

Forecasting Induced Seismicity Rates in Geothermal Fields with QuakeCastNet, a Temporal-Fusion-Transformer-Based Model

Zhengfa Bi¹ and Nori Nakata^{1,2}

¹Energy Geosciences Division, Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, 94720, CA, USA

²Department of Earth, Atmospheric and Planetary Sciences, MIT, Cambridge, MA 02139, USA

Key Points:

- We develop a deep learning model to forecast seismicity rates driven by fluid injection.
- The model provides probabilistic forecasts with uncertainty quantification and parameter interpretability.
- Applications to the Geysers and Utah FORGE datasets demonstrate robust performance and generalizability.

Corresponding author: Zhengfa Bi, zfb1@lbl.gov

Abstract

Induced seismicity presents significant challenges to the sustainable development of geothermal energy, as fluid injection and reservoir stimulation frequently trigger seismic events. The underlying physical mechanisms involve complex interactions between the solid Earth and injected fluid, which makes the numerical flow modeling and mechanical failure analysis difficult, and leads to inaccurate forecasting of induced seismicity. This study introduces a deep learning framework that integrates multi-horizon forecasting with interpretability to predict seismicity in geothermal fields. By dynamically selecting relevant seismicity and operational data through variable selection networks, the model unifies key drivers within a single predictive architecture. The model, which consists of self-attention mechanisms and probabilistic forecasting, identifies key seismicity drivers and quantifies uncertainty to evaluate fluid injection on future seismicity. Case studies from two geothermal fields, the Geysers and Utah FORGE, reveal site-specific seismic response governed by hydromechanical conditions, offering insights into fluid-induced stress evolution. Beyond risk mitigation, this approach provides a data-driven pathway for probing subsurface rheology through seismic response to geothermal operations.

Plain Language Summary

Geothermal energy is a promising renewable resource, but its development is often hindered by induced earthquakes caused by fluid injection into underground reservoirs. Predicting when these seismic events will occur remains a major challenge because the physical processes involved are highly complex and not fully understood. In this study, we introduce a data-driven approach that improves the forecasting of seismic activity in geothermal fields. Our model uses historical earthquake data, planned injection operations, and other relevant factors to make accurate predictions while identifying key drivers of seismicity. By incorporating advanced deep learning techniques, the model can assess uncertainty in its forecasts and evaluate how changes in fluid injection might influence future seismic activity. Tests on two geothermal sites show that the model effectively captures the nonlinear relationships between injection and seismic response, providing probabilistic estimates of earthquake hazards. Beyond improving risk assessment, this approach offers new insights into how fluid injection affects underground stress conditions to better manage geothermal resources and mitigate seismic risks.

1 Introduction

Energy and resource extraction reshape subsurface stress and pore pressure, potentially triggering seismic events (Ellsworth, 2013; Chen et al., 2017; Grigoli et al., 2017; Schultz et al., 2020). Incorporating these factors into quantitative models is essential for understanding earthquake processes and mitigating seismic hazards (McGarr, 2014; Chen & Nakata, 2017; Goebel & Brodsky, 2018; Lee et al., 2019; Hager et al., 2021; Ritz et al., 2022). As geothermal energy becomes a key pillar of global carbon neutrality efforts (Gaucher et al., 2015; Chang & Nakata, 2022), the challenge of managing induced seismicity and associated ground motion in geothermal fields grows increasingly critical. Large seismic events raise public safety concerns, drive economic repercussions, and necessitate stringent mitigation strategies. These concerns are magnified with the expansion of Enhanced Geothermal Systems (EGS) (Tester et al., 2006; Ziagos et al., 2013) designed to extract energy from low-permeability reservoirs. To create a geothermal reservoir at preferable depth and temperature, we stimulate fractures using, for example, hydraulic fracturing and shearing, potentially with proppant. This process alters the subsurface stress significantly. Large induced earthquakes can potentially pause and/or terminate the geothermal resource exploration, and hence the accurate forecasting of induced seismicity becomes an urgent task. Such perturbations even extend beyond geothermal operations to include reservoir impoundment, oil and gas production,

and underground wastewater disposal (Spica, Nakata, et al., 2018; Spica, Perton, et al., 2018).

Many efforts over the past two decades have been made to forecast induced seismicity using statistical and physics-based approaches, each grounded in assumptions or prior geological knowledge of seismicity mechanisms. Among statistical methods, the epidemic-type aftershock sequence (ETAS) model (Ogata, 1988, 1998) has been widely adopted to represent seismicity as a stochastic point process in time, space, and magnitude (Vere-Jones, 2010). ETAS captures both background activity driven by tectonic loading and triggered aftershocks, and has been applied extensively to tectonic and induced sequences (Bachmann et al., 2011; Brodsky & Lajoie, 2013; Llenos & Michael, 2016; Kothari et al., 2020; Aochi et al., 2021; Sedghizadeh et al., 2024). For induced seismicity, however, the standard ETAS formulation requires modification to incorporate anthropogenic forcing such as fluid injection and poroelastic stress perturbations (Weingarten et al., 2015; Alghannam & Juanes, 2020; Shcherbakov, 2024).

To overcome these limitations, alternative models that directly relate seismicity rate to fluid injection have been proposed, assuming a linear relationship between seismicity and injection (Hainzl & Ogata, 2005; Bachmann et al., 2011; Brodsky & Lajoie, 2013; Holtzman et al., 2018). These approaches have been applied to geothermal fields, hydraulic fracturing, and wastewater injection sites worldwide, demonstrating that modifying injection protocols can influence seismic hazard (Improta et al., 2015; Garcia-Aristizabal, 2018; Langenbruch et al., 2018; Dempsey & Riffault, 2019). Recent advancements propose convolutional formulations that treat fluid injection as an external forcing function, where the seismicity rate is expressed as the response to injection through a kernel function that encapsulates the transient dynamics of the subsurface (Kim & Avouac, 2023). This perspective has been further generalized by framing induced seismicity as the response of a nonequilibrium thermodynamic system to external perturbations, with rheology-dependent kernels encoding viscoelastic relaxation properties of the subsurface (Shcherbakov, 2024; White & Nakata, 2025). Other developments include the use of rate-and-state friction models, which explicitly account for stress redistribution effects and have improved the realism of induced seismicity forecasts (Segall & Lu, 2015; Deng et al., 2016; Norbeck & Rubinstein, 2018; Zhai et al., 2019). Despite these advancements, geomechanical models often emphasize deterministic stress-strain evolution while overlooking the stochastic nature and complex interactions that drive seismicity.

Physics-based modeling provides complementary insights by explicitly simulating stress and strain changes driven by geothermal operations. Several mechanisms have been identified as potential drivers of induced seismicity. One prominent mechanism is pore pressure diffusion, which reduces effective normal stress and promotes fault slip (Shapiro & Dinske, 2009; Bachmann et al., 2012). Another is poroelastic stress transfer, where fluid injection redistributes stresses beyond the fluid diffusion front and potentially triggers seismic events (Atkinson et al., 2016; Bao & Eaton, 2016). Additionally, aseismic slip propagation has been observed to induce earthquakes at distances far from injection sites, sometimes outpacing the diffusion of pore pressure (Guglielmi et al., 2015; Bhattacharya & Viesca, 2019; Yeo et al., 2020; Wynants-Morel et al., 2020). While these models provide valuable physical interpretations, they struggle to quantify the relative contributions of different processes and to resolve the coupled interactions that drive seismic activity. Geomechanical modeling plays an important role in understanding the physical mechanisms of induced seismicity, but it remains constrained by the randomness of earthquake occurrence and the heterogeneous subsurface properties and structures. Although numerical modeling of induced seismicity using thermal-hydrological-mechanical-chemical (THMC) modeling is another way to explore (Smith et al., 2024), it requires high computational cost and detailed subsurface knowledge for accurate modeling, which limits its feasibility for routine forecasting.

Beyond traditional methods, modern data-driven approaches have shown promise. For example, random forest models using physics-based features have been applied to forecast

induced seismicity in Oklahoma, which identifies stress rate changes as dominant drivers of seismicity rate variations (Qin et al., 2022). A spatiotemporal ConvLSTM network has been proposed to capture both spatial correlations and temporal dependencies across global seismic maps, demonstrating improved predictive accuracy over classical methods (Zhang & Wang, 2023). Real-time forecasting efforts using big data streams, such as geo-acoustic and electromagnetic signals, have also shown promise and achieve up to 70% accuracy in weekly predictions over a seven-month campaign in China (Saad et al., 2023). Meanwhile, deep learning has been applied to synthetic megathrust simulations to identify precursory signals and predict complete ruptures (Blank & Morgan, 2021), and foundation models like transformers and graph neural networks have been evaluated for earthquake nowcasting over large spatiotemporal domains (Jafari et al., 2024). In the different time scale, WaveCastNet forecasts seismic wavefields using a part of the wavefields as input for earthquake early warning (Lyu et al., 2024).

In this study, we propose an interpretable deep learning framework, QuakeCastNet, designed for forecasting induced seismicity in geothermal fields. Based on the Temporal Fusion Transformer (TFT) (Lim et al., 2021), we design the framework to unify attention mechanisms with multi-horizon quantile forecasting, enabling the model to capture interactions among various contributing factors with probabilistic estimates. Multi-horizon forecasting enables simultaneous prediction across multiple future time steps, enabling the model to characterize evolving temporal dynamics beyond a single-step forecast. Attention layers provide interpretable insights into the temporal influence of key covariates without imposing restrictive assumptions. Beyond forecasting, these capabilities provide a data-driven approach to analyze seismic responses to fluid injection and reveal the key operational drivers of seismicity in geothermal fields.

