthquake, leaving very little time margin for the combined earthquake characterization and tsunami modeling pipeline to issue an actionable warning to coastal populations.

VI. BENCHMARKS AND EVALUATION

Evaluating machine learning methods in seismology requires careful attention to the rarity and non-stationarity of large earthquakes, which complicates standard machine learning evaluation practices such as random train-test splitting, since a random split can allow information from a given earthquake sequence to leak between training and test sets in ways that inflate apparent forecasting performance relative to genuine prospective, real-time evaluation. Table II summarizes representative datasets and benchmarks used in seismic machine learning research.

Standardized, large-scale labelled waveform datasets combining seismic recordings with manually verified phase arrival annotations have substantially accelerated progress in automated phase picking research by providing a common benchmark for comparing model architectures. For aftershock forecasting and broader seismicity forecasting research, evaluation increasingly emphasizes genuine prospective testing, in which a candidate forecasting method's performance is assessed on entirely new earthquake sequences occurring after the method was finalized, rather than solely on retrospective testing against historical sequences that may have already informed model design choices, a distinction of particular importance given the scientific stakes and public communication risks associated with over-claiming forecasting capability based on retrospective evaluation alone.

TABLE II. REPRESENTATIVE DATASETS AND BENCHMARKS IN SEISMIC MACHINE LEARNING

Dataset/Benchmark	Task	Data Type
STEAD	Phase picking, event detection	Labeled seismic waveforms
Southern California Seismic Network catalogs	Aftershock/seismicity forecasting	Earthquake catalog data
Japan Meteorological Agency catalogs	Early warning, forecasting evaluation	Waveforms and catalogs
Los Alamos laboratory earthquake dataset	Precursor signal research (lab-scale)	Acoustic emission waveform data
Global CMT catalog	Long-term hazard assessment	Moment tensor earthquake catalog

VII. COMPARATIVE ANALYSIS AND DISCUSSION

Machine learning methods for real-time detection tasks, such as phase picking and early warning, have demonstrated the clearest and most broadly reproduced practical improvements over traditional methods, benefiting from comparatively large labelled training datasets and a well-defined, objectively verifiable evaluation criterion, namely comparison of predicted phase arrival times against manually verified ground truth. Aftershock and broader seismicity forecasting methods show more mixed but generally positive evidence of improvement over statistical baselines such as the ETAS model, though evaluation in this space is complicated by the comparative rarity of large earthquake sequences available for genuine prospective testing, and by an ongoing and legitimate scientific debate regarding how much of any observed retrospective forecasting improvement will persist under genuine prospective evaluation on entirely new earthquake sequences.

Slow-slip and induced seismicity detection methods occupy a comparatively data-constrained position, given the rarity and subtlety of the underlying signals, and current evidence, while promising, remains more preliminary than for the more mature phase-picking and early-warning applications discussed above. Across all of these applications, a consistent methodological theme is the importance of combining machine learning forecasting output with well-calibrated uncertainty quantification and, where possible, complementary physical or statistical modeling, rather than treating machine learning forecasting output as a standalone, fully deterministic prediction, given the genuine and substantial scientific uncertainty that remains regarding the underlying earthquake generation process.

TABLE III. COMPARATIVE MATURITY OF MACHINE LEARNING APPLICATIONS IN SEISMOLOGY

Application	**Evidence of Improvement Over Baseline**	**Evaluation Maturity**	**Key Limitation**
Phase picking / detection	Strong, widely reproduced	High (objective ground truth)	Generalization to novel tectonic settings
Aftershock forecasting	Moderate to positive	Improving (prospective testing growing)	Data scarcity for large sequences
Earthquake early warning	Positive for speed/accuracy trade-off	Moderate to high (operational testing)	Very short available decision window
Slow-slip/induced seismicity detection	Preliminary but promising	Lower (fewer verified events)	Signal rarity and subtlety

VIII. ILLUSTRATIVE CASE STUDY

Consider a representative machine-learning-based earthquake early warning pipeline deployed across a regional network of seismic monitoring stations. When an earthquake begins rupturing, the nearest monitoring stations detect the first-arriving primary waves within seconds, and a deep learning phase-picking model, trained on large historical labelled waveform datasets as described in Section IV-B, rapidly and precisely identifies the arrival time of these initial waves at each station, substantially faster and, for lower-amplitude signals, more reliably than traditional threshold-based triggering algorithms.