QuakeCastNet is adapted to process the past seismicity rate in combination with past and scheduled future inputs relevant to geothermal operations (Figure 1). It produces probabilistic predictions across multiple quantiles with uncertainty estimates under varying operational conditions. We apply the model on two geothermal systems with distinct geological settings and operational histories, the Geysers and the Utah FORGE fields. By tailoring QuakeCastNet to the challenges of induced seismicity forecasting and demonstrating superior performance relative to conventional approaches, this study aims to provide not only a methodological advance but also a practical tool to support sustainable geothermal energy production and strengthen risk assessment.

2 Methodology

The proposed model leverages a deep transformer to integrate multiple types of time series, dynamically select relevant features, and capture complex temporal dependencies, which enable accurate and interpretable forecasting of induced seismicity rates.

2.1 Temporal Transformer Framework

The Temporal Fusion Transformer (TFT) is a widely used architecture for multi-horizon time series forecasting to integrate heterogeneous data and capture complex temporal dynamics (Lim et al., 2021). Compared to conventional models such as Long Short-Term Memory (LSTM) or Sequence-to-Sequence architectures, TFT offers broader applicability for complex forecasting scenarios. It supports both static and time-varying covariates, handles known future inputs, and incorporates quantile forecasting for probabilistic output. We extend this network as discussed below to align with the characteristics of seismicity and injection data to achieve reliable and interpretable forecasting. Our model processes three types of input: static metadata, time-dependent past observations, and planned future schedules. The static metadata comprises the mean and standard deviation of historical seismicity rates within each input window. Additionally, site-specific geology and well attributes, such as locations and depths, can also be incorporated as static metadata to provide




Figure 1. Overview of the QuakeCastNet architecture for induced seismicity forecasting. Details of each element are discussed in the Methodology section. The architecture effectively captures complex temporal dynamics, providing interpretable insights into the primary drivers of seismicity.

physical context. These covariates remain constant for each input and are passed through the static covariate encoder to inform and modulate temporal dynamics. Time-dependent past observations include historical seismicity rates and injection-related variables, while planned future schedules represent known operational inputs such as future injections or pressures. All inputs are processed through embeddings for categorical and contextual variables and linear transformations for continuous variables. The model architecture then integrates these components via a static covariate encoder, variable selection networks, sequence-to-sequence LSTM layers, and an interpretable self-attention decoder (Figure 1). The following sections provide a detailed exposition of these elements.

2.1.1 Static Covariate Encoder and Gating Mechanism

The model employs a Gated Residual Network (GRN) to encode static covariates that constrain temporal dynamics across the network, enabling downstream components to leverage contextual information. When the properties of the covariates shift due to evolving operational or geological conditions, the GRN adapts the latent representations accordingly. The static covariates are further reintroduced at multiple stages through a static enrichment layer, ensuring that evolving contextual signals continue to inform the forecasts and that predictions remain physically consistent. The GRN employs gating mechanisms to achieve dynamic feature selection, enabling the efficient processing of heterogeneous data with varying complexity. The GRN is formulated as:

$$\text{GRN}_\omega(a, c) = \text{LayerNorm}(a + \text{GLU}_\omega(\eta_1)), \quad (1)$$

$$\eta_1 = W_{1,\omega}\eta_2 + b_{1,\omega}, \quad \eta_2 = \text{ELU}(W_{2,\omega}a + W_{3,\omega}c + b_{2,\omega}), \quad (2)$$

where a and c are the primary input and the optional context vector, respectively and ω denotes weight sharing. The Exponential Linear Unit (ELU) activation maintains sensitivity to informative positive signals while suppressing noisy negative inputs. This dynamic control over magnitudes allows the GRN to adjust its expressiveness based on the relevance of the inputs. To control the non-linear contribution, we employ Gated Linear Units (GLUs) as gating mechanisms:

$$\text{GLU}_\omega(\gamma) = \sigma(W_{4,\omega}\gamma + b_{4,\omega}) \odot (W_{5,\omega}\gamma + b_{5,\omega}), \quad (3)$$

where $\sigma(\cdot)$ is the sigmoid activation, $W(\cdot)$ and $b(\cdot)$ are trainable weights and biases, and $\odot$ represents element-wise multiplication. GLUs suppress non-essential layers by producing outputs close to zero. Dropout is further applied before the gating layer to enhance generalization.

2.1.2 Variable Selection Network

Time-varying past and known future inputs are processed through separate variable selection networks, which dynamically assign feature weights at each time step to filter out irrelevant or noisy inputs. Let $\xi_t^j \in \mathbb{R}^{d_{\text{model}}}$ denote the transformed input of the j -th variable at time t , and $\mathcal{X}_t = [(\xi_t^1)^\top, \dots, (\xi_t^{m_\chi})^\top]^\top$ represent the concatenated vector of all past inputs. Variable selection weights are computed by passing $\mathcal{X}_t$ and an optional context vector c_s from static covariate encoder through GRN followed by a Softmax layer:

$$v_\chi^t = \text{Softmax}(\text{GRN}_{v_\chi}(\mathcal{X}_t, c_s)), \quad (4)$$

where $v_\chi^t \in \mathbb{R}^{m_\chi}$ contains the selection weights for all m_χ variables at time t . For static variables, the context vector c_s is omitted. Each input ξ_t^j is further refined using the corresponding GRN, where weights of GRN_ξ^j are shared across all time steps and $\tilde{\xi}$ denotes weighted representation of past inputs. The final weighted representation of past inputs is obtained as:

$$\tilde{\xi}_t = \sum_{j=1}^{m_\chi} v_\chi^{t,j} \tilde{\xi}_t^j = \sum_{j=1}^{m_\chi} v_\chi^{t,j} \text{GRN}_\xi^j(\xi_t^j), \quad (5)$$

where $v_\chi^{t,j}$ is the j -th element of v_χ^t modulating the feature contribution.

2.1.3 Sequence-to-Sequence LSTM

To capture the temporal dependencies for seismicity forecasting, TFT integrates sequence-to-sequence LSTM layers with interpretable multi-head attention mechanisms. The sequence-to-sequence LSTM layers process past observations and known future inputs separately to learn localized temporal patterns while initializing hidden states derived from static covariates. This design ensures that static constraints and long-term operational trends are embedded into the temporal dynamics.

2.1.4 Interpretable Multi-Head Self-attention

The interpretable multi-head self-attention decoder enhances the model's ability to identify long-term dependencies by emphasizing significant time steps and feature interactions. It extends the transformer-based attention mechanism, where standard self-attention computes weighted representations of input values $\mathbf{V} \in \mathbb{R}^{N \times d_V}$,

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = A(\mathbf{Q}, \mathbf{K})\mathbf{V} = \text{Softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d_{\text{attn}}}}\right)\mathbf{V}, \quad (6)$$

where $A(\mathbf{Q}, \mathbf{K})$ denotes the scaled dot-product attention, and d_{attn} represents the dimensionality of the query and key vectors. Multi-head attention extends this by projecting queries, keys, and values into multiple subspaces, enabling to learn diverse temporal patterns. To improve interpretability, we constrain multi-head attention to share values V across all heads and aggregate outputs additively:

$$\text{MultiHead}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \tilde{A}(\mathbf{Q}, \mathbf{K})\mathbf{V}\mathbf{W}_V, \quad (7)$$

where the unified attention weights are given by:

$$\tilde{A}(\mathbf{Q}, \mathbf{K}) = \frac{1}{m_H} \sum_{h=1}^{m_H} A(\mathbf{Q}\mathbf{W}_Q^{(h)}, \mathbf{K}\mathbf{W}_K^{(h)}). \quad (8)$$

This formulation allows each head to learn distinct temporal patterns while attending to a common feature space. The shared attention weights $\tilde{A}(\mathbf{Q}, \mathbf{K})$ provide a unified representation that enhances the analysis of significant temporal dependencies and seismicity drivers. For example, the model highlights the dominant role of injection rate variability in driving seismicity, offering actionable insights into how operational adjustments can mitigate seismic risks. This mechanism also enables the identification of critical precursors of induced earthquakes, such as abrupt increases in injection rates or shifts in hydrothermal pressure.

2.2 Uncertainty Modeling and Quantile Forecasting

To quantify uncertainty in seismicity forecasting, our model produces probabilistic predictions by generating multiple quantile forecasts for each future time step. The training objective is the quantile loss, which penalizes deviations between observed seismicity rates and the predicted quantiles.

Let y_t denote the observed seismicity rate at time t , computed as the number of earthquakes per unit time. Let $\hat{y}(q, t - \tau, \tau)$ represent the model's forecast of the seismicity rate at quantile level q for horizon τ , conditioned on historical seismicity and covariates including injection rate and treating pressure. The model parameters are denoted by θ , and Ω is the training data domain containing N samples. The set of output quantiles is Q ($Q = \{0.02, 0.1, 0.25, 0.5, 0.75, 0.9, 0.98\}$). In this formulation, uncertainty is evaluated separately at each forecast horizon, and the framework does not prescribe a systematic widening of intervals within a single prediction window. The resulting quantile bands reflect horizon-specific variability learned from the training data, rather than a smooth temporal growth of uncertainty away from the forecast origin in each window. The total loss is defined as:

$$\mathcal{L}(\Omega, \theta) = \sum_{y_t \in \Omega} \sum_{q \in Q} \frac{1}{N \tau_{\max}} \sum_{\tau=1}^{\tau_{\max}} \mathcal{Q}(y_t, \hat{y}(q, t - \tau, \tau), q), \quad (9)$$

where the quantile loss is given by

$$\mathcal{Q}(y, \hat{y}, q) = q(y - \hat{y})_+ + (1 - q)(\hat{y} - y)_+, \quad (10)$$

with $(x)_+ = \max(0, x)$. The forecasting horizon length is denoted by $\tau_{\max}$.