As data from an increasing number of nearby stations becomes available over the following seconds, a second machine learning model estimates the earthquake's location and magnitude from the pattern of arrival times and waveform amplitudes recorded so far, updating this estimate continuously as additional station data arrives, reflecting the fundamental early-warning trade-off between estimation accuracy and available warning time discussed in Section IV-C. Once the estimated magnitude exceeds a specified alerting threshold for a given region, an automated warning is issued to that region, timed to arrive, ideally, before the slower, more damaging secondary and surface waves reach populated areas at greater distance from the epicentre.

This case study illustrates the tightly coupled, multi-stage nature of a practical earthquake early warning pipeline, in which machine learning components for phase picking and rapid magnitude estimation are integrated within a broader, carefully engineered operational system responsible for network communication, alert dissemination, and integration with downstream automated protective actions, and in which the overall system's practical value depends as much on end-to-end engineering reliability and integration as on the accuracy of any individual machine learning component considered in isolation.

IX. IMPLEMENTATION CONSIDERATIONS AND TOOLING

A. Real-Time Processing Infrastructure

Deploying machine learning models for earthquake early warning and real-time seismic monitoring requires low-latency data processing infrastructure capable of ingesting continuous waveform data streams from potentially hundreds or thousands of monitoring stations and running model inference within a strict latency budget, given that even a fraction of a second of additional processing delay can meaningfully reduce the warning time available before damaging shaking arrives at populated areas closer to the anticipated shaking front.

B. Data Quality and Network Heterogeneity

Seismic monitoring networks often combine instruments of varying age, sensitivity, and installation quality, introducing data heterogeneity that machine learning models must be robust to, particularly when a model trained primarily on data from one network or region is deployed or evaluated on a different network with different instrument characteristics and background noise conditions, a domain-generalization concern directly relevant to the broader discussion of out-of-distribution robustness raised elsewhere in this survey series and of particular practical importance for regions with sparser or lower-quality seismic monitoring infrastructure than the well-instrumented networks from which many published training datasets are drawn.

X. ETHICAL AND SOCIETAL CONSIDERATIONS

Public communication regarding the capabilities and limitations of machine learning in earthquake forecasting carries particular ethical weight given the potential public safety consequences of both overstating and understating current capability. Overstating deterministic prediction capability risks encouraging inappropriate public complacency or panic based on a specific claimed prediction that current science cannot actually support, while understating the genuine, more modest but real value of probabilistic forecasting and early warning capability risks discouraging beneficial adoption of tools that have demonstrated measurable practical value for tasks such as rapid detection and aftershock forecasting.

A further consideration concerns equitable access to earthquake early warning and forecasting capability, since the seismic monitoring infrastructure and computational resources required to deploy these systems are unevenly distributed globally, with many of the regions at highest seismic risk, including several rapidly urbanizing areas in seismically active parts of the developing world, currently lacking the dense seismic monitoring networks required to support effective machine-learning-based early warning, motivating international investment and technology transfer efforts aimed at extending these capabilities more equitably to populations at elevated seismic risk regardless of a region's existing economic or infrastructure development level.

XI. CHALLENGES AND OPEN PROBLEMS

A central and recurring challenge across machine learning applications in seismology is data scarcity for the largest, most damaging, and most consequential earthquakes, since these events are, by their nature, rare, meaning that even a large historical seismic catalog typically contains relatively few examples of the largest-magnitude events most relevant to worst-case hazard planning, constraining the ability of data-driven methods to learn reliable patterns specifically for these high-consequence, low-frequency events relative to the much larger number of small, routine earthquakes available for training.