3 Result

In this section, two case studies from the Geysers and Utah FORGE geothermal fields demonstrate the model's ability to capture relationships between fluid-injection and seismic response.

3.1 Application in the Geysers Dataset

This section focuses on induced seismicity forecasting in the northwestern region of the Geysers geothermal field, using a seismic catalog from the Northern California Earthquake Data Center and injection data from two wells, Prati-9 and Prati-29 (Figure 2a). The Geysers is the world's largest geothermal field and has been operational since the 1960s and generates approximately 850 MW of electricity annually through steam production driven by extensive water injection (Martínez-Garzón et al., 2014; Kwiitek et al., 2015). Over 70 injection wells reaching depths of up to 5 km are used to enhance steam production and mitigate reservoir pressure depletion. The vapor-dominated reservoir consists of fractured, hydrothermally altered greywacke sandstone with low porosity and temperatures exceeding 350°C below 2.75 km in the northwestern part of the field.

Fluid injection at this region is gravity-driven, resulting in a negative gauge pressure at the wellhead after the permeability enhancement based on thermal contraction, in contrast to the high-pressure methods typically used for EGS stimulation (Garcia et al., 2016;




Figure 2. Seismicity and fluid injection data characteristics at the Geysers geothermal field. (a) Spatial distribution of seismic events recorded from January 2009 to August 2015, along with the locations of the Prati-9 and Prati-29 injection wells; (b) Temporal variations in seismicity rates and injection rates for Prati-9 and Prati-29 indicate their seasonal fluctuations due to the energy demands; (c) Earthquake distribution over time. The magnitude completeness is 1.37, and we use earthquakes with this or larger magnitudes for the analysis.




Figure 3. Comparison of observed seismicity rate (blue) and injection rate (green) with prediction (orange) and uncertainty bounds in the Geysers field, evaluated across different forecast times (gray). Panels (a–d) correspond to forecasts initiated at different times. Shaded orange regions denote forecast uncertainty, with the outer, middle, and inner bands corresponding to the 0.02–0.98, 0.10–0.90, and 0.25–0.75 quantile intervals, respectively.

Rutqvist et al., 2016). Induced seismicity in the region is primarily linked to water injection rather than steam production and is attributed to thermal stresses caused by subsurface rock expansion and contraction. The Geysers provides a unique environment for advancing methods to model the interplay between external forcing and seismic responses, with implications for geothermal energy development and seismic hazard mitigation.

3.1.1 Fluid Injection-Induced Seismicity Dynamics

The earthquake catalog (Martínez-Garzón et al., 2014; Leptokaropoulos et al., 2018) spans from January 2009 to August 2015 and includes 1254 seismic events. occurred up to a maximum distance of 600 m from the open hole of Prati-9 injection well (Figure 2b). These events are recorded by the Berkeley-Geyers (BG) seismic network, consisting of 31 three-component short-period geophones. Only events with magnitudes $M_w \geq 1.37$ are considered to ensure consistency with the magnitude of completeness.

Injection at Prati-9 began in November 2007 and continued for approximately 18 years, while injection at Prati-29 started in April 2010 and lasted for 3 years. Continuous fluid injection at Prati-9 totals 7.21 Mm^3 , while Prati-29 injects 3.29 Mm^3 , with both wells exhibiting average injection rates of $8.79 \times 10^4 \text{ m}^3/\text{month}$ (Kwiatek et al., 2015). The seismicity cloud surrounding Prati-9 (Figure 2c) shows a clear temporal correlation with injection activities, with seismicity pulsating in response to changes in injection rates. The average distance of seismic events from Prati-9 is 376 m, while only sparse activity (approximately 30 events) occurs near Prati-29, at an average distance of 800 m. Due to the limited seismicity near Prati-29, detailed analyses focuses on events near Prati-9, though the influence of total injection rates from both wells is considered.

3.1.2 Data Conditioning for Time-Series Modeling

Seismicity rates are derived by counting earthquake events within a uniform daily time grid, while fluid injection rates are interpolated onto this grid to ensure temporal alignment. To reduce short-term fluctuations, we apply a moving average with a one-month window and a one-day sliding step to smooth both time series. From a modeling perspective, the smoothing prevents the network from overfitting to fine-scale features and enhances its ability to capture the broader temporal structure. Given the large magnitude difference, we normalize each variable separately using z-score normalization while preserving the mean and standard deviation for subsequent de-normalization. To prevent information leakage, the dataset is divided into independent training and validation sets at the forecast time, with the validation set comprising only data from that point onward. The seismicity rates targeted for forecasting are excluded from both the inputs and supervision targets of the training set, which ensures that forecasts are strictly retrospective. Both training and validation sequences are segmented into 960-day windows with a sliding step of 4 days.

Each segment comprises up to 800 days of historical seismicity rate together with full injection rate and its gradient, with the final 320 days of seismicity rate used as the forecast horizon and training label. The injection rate gradient is computed using a finite-difference scheme, with first-order approximation. To capture a diverse range of temporal dependencies, the length of the historical input window is varied between 320 and 800 days, while the forecast horizon ranges from 160 to 320 days, with both randomly perturbed independently. The chosen horizon length is sufficiently long to account for potential temporal lags in seismic responses to injection activities, while the shorter bound of the window is set to prevent the model from overfitting to local fluctuations. The injection rate and its gradient are related by time derivatives, but they hold different spectral contents. For conventional inversion, time derivatives of variables provide additional constraints (Virieux & Operto, 2009), and hence we use it for the inputs of QuakeCastNet. Moreover, static covariates are incorporated to provide contextual information for each segment. These include the mean and standard deviation of historical seismicity rates and the input window length, enabling the model to better contextualize variations in temporal dynamics.

3.1.3 Training and Forecast Validation

The model is trained to forecast seismicity rates by capturing complex temporal dependencies and uncertainty distributions. The optimization process employs an adaptive

learning rate strategy that dynamically adjusts the learning rate based on model performance to enhance convergence. With a reduction-on-plateau mechanism, the learning rate is initialized at 0.03 and decreases when the model fails to improve within 4 consecutive evaluations. The loss gradient clipping at 0.1 stabilizes training by mitigating extreme gradient values. The model is trained for up to 200 epochs but incorporates early stopping with a patience of 10 to prevent unnecessary computations once convergence is reached.

Hyperparameter optimization is performed to refine model performance while ensuring computational efficiency. A search space spanning hidden dimensions from 8 to 128, attention head sizes from 1 to 4, and dropout rates between 0.1 and 0.3 is explored, while the learning rate is tuned within the range of 0.01 to 0.1. The tuning process conducted over 200 trials with a maximum training limit of 60 epochs per trial evaluates configurations to identify an optimal balance between learning speed and stability. The loss gradient clipping value is adjusted within the range of 0.01 to 1.0 to determine an appropriate level of stabilization. The selection of hyperparameters is guided by empirical performance metrics and stability considerations, ensuring the trained model generalizes well to unseen seismicity patterns.

To evaluate predictive performance, we conduct forecasts using a progressively expanding training window and a fixed 320-day forecasting horizon (Figure 3a-3d). Although this horizon is chosen for evaluation, the model is trained within a multi-horizon framework that enables flexible prediction across varying lead times. The evaluation is repeated across different forecast times (dashed gray line) to ensure that the assessment reflects performance under varying temporal contexts. The results show that the model captures the broad evolution of earthquake occurrence, with predicted seismicity rates (orange line) closely following the observed temporal dynamics (blue line). By incorporating uncertainty quantification, the forecasts provide probabilistic estimates of seismic hazard that enhance their utility for risk evaluation. The predictive uncertainty is expressed through the 5% and 95% percentiles of the forecast distribution, which form an interval (orange band) encompassing most local peaks and troughs. This indicates that the model characterizes the probability distribution of seismic responses in a statistical sense.

The forecasts faithfully reproduce the dominant long-term trends in seismicity, while finer-scale fluctuations remain only partially resolved. This indicates that the model can track the sustained evolution of seismic activity over extended horizons and to reflect how seismicity responds to operational forcing at broader temporal scales. This reflects the intrinsic stochasticity of earthquake occurrence and the difficulty of modeling abrupt rate changes. Deviations are most evident during phases of rapid seismicity shifts, where abrupt perturbations exceed the temporal resolution of the learned dynamics.

The observed seismicity curve often resembles the injection rate (dashed green line), suggesting a close relationship. However, the forecasts indicate that seismicity is not simply governed by injection rate alone, and the nature of this relationship requires further investigation. Instead, the model identifies patterns where predicted seismicity rates follow the observed seismicity more closely than the injection records alone. This demonstrates that the model captures not only the immediate stimulation effects but also more gradual responses that emerge from the cumulative influence of prior injections. These responses reflect general lagged effects. Forecast accuracy is reduced during sudden and isolated seismic bursts that exhibit only weak connections to external forcing.

3.2 Application in Utah FORGE Dataset

This section focuses on induced seismicity forecasting at the Utah FORGE geothermal site, utilizing high-resolution seismic and injection data collected from the 2024 hydraulic stimulation of wells 16A(78)-32 and 16B(78)-32 (Figure 4a). Utah FORGE is a DOE-funded dedicated field laboratory for Enhanced Geothermal Systems (EGS) research, targeting a granitic basement at approximately 2.5 km depth. The objective is to create a sustainable




Figure 4. Seismicity and fluid injection characteristics at the Utah FORGE geothermal site. (a) Spatial distribution of induced seismic events recorded during the 2024 hydraulic stimulation phases, along with the locations of the 16A(78)-32 and 16B(78)-32 injection wells; (b) Temporal variations in seismicity rates, injection rates, and treating pressures, illustrating their correlation; (c) Seismic magnitude distribution and cumulative events over time; While the magnitudes in the Geysers dataset are mostly concentrated above 0, the magnitudes in this dataset are distributed within the range of 0 to -1.

heat exchange reservoir by enhancing permeability through hydraulic stimulation. Unlike the Geysers geothermal field where fluid injection is gravity-driven, Utah FORGE employs high-pressure hydraulic fracturing methods to control fracture growth and connectivity. The interplay between injection rate, pressure, and induced seismicity provides a benchmark for advancing forecasting approaches to understand fluid-driven seismicity mechanisms.

3.2.1 Induced Seismicity and Injection Characteristics

The 2024 hydraulic stimulation at Utah FORGE spanned 10 stages, during which seismic events are continuously monitored to assess the response of the reservoir to injection activities. The earthquake catalog sourced from the Geothermal Data Repository (GDR) for the Utah FORGE project contains over 2,900 induced seismic events, with magnitudes ranging predominantly from M_L -1.0 to 0.0, and a maximum recorded event of M_L 2.0. Fluid injection is carried out in stages at wells 16A(78)-32 and 16B(78)-32, involving the injection of approximately 118,000 barrels (4.96 million gallons) of water over two weeks. The injection rate varied between 35 and 50 bbl per minute, with corresponding treating pressures reaching up to 50 MPa (7,000 psi) (Figure 4). The integration of treating pressure data provides additional insights into the stress evolution within the reservoir and its impact on induced seismicity.

Seismicity is primarily concentrated around the injection wells, aligning with the extent of stimulated fracture networks (Figure 4a). The spatial extent of seismicity varied across injection stages, with most events occurring within a few hundred meters of the wellbore. The seismicity cloud extended preferentially along the direction of the maximum horizontal stress (NNE-SSW), consistent with expected fracture orientations in the granitic host rock (Xing et al., 2020). Temporal trends indicate a strong correlation between seismic activity and injection parameters, with increased seismicity observed during high-rate injection phases and periods of pressure buildup (Figure 4b). Notably, seismic events continue to




Figure 5. Comparison of observed seismicity rate (blue), injection rate (green dash), treating pressure (purple), and prediction (orange) with uncertainty bound in the Utah FORGE field, evaluated across various forecast times (gray dash). Panels (a–d) correspond to forecasts initiated at different times.

occur post-injection, indicating the delayed response of the reservoir to fluid diffusion and stress redistribution.

3.2.2 Data Preprocessing and Feature Engineering

Seismicity rates are calculated by counting events within a uniform time grid, while injection rates and treating pressures are interpolated onto this grid to ensure temporal alignment. To suppress short-term fluctuations of seismicity rates, injection rates and treating pressures, we apply a moving average with a 2-minute window to smooth all time series. Each input variable is normalized using z-score normalization, preserving the mean and standard deviation for subsequent de-normalization. The dataset is partitioned into independent training and validation sets based on the forecast time, with no overlap between them. The validation set is constructed to ensure that target seismicity rates are not included in any training inputs or labels, thus guaranteeing retrospective forecasts. Time series are further segmented into 32-hour windows with a sliding step of 8 minutes.

Each segment contains up to 26 hours of historical seismicity rate and full injection rate and treating pressure as inputs, with the following 6 hours of seismicity rate as target. To improve generalization, the input window length is randomly varied between 6 and 26 hours, while the forecast horizon ranges between 3 and 6 hours. As in the longer-term case of the Geysers geothermal field, the mean and standard deviation for each input time series are used as static covariates to adapt different temporal dynamics.

3.2.3 Forecasting of Injection-Induced Seismicity

The optimization process employs an adaptive learning rate strategy, starting at 0.03 and reducing when model improvement stagnates over four consecutive evaluations. The loss gradient clipping at 0.1 stabilizes training by mitigating extreme gradients. The model is trained for up to 100 epochs with early stopping to prevent redundant computations upon convergence. As in the Geysers dataset application, hyperparameter tuning enhances performance while maintaining computational efficiency.




Figure 6. Iterative forecasting of seismicity rate in the Geysers (a) and the Utah FORGE (b) geothermal fields. The model is recursively applied across overlapping time windows, using previously estimated seismicity as input for subsequent step.

Model performance is evaluated with a progressively expanding time window across different forecast times and a fixed 4-hour forecasting horizon (Figure 5). The results indicate that the model captures the overall trends in seismicity and accurately reproduces the evolution of seismicity rate peaks, although finer-scale fluctuations remain challenging to resolve. Incorporating uncertainty quantification further enhances the robustness of the forecasts, strengthening their value for seismic hazard assessment.

4 Discussion

In this section, we reflect on the implications of our findings for forecasting induced seismicity and clarify the main challenges that remain. Our results show that iterative forecasts can be extended to long horizons with stable performance in the near term, but accumulating uncertainties highlight limits for operational deployment. Analysis of feature importance and attention weights indicates that the model captures site-specific physical regimes, with delayed responses at the Geysers contrasting with more immediate injection–seismicity links at Utah FORGE. Comparisons with variable combinations and alternative methods confirm that injection-related covariates are indispensable and that deep learning provides clear gains over traditional approaches.

4.1 Iterative Long-Horizon Forecasting

To extend the temporal coverage beyond a fixed horizon, we implement an iterative strategy where the predicted results are recursively used as input for subsequent forecasts (Figure 6). We divide the testing set that corresponds to the period after the forecast time (vertical gray dashed line) into sequential overlapping segments, each comprising 800 timesteps of historical data used to predict the next 320 timesteps. After each forecast, the outputs are concatenated with the portion of previous seismicity rates to form the new

Table 1. Feature Importance Analysis for the two Geothermal Fields.

Category	Variable	Dataset	
		Geysers	FORGE
Past Inputs	Injection Rate	8.0	42.6
	Injection Gradient	45.5	11.9
	Treating Pressure	–	30.0
	Pressure Gradient	–	6.7
	Past Seismicity Rate	46.5	8.8
Future Inputs	Injection Rate	15.0	46.7
	Injection Gradient	85.0	13.0
	Treating Pressure	–	32.9
	Pressure Gradient	–	7.4

context for the following step. This recursive process enables long-term forecasting across the entire set while preserving temporal consistency with smooth connection among different time windows.

Our model captures not only near-future dependencies but also remains robust under evolving seismic and injection conditions over time. In the Geysers dataset, they reproduce the major injection-related changes in seismicity, including the rising and falling phases of pronounced peaks observed in October 2012 and October 2013. The evidence implies that the network learns associations between injection dynamics and induced seismicity, rather than relying on short-term trend continuation.

To evaluate forecast reliability over time, we compute cumulative uncertainty bounds (gray shaded region) from averaged uncertainties (orange line) using an exponentially decaying kernel. The cumulative upper and lower bounds are calculated as weighted sums of deviations from the mean prediction. In the Geysers data application, cumulative uncertainty grows slowly at the early stage, reflecting stable predictions. Beyond approximately 12 months into the forecast horizon, accumulation error begins to rise more rapidly. The uncertainties accumulating over time are expected, and they are reduced as time passes and we assimilate additional data.

In the Utah FORGE data application, cumulative uncertainty grows slowly until about 12 hours after forecasting origin, when seismicity rates start to decline. At this stage, accumulation uncertainty increases rapidly, which is attributed to elapsed forecast time and the cessation of injection. These phases occur when seismicity no longer covaries closely with fluid-injection, which represents a transition in the temporal dynamics as reservoir conditions evolve. This underscores that the forecasts are sensitive to underlying physical changes in the reservoir system, rather than a mere statistical fit to short-term correlations.

In contrast, the Utah FORGE dataset exhibits a different forecasting pattern. Uncertainty remains relatively low near the forecasting origin, and the model captures short-term variations less precisely than in the Geysers. However, as the horizon extends, the Utah FORGE forecasts show a closer fit to the longer-term decline in seismicity rates. This contrast suggests that the model is more effective at reproducing fine-scale details in the Geysers in the very near term, while better capturing broader temporal trends in Utah FORGE at later times. A sharp spike in seismicity around 00:00 on April 7 cannot be explained by changes in injection rate or treating pressure, which remain stable during this period. Although not captured by the current input variables, such events indicate the need to estimate the probabilities of low-likelihood but high-impact occurrences, representing an important direction for future work.