A second challenge concerns generalization across tectonic settings: a model trained on seismic data from one region, with its own characteristic fault geometry, crustal structure, and background seismicity rate, may not generalize reliably to a different region with distinct tectonic characteristics, motivating either region-specific model training, which requires sufficient regional training data that may not be available for less well-instrumented regions, or research into more broadly generalizable model architectures and training strategies capable of transferring effectively across tectonic settings.

A third challenge, discussed throughout this paper, concerns the risk of conflating retrospective forecasting accuracy with genuine prospective predictive capability, given the comparative rarity of opportunities for rigorous prospective evaluation on the specific rare, high-consequence events of greatest ultimate interest.

A fourth challenge concerns integrating machine learning forecasting output into effective, trusted decision-making processes for emergency managers and the public, who must interpret probabilistic forecasting information correctly and act on it appropriately despite the substantial remaining uncertainty in any seismic forecast, a communication challenge that is at least as important to the ultimate practical value of these systems as the underlying algorithmic accuracy itself, and one that requires close collaboration between machine learning researchers, seismologists, and emergency management and risk communication specialists rather than being addressed by algorithmic research alone.

A fifth, more subtle challenge concerns benchmark and dataset representativeness. Many widely used labeled seismic datasets are drawn disproportionately from a small number of well-instrumented seismic networks, typically located in wealthier regions with long-established monitoring infrastructure, meaning that models achieving strong benchmark performance on these datasets may not generalize as effectively to the seismic characteristics, instrument types, and noise conditions of less well-represented regions, some of which face substantial seismic hazard but have historically contributed comparatively little data to the datasets most commonly used for training and evaluating seismic machine learning models, a representativeness concern with direct implications for the equitable global deployment of these techniques discussed further in Section X.

XII. FUTURE RESEARCH DIRECTIONS

One promising direction involves increased integration of multiple observational data sources, including seismic waveform data, GPS-based ground deformation measurements, and satellite-based radar interferometry, into unified machine learning forecasting frameworks capable of jointly exploiting complementary information from each data source, potentially improving detection sensitivity for subtle signals such as slow-slip events that may not be reliably detectable from seismic data alone.

A second direction explores physics-informed and hybrid statistical-machine-learning architectures more deeply, aiming to combine the pattern-recognition strengths of deep learning with the interpretability, physical grounding, and demonstrated baseline reliability of established statistical seismicity models such as ETAS, an approach that may help address legitimate scientific and public communication concerns regarding the interpretability and physical plausibility of purely data-driven forecasting models applied to a domain with such significant public safety stakes.

A further direction concerns expanding rigorous prospective evaluation infrastructure for seismic forecasting research more broadly, potentially through formalized, community-wide prospective forecasting experiments in which multiple research groups' forecasting methods are evaluated in real time against genuinely new earthquake sequences as they occur, providing a more scientifically rigorous and publicly trustworthy basis for assessing genuine forecasting progress than retrospective benchmark comparison alone. Finally, continued investment in extending dense, high-quality seismic monitoring infrastructure to currently under-instrumented but seismically active regions is likely to remain an important complement to algorithmic research progress, since even the most sophisticated machine learning method cannot meaningfully improve forecasting or early warning capability in a region lacking the underlying sensor infrastructure needed to generate the real-time data such methods require.

XIII. CONCLUSION

Machine learning has become a valuable tool across a broad range of seismic forecasting and disaster preparedness tasks, even as deterministic earthquake prediction remains, and is likely to remain, beyond the reach of current scientific understanding. This survey reviewed the taxonomy of machine learning approaches spanning real-time seismic detection and phase picking, earthquake early warning, aftershock sequence forecasting, and slow-slip and induced seismicity detection, examined applications across early warning deployment, aftershock communication, induced seismicity regulation, infrastructure risk assessment, and tsunami warning integration, and presented a case study illustrating an end-to-end machine-learning-based early warning pipeline. While the field faces substantial and, in some cases, likely persistent challenges related to data scarcity for rare high-consequence events, generalization across tectonic settings, and the risk of overstating predictive capability, the demonstrated practical value of machine learning for rapid detection, aftershock forecasting, and early warning, combined with continued progress in physics-informed hybrid modeling and expanded seismic monitoring infrastructure, suggests that machine learning will continue to play an increasingly important, if appropriately bounded, role in reducing the human and economic toll of earthquakes worldwide.