Figure 7. Feature importance and temporal attention patterns in seismicity forecasting. Percentiles of attention weights $\alpha(t, p, \tau = 0)$ for one-step-ahead forecasting in the Geysers (a) and the Utah FORGE (b) datasets, respectively. The attention distribution highlights varying levels of predictive significance across past time steps. Mean attention weights across different forecasting horizons in the Geysers (c) and the Utah FORGE (d) datasets, showing the evolving impact of historical information on predictions. The dashed vertical line denotes the forecasting step.

4.2 Interpretability of Attention Weights

To explain the model’s decision-making process, we quantify relative contributions using percentiles of attention weights aggregated across the entire dataset. These weights can potentially provide us with the information of which parameter is important for controlling future seismicity and help us to interpret the results based on our understanding of physics and statistics. We compute the percentiles along the temporal axis for each input variable, and normalize their averaged values to yield an importance score (Table 1).

In the Geysers dataset, which is long-term forecasting, the model assigns dominant attention to historical injection gradient (45.5%) and past seismicity rate (46.5%), but not injection rate (8.0%). This highlights the cumulative influence of historical injection dynamics and seismic feedback mechanisms. Notably, the scheduled injection gradient also ranks highly (85.0%), indicating the model’s reliance on anticipated trends in injection variability. The high importance of the injection rate gradient suggests that the system is more sensitive to abrupt changes in injection than to the total injected volume, even though the injection rate is already normalized by time. This finding is consistent with Weingarten et al. (2015), who reported that, for induced seismicity associated with wastewater disposal, temporal changes in injection rate exert stronger control on seismicity than cumulative injection volume, which is likely explained by the propagation of pore-pressure increases. We interpret this sensitivity as a consequence of the reservoir’s high permeability, which enables rapid pressure diffusion and fault stimulation under near-critical stress conditions. The high importance of past seismicity rate indicates that QuakeCastNet captures recurrent patterns, which may arise from long-term tectonic loading or seasonal variability. Seasonal injection cycles induce corresponding seasonal fluctuations in seismicity, creating shared periodicity

between these variables. This overlap leads the model to allocate importance across injection and seismicity covariates, complicating efforts to evaluate their contributions under the present setup.

The Utah FORGE dataset reveals a distinct pattern shaped by its shorter forecasting window and injection schedule. The future injection rate (46.7%) and treating pressure (32.9%) emerge as the most significant features, which suggests the impact of operational conditions on induced seismicity. The relatively low contribution from past seismicity rate (8.8%) indicates that short-term induced responses are less influenced by historical seismic activity and more by ongoing fluid-injection operations. This low contribution confirms that the seismicity pattern in each stage is unique, and we need operational information to forecast seismicity better. Interestingly, the gradients show smaller contributions to forecast seismicity compared to the Geysers case. One reason is that for the Utah FORGE case, the data duration is short, and we use a short time window to smooth the time-series data. Hence, the time-lapse changes are well captured in the injection rate already, and the relative contribution of the injection rate gradient becomes smaller. The seismic response is closely tied to immediate injection controls, especially in a low-permeability reservoir. In such settings, seismicity tends to reflect direct pressure perturbations rather than long-term memory effects, contrasting with the persistent clustering observed in long-lived geothermal systems.

The contrast between the two fields demonstrates that the model goes beyond identifying statistical correlations, as its ability to generalize to unseen data using only operational inputs indicates reliance on learned site-specific physical regimes rather than short-term trend fitting. This suggests two possible interpretations. First, the difference may reflect contrasting reservoir permeabilities: at the Geysers, pore pressure changes may propagate more broadly due to higher permeability, while at Utah FORGE, pore-pressure gradients resulting from changes in injection rate may dissipate more quickly during propagation. An alternative interpretation is that treating pressure and injection provide similar physical effects in our model, and their importance is split across these correlated features. These provide insight into how geological and operational regimes shape the key drivers of induced seismicity.

QuakeCastNet also provides insights into the time lag between seismicity and operation parameters. Understanding how past events influence future seismicity requires further analysis of persistent temporal patterns to reveal the underlying mechanisms governing seismic evolution. These patterns are captured by the self-attention weights $\alpha(t, p, \tau)$, which quantify the contribution of past input features p , including both seismic and injection variables, at time t to forecasts at horizon τ . Distributions of $\alpha(t, p, \tau)$ across time steps reveal the relevance of past observations. For one-step-ahead forecasts, attention distributions over quantiles (10_{th}, 50_{th}, and 90_{th} percentiles) highlight memory effects in each field. In the Geysers dataset (Figure 7a), high attention is assigned to past observations spanning roughly 180 days before the forecast origin, with weights increasing steadily toward the present. This delayed response reflects the time required for pore-pressure perturbations to diffuse from the injection point to critically stressed faults and/or for stress to accumulate sufficiently to induce slip. The rising attention-weight trend near the present also indicates the increasing impact of recent seismic and injection history. The central quantile band (0.25–0.75) captures the typical range of realizations, while the wider intervals (0.10–0.90 and 0.02–0.98) represent increasingly rare but plausible fluctuations. In contrast, the Utah FORGE dataset (Figure 7b) exhibits a more immediate response, with attention weights peaking near the forecast time and decaying rapidly within 6 hours.

Figures 7c and 7d show how past data influences predictions across timescales over forecast horizons. In the Geysers dataset, attention weights rise steadily from roughly 180 days before the forecast origin and remain high up to the present, which indicates that forecasts are informed by a sustained memory. This persistent influence reinforces the interpretation of a delayed seismic response, reflecting memory effects driven by stress redistribution and




Figure 8. Comparison of forecasting performance in the different geothermal fields. (a–b) Forecasts from ETAS models (red) and QuakeCastNet (orange) in the Geysers. (c–d) Forecasts from ARX (red), GBDT (gray), and QuakeCastNet in Utah FORGE.

Table 2. Quantitative evaluation metrics obtained from the QuakeCastNet applied to the Geysers and Utah FORGE datasets using different combinations of input variables. The lowest error values and the highest correlation are highlighted in bold. The input variables are abbreviated as SR for seismicity rate, I for injection rate, P for treating pressure, and G for gradient.

Input Variables	MSE ($\times 10^{-1}$)	RMSE	MAE	Pearson Corr.	Pearson Dist.	MAPE (%)	SMAPE (%)
The Geysers							
SR	1.255	0.354	0.288	0.293	0.707	31.22	39.69
SR + I	0.279	0.167	0.128	0.818	0.182	21.73	19.71
SR + IG	0.261	0.162	0.121	0.848	0.152	20.63	18.57
FORGE							
SR	2.060	0.454	0.383	0.445	1.445	67.33	57.84
SR + I	0.263	0.162	0.129	0.831	0.169	18.95	18.60
SR + P	0.596	0.244	0.187	0.680	0.320	31.30	28.39
SR + IP	0.197	0.140	0.119	0.853	0.147	17.33	17.54
SR + IPG	0.179	0.134	0.108	0.901	0.100	17.180	16.026

modulated by subsurface rheology. In contrast, the Utah FORGE dataset shows a localized attention distribution, with weights peaking near the forecast time across all horizons. This demonstrates that seismicity is dependent on immediate stress perturbations, with limited influence from past injection patterns beyond a few hours. This time-lapse interpretation is consistent with the above interpretation for the pressure pulse related to the feature importance analysis (Table 1). The contrast highlights differences in the timescales of induced seismicity, driven by site-specific geological and hydromechanical conditions.

4.3 Impact of Input Variables on Model Performance

We assess the impact of input variables on model performance for the Geysers and Utah FORGE datasets (Table 2), using six evaluation metrics: mean squared error (MSE), root mean squared error (RMSE), mean absolute error (MAE), Pearson correlation, Pearson distance, mean absolute percentage error (MAPE), and symmetric mean absolute percentage error (SMAPE). In the Geysers data application, the model achieves improved performance when injection rate is included alongside seismicity rate (SR + I). Using seismicity rate alone yields relatively poor results across all metrics, with low correlation and high error values (SR). By contrast, the addition of injection rate reduces the errors by more than a

factor of four (MSE from 1.255 to 0.279) and raises the correlation markedly (from 0.293 to 0.818). This demonstrates that seismicity cannot be adequately forecast from its own history alone, but instead requires information on contemporaneous injection dynamics. Moreover, incorporating the gradient of injection rate further enhances predictive accuracy to capture the impact of rapid operational fluctuations that absolute injection values alone cannot represent (SR+IG). However, the improvement is smaller than suggested by the high importance of gradient in Table 1. This discrepancy reflects the compensatory role of correlated covariates. When gradient is excluded, much of its effect can be captured by injection rate and past seismicity, which together provide partial surrogates for the underlying dynamics. Thus, variable importance and performance gains are related but not in a one-to-one correspondence. Note that this evaluation is related to the importance of Table 1, but provides two different information.

In the Utah FORGE data application, the model achieves its best performance across all metrics when incorporating the full set of input variables. Injection rate substantially improves predictive skill (SR+I) relative to using seismicity rate alone (SR). Treating pressure (SR+P) also contributes as it governs the effective stress on fractures. The pressure fluctuations can trigger seismic slip on patches where faults are weak, before cumulative fluid volumes reach thresholds to make the wide range of the faults unstable. The importance of treating pressure (Table 1) further reinforces this interpretation, as it governs the effective stress on fractures. However, treating pressure yields inferior performance compared to injection rate (SR+I), suggesting that injection pressure alone may inadequately describe spatiotemporal pore-pressure increase, which triggers earthquakes without information on injection volume. Thus, while each variable contains partial information about the reservoir’s stress state, their integration (SR+IP) enables the model to capture complex interactions that drive induced seismicity. Moreover, their gradients (SR+IPG) offer an additional, though more minor, improvement. This validates the model’s ability to leverage multi-source temporal information and offer physical insights into the mechanisms of fluid-triggered seismicity. The performance across different variable combinations mirrors the pattern observed in the Geysers. Adding more variables consistently improves accuracy, with injection rate providing the most substantial gain. This suggests that the amount of fluid injected is the critical factor for forecasting seismicity, while injection pressure likely modulates fracture opening and closing through its relationship with S_{hmin} . However, the total fluid volume entering the fractures exerts a stronger control on seismic response, which shows that injection-driven processes are the dominant mechanism governing short-term induced seismicity at Utah FORGE.