REFERENCES

[1] B. Rouet-Leduc, C. Hulbert, N. Lubbers, K. Barros, C. J. Humphreys, and P. A. Johnson, "Machine learning predicts laboratory earthquakes," Geophysical Research Letters, vol. 44, 2017.
[2] P. A. Johnson, B. Rouet-Leduc, L. J. Pyrak-Nolte, G. C. Beroza, C. J. Marone, C. Hulbert, A. Howard, P. Singer, D. Gordeev, D. Karaflos, et al., "Laboratory earthquake forecasting: A machine learning competition," Proceedings of the National Academy of Sciences, vol. 118, 2021.

[3] T. Perol, M. Gharbi, and M. Denolle, "Convolutional neural network for earthquake detection and location," Science Advances, vol. 4, 2018.
[4] W. Zhu and G. C. Beroza, "PhaseNet: A deep-neural-network-based seismic arrival-time picking method," Geophysical Journal International, vol. 216, 2019.
[5] S. M. Mousavi, W. L. Ellsworth, W. Zhu, L. Y. Chuang, and G. C. Beroza, "Earthquake transformer: An attentive deep-learning model for simultaneous earthquake detection and phase picking," Nature Communications, vol. 11, 2020.
[6] S. M. Mousavi, Y. Sheng, W. Zhu, and G. C. Beroza, "STanford EArthquake Dataset (STEAD): A global data set for seismic machine learning," IEEE Access, vol. 7, 2019.
[7] P. M. R. DeVries, F. Viegas, M. Wattenberg, and B. J. Meade, "Deep learning of aftershock patterns following large earthquakes," Nature, vol. 560, 2018.
[8] K. J. Bergen, P. A. Johnson, M. V. de Hoop, and G. C. Beroza, "Machine learning for data-driven discovery in solid Earth geoscience," Science, vol. 363, 2019.
[9] Y. Y. Kagan and D. D. Jackson, "Long-term earthquake clustering," Geophysical Journal International, vol. 104, 1991.
[10] Y. Ogata, "Statistical models for earthquake occurrences and residual analysis for point processes," Journal of the American Statistical Association, vol. 83, 1988.
[11] R. M. Allen and D. Melgar, "Earthquake early warning: Advances, scientific challenges, and societal needs," Annual Review of Earth and Planetary Sciences, vol. 47, 2019.
[12] R. M. Allen, P. Gasparini, O. Kamigaichi, and M. Bose, "The status of earthquake early warning around the world: An introductory overview," Seismological Research Letters, vol. 80, 2009.
[13] D. Kuyuk and R. Allen, "A global approach to provide magnitude estimates for earthquake early warning alerts," Geophysical Research Letters, vol. 40, 2013.
[14] J. Munchmeyer, D. Bindi, U. Leser, and F. Tilmann, "The transformer earthquake alerting model: A new versatile approach to earthquake early warning," Geophysical Journal International, vol. 225, 2021.
[15] S. M. Mousavi and G. C. Beroza, "Deep-learning seismology," Science, vol. 377, 2022.
[16] M. J. Werner, A. Helmstetter, D. D. Jackson, and Y. Y. Kagan, "High-resolution long-term and short-term earthquake forecasts for California," Bulletin of the Seismological Society of America, vol. 101, 2011.
[17] A. Mignan and M. Broccardo, "Neural network applications in earthquake prediction (1994-2019): Meta-analytic and statistical insights on their limitations," Seismological Research Letters, vol. 91, 2020.
[18] A. Mignan and M. Broccardo, "One neuron versus deep learning in aftershock prediction," Nature, vol. 574, 2019.
[19] Y. Bengio, Y. LeCun, and G. Hinton, "Deep learning for AI," Communications of the ACM, vol. 64, 2021.
[20] A. Ross, S. M. Mousavi, D. T. Trugman, D. Bolton, and G. C. Beroza, "PhaseLink: A deep learning approach to seismic phase association," Journal of Geophysical Research: Solid Earth, vol. 124, 2019.