4.4 Comparison with Other Forecasting Methods

We benchmark our forecasts against widely used stochastic Epidemic Type Aftershock Sequence (ETAS) models (Shcherbakov, 2024) in the Geysers dataset (Figure 8). ETAS is a stochastic temporal point process and its raw simulations generate oscillatory events rather than continuous sequences. To enable a fair comparison with our forecasts, we aggregate event counts within local windows and then apply a sliding window smoothing step.

The ETAS models are trained on past seismicity up to the forecast time without injection-related covariates. In injection-influenced systems, ETAS can partly capture the temporal clustering of seismicity, yet its forecasts diverge in magnitude from the observed rates (Figure 8a and 8b). The ETAS is conceived for tectonic sequences and models seismicity as a self-exciting point process with a stationary background rate without exogenous drivers (Trugman et al., 2016). It can replicate temporal clustering yet misestimate the level of activity in injection-influenced systems when operational conditions shift. This mismatch has been documented across multiple settings where nonstationary forcing from wastewater or geothermal injection controls rates and spatial patterns (Asanuma et al., 2014; Llenos & Michael, 2019), prompting extensions that allow time-varying background rates or limited covariate modulation (Bachmann et al., 2011; Mancini et al., 2021). However, these




Figure 9. The distribution of quantitative evaluation metrics for different forecasting methods (ARX, GBDT, and QuakeCastNet) applied to the Utah FORGE dataset across multiple forecasting windows. For each method, the box denotes the interquartile range, the horizontal line marks the median, and whiskers extend to 1.5 times the interquartile range. Individual points represent results from individual forecast windows.

variants still require strong parametric assumptions about how operations map into rate changes and may not capture delayed responses associated with pressure diffusion or regime shifts as reservoir conditions evolve. In contrast, the QuakeCastNet bypasses these restrictions by learning the dependencies between injection and seismic evolution directly from data without relying on explicit constraints. At longer forecast horizons, our model aligns more closely with observations by using external injection-related information. This indicates that injection variability is indispensable for establishing predictive capability, while the seismic persistence determines how that information is expressed in long-term forecasts.

To further quantify the marginal gains of deploying the QuakeCastNet, we benchmark its forecasts against two baseline methods in the Utah FORGE dataset: an autoregressive model with exogenous variables (ARX) and a gradient boosted decision tree (GBDT). The ARX model extends linear autoregression by incorporating injection-related covariates to capture lagged linear relationships between operational parameters and seismicity. In contrast, the GBDT model is a nonlinear ensemble method that sequentially combines decision trees to approximate complex input-output mappings.

The comparison (Figure 8c and 8d) highlights the limitations of both approaches. ARX achieves reasonable correlations when the relationship between injection and seismicity is approximately linear, but its errors remain substantially higher than those of the QuakeCastNet in cases where the seismic response exhibits nonlinear variability. GBDT improves upon ARX in some scenarios by modeling nonlinear effects, yet it tends to overfit local fluctuations and produces larger errors at longer horizons. By contrast, the QuakeCastNet consistently delivers the lowest errors (MSE, RMSE, and MAE) and the highest correlations across nearly all settings (Figure 9). In several evaluation windows the QuakeCastNet reduces RMSE by more than 40% relative to ARX and GBDT, while also achieving lower MAPE and SMAPE scores. These results confirm that while linear and tree-based models can capture parts of the signal, they fall short in describing the temporal dependencies and multi-horizon dynamics characteristic of injection-induced seismicity. The QuakeCastNet employs autoregressive memory and covariate interactions within a unified deep learning framework, achieving superior predictive skill to offer robustness for risk-oriented forecasting.




Figure 10. Calibration diagnostics of predicted uncertainties. Empirical coverage $\hat{p}(\alpha)$ versus nominal quantile level α is shown for (a) the Geysers and (b) the Utah FORGE datasets. The dashed gray lines denote perfect calibration $p(\alpha) = \alpha$. Colored lines represent forecasts from different horizon windows: for the Geysers, spanning July 2012 to May 2013; for Utah FORGE, spanning day 6, hour 0 to hour 11. Shading indicates the maximum Kolmogorov–Smirnov (KS) deviation from perfect calibration.

4.5 Challenges and Future Directions

4.5.1 Forecast Uncertainties and System Dynamics

Despite advances in deep learning–based seismicity forecasting, several challenges remain. In this study, uncertainty bands are derived from quantile-based probabilistic forecasts. The bands are not uniform, but instead reflect the temporal context of each interval, shaped by variability in past seismicity rates, injection operations, and associated covariates. Periods characterized by stronger fluctuations in seismicity or external drivers tend to yield wider uncertainty ranges. Although longer histories can stabilize forecasts, their influence is not linear, as the predictive distribution adapts dynamically to changing reservoir conditions. Thus, forecast uncertainty should be interpreted not as a fixed property of the model, but as conditional on the historical inputs and the evolving state of the reservoir system.

Model performance depends on the coupling between fluid injection and seismicity, which is shaped by field settings such as, for example, geology, rock types, depth, well locations, and fluid and proppant types. In this study, both the Geysers and Utah FORGE datasets exhibit relatively strong correlations between operational data and seismicity rates. However, predicting individual high-magnitude events that occur after injection has ceased remains a challenge. For example, in the Utah FORGE case, a pronounced peak in seismicity rate is observed around April 6, 23:00 (Figure 6b), which our model cannot detect. The nature of non-periodic features in seismicity, as well as limited information on injection, may degrade forecast accuracy. Addressing these challenges requires improved modeling of delayed triggering processes and the evolution of post-injection stress. A promising direction is to integrate geomechanical models that bridge empirical observations with physical principles. It offers mechanistic insights into stress redistribution, fault reactivation, and long-term reservoir stability. This would further enhance the reliability of data-driven forecasting in managing injection-induced seismic risk.

4.5.2 Calibration Diagnostics for Uncertainties

To assess the reliability of estimated uncertainties, we compute empirical coverage $p(\alpha)$ for each nominal quantile level α using sliding windows after the forecast time. For a given




Figure 11. Sensitivity of forecasting performance to historical window length in (a) the Geysers and (b) the Utah FORGE geothermal fields. The forecast window is fixed at 320 time samples.

quantile forecast $q_\alpha(t)$, empirical coverage is defined as

$$p(\alpha) = \frac{1}{T} \sum_{t=1}^T \mathbf{1}\{y_t \leq q_\alpha(t)\}, \quad (11)$$

where y_t denotes the observed seismicity rate and T is the number of predicted points. This diagnostic evaluates whether the observed frequency of outcomes within a given quantile interval matches the nominal level α .

In the Geysers dataset, most curves follow the identity line closely (Figure 10), with relatively small Kolmogorov–Smirnov (KS) deviations ($KS \approx 0.10$ – 0.20), where the KS statistic measures the maximum vertical distance between the empirical coverage curve and the ideal identity line.] This indicates that the predicted intervals provide a reasonable probabilistic description of seismicity rates beyond the training period. In Utah FORGE dataset, calibration is also broadly satisfactory, but deviations increase for mid-quantile levels ($\alpha \sim 0.4$ – 0.6), where empirical coverage systematically falls below the nominal level. These diagnostics reveal that the quantile-based uncertainty estimates can be trusted at a coarse level for both sites, while highlighting scenarios where calibration deteriorates. Future work needs to incorporate post-hoc calibration methods or distributional consistency checks to further improve the reliability of uncertainty quantification in induced seismicity forecasting.

4.5.3 Constraints from Historical Window Selection

Window lengths are one of the important parameters for QuakeCastNet. Here, we conduct an experiment to understand the effect of the segment length of each historical

window while fixing the number of training samples. This analysis is applied to the Geysers and Utah FORGE datasets (Figure 11a and 11b).

The results reveal a clear dependence of forecasting skill on the length of the historical context. When the historical time window is relatively short (≤ 200 timesteps), predictive accuracy deteriorates rapidly across all evaluation metrics, which reflects the inability of the model to capture temporal dependencies from insufficient history. At the other extreme, excessively long windows (> 800 timesteps) also degrade performance due to the inclusion of distant information that contributes noise rather than relevant predictive signals. Optimal forecasting skill emerges when the historical length is approximately twice the prediction horizon (around 700 timesteps), which indicates a balance between retaining sufficient historical context and avoiding dilution by irrelevant long-term fluctuations. For example, the model requires at least one full horizon of history to capture recurrent patterns and lagged responses, while an additional horizon provides redundancy that improves generalization. These underscore that reliable forecasts depend on selecting an appropriate historical window, and this kind of experiments using a validation dataset helps us determine the window length. Ensuring that the model is provided with enough but not excessive historical data in each window is essential for accurate forecasting of induced seismicity.

5 Conclusions

This study presents an interpretable deep learning framework, QuakeCastNet, for forecasting induced seismicity in geothermal fields. Building on the Temporal Fusion Transformer, the approach performs quantile forecasting with variable selection and attention mechanisms, enabling accurate predictions with physically meaningful explanations. Applications to the Geysers and Utah FORGE datasets demonstrate consistent gains over conventional baselines. Note that the time scale is very different (years vs. days) between these datasets, but QuakeCastNet works well on both scales, which demonstrates the robustness of the network, although the interpretability of the model provides different physical insights into how the injected fluid interacts with the solid Earth. The interpretability highlights contrasting controls across sites. At the Geysers, forecasting is driven mainly by the historical injection gradient and past seismicity rate, with a relatively minor role for absolute injection volumes. This reflects the influence of recurrent seismic feedback and changes in the injection rate, consistent with rapid pressure diffusion in a high-permeability reservoir. The model also captures seasonal injection cycles, which overlap with seismic recurrence and split importance across correlated covariates. By contrast, Utah FORGE shows a short-term response governed by immediate operational conditions. Future injection rate and treating pressure dominate, while past seismicity contributes little. This indicates that seismicity in a low-permeability reservoir is controlled by rapid pressure perturbations, with minimal memory effects, and that injection rates already capture the variability without explicit gradient information. Forecasting remains limited by the difficulty of capturing non-periodic seismic features, post-injection events, and partially calibrated uncertainty estimates, and its skill is further constrained by sensitivity to historical window length. Looking ahead, integrating this data-driven framework with geomechanical models and spatially resolved fault information could provide a more mechanistic basis for forecasts. Such hybrid approaches may enable more reliable assessment of delayed triggering processes and long-term reservoir evolution to support more sustainable subsurface energy operations.

Acknowledgments

Funding provided by DOE EERE Geothermal Technologies Office (GTO) to Utah FORGE and the University of Utah under Project DE-EE0007080 Enhanced Geothermal System Concept Testing and Development at the Milford City, Utah Frontier Observatory for Research in Geothermal Energy (Utah FORGE) site. We are also grateful to Utah FORGE, Calpine and GTO to maintain the datasets used in this study. The authors also acknowledge the support of the Core Service Anthropogenic Hazards and the IS-EPOS Platform for

making seismicity-related data available. Development of the induced seismicity forecasting framework also made use of the open-source PyTorch Forecasting library (Beitner, 2025), which provides foundational implementation of the Temporal Fusion Transformer.

Open Research

Raw injection data and the seismicity catalog for the Geysers geothermal field are available through the IS-EPOS platform of the Core Service Anthropogenic Hazards at <https://tcs.ah-epos.eu/>, accessible after registration and provision of institutional affiliation (IS-EPOS Platform, 2022). The injection and seismic datasets for the Utah FORGE geothermal field are available via the U.S. Department of Energy’s Geothermal Data Repository hosted on OpenEI at <https://gdr.openei.org> (Utah FORGE, 2023). All code and Jupyter notebooks used for data preprocessing, model training, forecasting, and figure generation are archived in a publicly accessible Zenodo repository at <https://doi.org/10.5281/zenodo.15466652> (Bi & Nakata, 2025).

References

- Alghannam, M., & Juanes, R. (2020). Understanding rate effects in injection-induced earthquakes. *Nature communications*, 11(1), 3053.
- Aochi, H., Maury, J., & Le Guenan, T. (2021). How do statistical parameters of induced seismicity correlate with fluid injection? case of oklahoma. *Seismological Society of America*, 92(4), 2573–2590.
- Asanuma, H., Eto, T., Adachi, M., Saeki, K., Aoyama, K., Ozeki, H., & Häring, M. (2014). Seismostatistical characterization of earthquakes from geothermal reservoirs. In *Proceedings of thirty-ninth workshop on geothermal reservoir engineering*.
- Atkinson, G. M., Eaton, D. W., Ghofrani, H., Walker, D., Cheadle, B., Schultz, R., ... others (2016). Hydraulic fracturing and seismicity in the western canada sedimentary basin. *Seismological research letters*, 87(3), 631–647.
- Bachmann, C. E., Wiemer, S., Goertz-Allmann, B., & Woessner, J. (2012). Influence of pore-pressure on the event-size distribution of induced earthquakes. *Geophysical Research Letters*, 39(9).
- Bachmann, C. E., Wiemer, S., Woessner, J., & Hainzl, S. (2011). Statistical analysis of the induced basel 2006 earthquake sequence: introducing a probability-based monitoring approach for enhanced geothermal systems. *Geophysical Journal International*, 186(2), 793–807.
- Bao, X., & Eaton, D. W. (2016). Fault activation by hydraulic fracturing in western canada. *science*, 354(6318), 1406–1409.
- Beitner, J. (2025). *Pytorch forecasting*. GitHub. Retrieved from <https://github.com/sktime/pytorch-forecasting/releases/tag/v1.4.0> ([Software])
- Bhattacharya, P., & Viesca, R. C. (2019). Fluid-induced aseismic fault slip outpaces pore-fluid migration. *Science*, 364(6439), 464–468.
- Bi, Z., & Nakata, N. (2025, May). *Code and notebooks for forecasting induced seismicity in geothermal fields with interpretable deep learning*. Zenodo. Retrieved from <https://doi.org/10.5281/zenodo.17213298> doi: 10.5281/zenodo.17213298
- Blank, D., & Morgan, J. (2021). Can deep learning predict complete ruptures in numerical megathrust faults? *Geophysical Research Letters*, 48(18), e2021GL092607.
- Brodsky, E. E., & Lajoie, L. J. (2013). Anthropogenic seismicity rates and operational parameters at the salton sea geothermal field. *Science*, 341(6145), 543–546.
- Chang, H., & Nakata, N. (2022). Investigation of time-lapse changes with das borehole data at the brady geothermal field using deconvolution interferometry. *Remote Sensing*, 14(1), 185.
- Chen, X., & Nakata, N. (2017). Preface to the focus section on the 3 september 2016 pawnee, oklahoma, earthquake. *Seismological Research Letters*, 88(4), 953–955.

- Chen, X., Nakata, N., Pennington, C., Haffener, J., Chang, J. C., He, X., ... Walter, J. I. (2017). The pawnee earthquake as a result of the interplay among injection, faults and foreshocks. *Scientific Reports*, 7(1), 4945.
- Dempsey, D., & Riffault, J. (2019). Response of induced seismicity to injection rate reduction: Models of delay, decay, quiescence, recovery, and oklahoma. *Water Resources Research*, 55(1), 656–681.
- Deng, K., Liu, Y., & Harrington, R. M. (2016). Poroelastic stress triggering of the december 2013 crooked lake, alberta, induced seismicity sequence. *Geophysical Research Letters*, 43(16), 8482–8491.
- Ellsworth, W. L. (2013). Injection-induced earthquakes. *science*, 341(6142), 1225942.
- Garcia, J., Hartline, C., Walters, M., Wright, M., Rutqvist, J., Dobson, P. F., & Jeanne, P. (2016). The northwest geysers egs demonstration project, california: Part 1: characterization and reservoir response to injection. *Geothermics*, 63, 97–119.
- Garcia-Aristizabal, A. (2018). Modelling fluid-induced seismicity rates associated with fluid injections: examples related to fracture stimulations in geothermal areas. *Geophysical Journal International*, 215(1), 471–493.
- Gaucher, E., Schoenball, M., Heidbach, O., Zang, A., Fokker, P. A., van Wees, J.-D., & Kohl, T. (2015). Induced seismicity in geothermal reservoirs: A review of forecasting approaches. *Renewable and Sustainable Energy Reviews*, 52, 1473–1490.
- Goebel, T. H., & Brodsky, E. E. (2018). The spatial footprint of injection wells in a global compilation of induced earthquake sequences. *Science*, 361(6405), 899–904.
- Grigoli, F., Cesca, S., Priolo, E., Rinaldi, A. P., Clinton, J. F., Stabile, T. A., ... Dahm, T. (2017). Current challenges in monitoring, discrimination, and management of induced seismicity related to underground industrial activities: A european perspective. *Reviews of Geophysics*, 55(2), 310–340.
- Guglielmi, Y., Cappa, F., Avouac, J.-P., Henry, P., & Elsworth, D. (2015). Seismicity triggered by fluid injection–induced aseismic slip. *Science*, 348(6240), 1224–1226.
- Hager, B. H., Dieterich, J., Frohlich, C., Juanes, R., Mantica, S., Shaw, J. H., ... others (2021). A process-based approach to understanding and managing triggered seismicity. *Nature*, 595(7869), 684–689.
- Hainzl, S., & Ogata, Y. (2005). Detecting fluid signals in seismicity data through statistical earthquake modeling. *Journal of Geophysical Research: Solid Earth*, 110(B5).
- Holtzman, B. K., Paté, A., Paisley, J., Waldhauser, F., & Repetto, D. (2018). Machine learning reveals cyclic changes in seismic source spectra in geysers geothermal field. *Science advances*, 4(5), eaao2929.
- Improta, L., Valoroso, L., Piccinini, D., & Chiarabba, C. (2015). A detailed analysis of wastewater-induced seismicity in the val d’agri oil field (italy). *Geophysical Research Letters*, 42(8), 2682–2690.
- IS-EPOS Platform. (2022). *Induced seismicity episode: The geysers geothermal field, california*. [Dataset]. IS-EPOS: Thematic Core Service Anthropogenic Hazards. doi: 10.20366/ah-epos-2022.001
- Jafari, A., Fox, G., Rundle, J. B., Donnellan, A., & Ludwig, L. G. (2024). Time series foundation models and deep learning architectures for earthquake temporal and spatial nowcasting. *GeoHazards*, 5(4), 1247–1274.
- Kim, T., & Avouac, J.-P. (2023). Stress-based and convolutional forecasting of injection-induced seismicity: application to the otaniemi geothermal reservoir stimulation. *Journal of Geophysical Research: Solid Earth*, 128(4), e2022JB024960.
- Kothari, S., Shcherbakov, R., & Atkinson, G. (2020). Statistical modeling and characterization of induced seismicity within the western canada sedimentary basin. *Journal of Geophysical Research: Solid Earth*, 125(12), e2020JB020606.
- Kwiatk, G., Martínez-Garzón, P., Dresen, G., Bohnhoff, M., Sone, H., & Hartline, C. (2015). Effects of long-term fluid injection on induced seismicity parameters and maximum magnitude in northwestern part of the geysers geothermal field. *Journal of Geophysical Research: Solid Earth*, 120(10), 7085–7101.