[21] A. Michelini, S. Cianetti, S. Gaviano, C. Giunchi, D. Jozinovic, and V. Lauciani, "INSTANCE, the Italian seismic dataset for machine learning," Earth System Science Data, vol. 13, 2021.
[22] K. Nakano, K. Yamada, and Y. Sato, "Machine learning for slow slip event detection from GPS time series," Geophysical Research Letters, vol. 47, 2020.
[23] P. Rouet-Leduc, C. Hulbert, and P. A. Johnson, "Continuous chatter of the Cascadia subduction zone revealed by machine learning," Nature Geoscience, vol. 12, 2019.
[24] T. H. W. Goebel and E. E. Brodsky, "The spatial footprint of injection wells in a global compilation of induced earthquake sequences," Science, vol. 361, 2018.
[25] J. Grigoli, S. Cesca, A. P. Rinaldi, J. F. Manconi, J. A. Lopez-Comino, J. Clinton, R. Westaway, C. Cauzzi, T. Dahm, and S. Wiemer, "The November 2017 Mw 5.5 Pohang earthquake: A possible case of induced seismicity in South Korea," Science, vol. 360, 2018.
[26] V. Kuehn, D. Snelson, and S. Malone, "An overview of injection-induced seismicity monitoring," Seismological Research Letters, vol. 90, 2019.
[27] K. R. Milner, E. H. Field, W. H. Savran, M. T. Page, and T. H. Jordan, "Operational earthquake forecasting during the 2019 Ridgecrest, California, earthquake sequence with the UCERF3-ETAS model," Seismological Research Letters, vol. 91, 2020.
[28] F. Cotton, F. Pacor, and D. Bindi, "Earthquake ground motion modeling: State of the art and future directions," Bulletin of Earthquake Engineering, vol. 18, 2020.
[29] S. Wiemer and M. Wyss, "Mapping spatial variability of the frequency-magnitude distribution of earthquakes," Advances in Geophysics, vol. 45, 2002.
[30] D. Schorlemmer, S. Wiemer, and M. Wyss, "Variations in earthquake-size distribution across different stress regimes," Nature, vol. 437, 2005.
[31] E. H. Field, T. H. Jordan, and C. A. Cornell, "OpenSHA: A developing community-modeling environment for seismic hazard analysis," Seismological Research Letters, vol. 74, 2003.
[32] D. J. Wald, V. Quitoriano, T. H. Heaton, and H. Kanamori, "Relationships between peak ground acceleration, peak ground velocity, and modified Mercalli intensity in California," Earthquake Spectra, vol. 15, 1999.
[33] K. Noda and T. Yamazaki, "Deep learning approaches for rapid post-earthquake building damage assessment," Earthquake Engineering & Structural Dynamics, vol. 49, 2020.
[34] H. Kanamori, "Real-time seismology and earthquake damage mitigation," Annual Review of Earth and Planetary Sciences, vol. 33, 2005.
[35] R. J. Geller, D. D. Jackson, Y. Y. Kagan, and F. Mulargia, "Earthquakes cannot be predicted," Science, vol. 275, 1997.
[36] T. H. Jordan, Y.-T. Chen, P. Gasparini, R. Madariaga, I. Main, W. Marzocchi, G. Papadopoulos, G. Sobolev, K. Yamaoka, and J. Zschau, "Operational earthquake forecasting: State of knowledge and guidelines for utilization," Annals of Geophysics, vol. 54, 2011.
[37] S. Uhlemann, J. Chambers, N. Falck, and P. Meldrum, "Applications of machine learning to geotechnical monitoring for hazard early warning," Engineering Geology, vol. 279, 2021.
[38] D. T. Trugman and P. M. Shearer, "Application of a stochastic clustering model to characterize induced seismicity," Bulletin of the Seismological Society of America, vol. 109, 2019.
[39] M. Corradini, F. Cotton, and D. Bindi, "Uncertainty quantification for machine-learning-based seismic hazard estimation," Natural Hazards and Earth System Sciences, vol. 22, 2022.