- Langenbruch, C., Weingarten, M., & Zoback, M. D. (2018). Physics-based forecasting of man-made earthquake hazards in oklahoma and kansas. *Nature communications*, 9(1), 3946.
- Lee, K.-K., Ellsworth, W. L., Giardini, D., Townend, J., Ge, S., Shimamoto, T., ... others (2019). Managing injection-induced seismic risks. *Science*, 364(6442), 730–732.
- Leptokaropoulos, K., Staszek, M., Lasocki, S., Martínez-Garzón, P., & Kwiatak, G. (2018). Evolution of seismicity in relation to fluid injection in the north-western part of the geysers geothermal field. *Geophysical Journal International*, 212(2), 1157–1166.
- Lim, B., Arik, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764.
- Llenos, A. L., & Michael, A. J. (2016). Characterizing potentially induced earthquake rate changes in the brawley seismic zone, southern california. *Bulletin of the Seismological Society of America*, 106(5), 2045–2062.
- Llenos, A. L., & Michael, A. J. (2019). Ensembles of etas models provide optimal operational earthquake forecasting during swarms: Insights from the 2015 san ramon, california swarm. *Bulletin of the Seismological Society of America*, 109(6), 2145–2158.
- Lyu, D., Nakata, R., Ren, P., Mahoney, M. W., Pitarka, A., Nakata, N., & Erichson, N. B. (2024). Wavecastnet: an ai-enabled wavefield forecasting framework for earthquake early warning. *arXiv preprint arXiv:2405.20516*.
- Mancini, S., Werner, M. J., Segou, M., & Baptie, B. (2021). Probabilistic forecasting of hydraulic fracturing-induced seismicity using an injection-rate driven etas model. *Seismological Society of America*, 92(6), 3471–3481.
- Martínez-Garzón, P., Kwiatak, G., Sone, H., Bohnhoff, M., Dresen, G., & Hartline, C. (2014). Spatiotemporal changes, faulting regimes, and source parameters of induced seismicity: A case study from the geysers geothermal field. *Journal of Geophysical Research: Solid Earth*, 119(11), 8378–8396.
- McGarr, A. (2014). Maximum magnitude earthquakes induced by fluid injection. *Journal of Geophysical Research: solid earth*, 119(2), 1008–1019.
- Norbeck, J., & Rubinstein, J. L. (2018). Hydromechanical earthquake nucleation model forecasts onset, peak, and falling rates of induced seismicity in oklahoma and kansas. *Geophysical Research Letters*, 45(7), 2963–2975.
- Ogata, Y. (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical association*, 83(401), 9–27.
- Ogata, Y. (1998). Space-time point-process models for earthquake occurrences. *Annals of the Institute of Statistical Mathematics*, 50, 379–402.
- Qin, Y., Chen, T., Ma, X., & Chen, X. (2022). Forecasting induced seismicity in oklahoma using machine learning methods. *Scientific Reports*, 12(1), 9319.
- Ritz, V. A., Rinaldi, A. P., & Wiemer, S. (2022). Transient evolution of the relative size distribution of earthquakes as a risk indicator for induced seismicity. *Communications Earth & Environment*, 3(1), 249.
- Rutqvist, J., Jeanne, P., Dobson, P. F., Garcia, J., Hartline, C., Hutchings, L., ... Walters, M. (2016). The northwest geysers egs demonstration project, california—part 2: modeling and interpretation. *Geothermics*, 63, 120–138.
- Saad, O. M., Chen, Y., Savvaidis, A., Fomel, S., Jiang, X., Huang, D., ... others (2023). Earthquake forecasting using big data and artificial intelligence: A 30-week real-time case study in china. *Bulletin of the Seismological Society of America*, 113(6), 2461–2478.
- Schultz, R., Skoumal, R. J., Brudzinski, M. R., Eaton, D., Baptie, B., & Ellsworth, W. (2020). Hydraulic fracturing-induced seismicity. *Reviews of Geophysics*, 58(3), e2019RG000695.
- Sedghizadeh, M., van den Berghe, M., & Shcherbakov, R. (2024). Leveraging the etas model to forecast mining microseismicity. *Geophysical Journal International*, 238(3), 1491–1504.
- Segall, P., & Lu, S. (2015). Injection-induced seismicity: Poroelastic and earthquake nucle-

- ation effects. *Journal of Geophysical Research: Solid Earth*, 120(7), 5082–5103.
- Shapiro, S. A., & Dinske, C. (2009). Scaling of seismicity induced by nonlinear fluid-rock interaction. *Journal of Geophysical Research: Solid Earth*, 114(B9).
- Shcherbakov, R. (2024). A stochastic model for induced seismicity at the geothermal systems: A case of the geysers. *Seismological Research Letters*, 95(6), 3545–3556.
- Smith, J. T., Sonnenthal, E. L., & Nakata, N. (2024). Thermal hydrological mechanical modelling of anticipated stimulation in utah. *Geothermal Reservoir Engineering*, SGP-TR-227.
- Spica, Z. J., Nakata, N., Liu, X., Campman, X., Tang, Z., & Beroza, G. C. (2018). The ambient seismic field at groningen gas field: An overview from the surface to reservoir depth. *Seismological Research Letters*, 89(4), 1450–1466.
- Spica, Z. J., Perton, M., Nakata, N., Liu, X., & Beroza, G. C. (2018). Site characterization at groningen gas field area through joint surface-borehole h/v analysis. *Geophysical Journal International*, 212(1), 412–421.
- Tester, J. W., Anderson, B., Batchelor, A., Blackwell, D., DiPippo, R., Drake, E., et al. (2006). The future of geothermal energy-impact of enhanced geothermal systems (egs) on the united states in the 21st century: An assessment. *Idaho Falls: Idaho National Laboratory*, 1e8.
- Trugman, D. T., Shearer, P. M., Borsa, A. A., & Fialko, Y. (2016). A comparison of long-term changes in seismicity at the geysers, salton sea, and coso geothermal fields. *Journal of Geophysical Research: Solid Earth*, 121(1), 225–247.
- Utah FORGE. (2023). *Utah forge: Geothermal field datasets*. [Dataset]. U.S. Department of Energy Geothermal Data Repository. doi: 10.15121/1452741
- Vere-Jones, D. (2010). Foundations of statistical seismology. *Pure and applied geophysics*, 167, 645–653.
- Virieux, J., & Operto, S. (2009). An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74. doi: 10.1190/1.3238367
- Weingarten, M., Ge, S., Godt, J. W., Bekins, B. A., & Rubinstein, J. L. (2015). High-rate injection is associated with the increase in us mid-continent seismicity. *Science*, 348(6241), 1336–1340.
- White, M. C., & Nakata, N. (2025). Induced seismicity and geothermal energy production in the salton sea geothermal field, california. *Scientific Reports*, 15(1), 1–14.
- Wynants-Morel, N., Cappa, F., De Barros, L., & Ampuero, J.-P. (2020). Stress perturbation from aseismic slip drives the seismic front during fluid injection in a permeable fault. *Journal of Geophysical Research: Solid Earth*, 125(7), e2019JB019179.
- Xing, P., McLennan, J., & Moore, J. (2020). In-situ stress measurements at the utah frontier observatory for research in geothermal energy (forge) site. *Energies*, 13(21), 5842.
- Yeo, I., Brown, M. R., Ge, S., & Lee, K. (2020). Causal mechanism of injection-induced earthquakes through the mw 5.5 pohang earthquake case study. *Nature communications*, 11(1), 2614.
- Zhai, G., Shirzaei, M., Manga, M., & Chen, X. (2019). Pore-pressure diffusion, enhanced by poroelastic stresses, controls induced seismicity in oklahoma. *Proceedings of the National Academy of Sciences*, 116(33), 16228–16233.
- Zhang, Z., & Wang, Y. (2023). A spatiotemporal model for global earthquake prediction based on convolutional lstm. *IEEE Transactions on Geoscience and Remote Sensing*, 61, 1–12.
- Ziagos, J., Phillips, B. R., Boyd, L., Jelacic, A., Stillman, G., & Hass, E. (2013). *A technology roadmap for strategic development of enhanced geothermal systems* (Tech. Rep.). Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States).