Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Quantifying Epistemic Uncertainty in Ground Motion Prediction Using Bayesian Neural Networks: Application to Peninsular India

Yellapragada Meenakshi

meena.yellapragada@gmail.com

Indian Institute of Technology Madras

STG. Raghukanth

Indian Institute of Technology Madras

Research Article

Keywords:

Posted Date: October 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10823123/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Quantifying Epistemic Uncertainty in Ground Motion Prediction Using Bayesian Neural Networks: Application to Peninsular India

Yellapragada Meenakshi[1] (corresponding author), S. T. G. Raghukanth[1]

1 Indian Institute of Technology Madras, Chennai, Tamil Nadu, India

Corresponding author: Meenakshi Yellapragada; Email: meena.yellapragada@gmail.com

Abstract

Ground-motion models (GMMs) for probabilistic seismic hazard analysis (PSHA) in data-sparse regions are constrained by limited recordings and incomplete sampling of magnitude, distance, and site conditions. Inadequate characterization of epistemic uncertainty can therefore produce overconfident ground-motion predictions and poorly constrained hazard estimates. This study develops a Bayesian neural network GMM (BNN-GMM) for Peninsular India that separately characterizes epistemic uncertainty and aleatory variability. The model predicts PGA, PGV, and 5%-damped pseudo-spectral acceleration at 20 periods from 0.01 to 3.0 s as functions of moment magnitude, Joyner–Boore distance, V_{S30}, and focal depth, using data compiled from NGA-East, regional strong-motion records, and the NEar-Source Strong-motion flat file. The network contains two Bayesian hidden layers with mean-field Gaussian parameter distributions and is trained through variational inference, with its architecture and hyperparameters selected by Bayesian optimization. On the test dataset, the model achieves an $R^2$of 0.817 and a mean squared error of 0.824 while reproducing physically consistent magnitude, distance, site, depth, and spectral scaling. Epistemic uncertainty is quantified from Monte Carlo samples of the posterior network parameters, whereas aleatory variability is estimated by mixed-effects decomposition of logarithmic residuals. Aleatory standard deviations range from 0.76 to 0.89 natural-log units, compared with scenario-dependent epistemic standard deviations of approximately 0.04–0.11. Validation against significant PI earthquakes and established GMMs demonstrates consistent attenuation and spectral trends. PSHA sensitivity analyses demonstrate that propagation of BNN-derived epistemic uncertainty through the hazard integral produces systematic increases in spectral acceleration estimates relative to an aleatory-only treatment, with the magnitude of this effect varying as a function of site, spectral period, and annual exceedance rate, and becoming most pronounced at long return periods for Mumbai and Hyderabad. This site and exceedance rate dependent behaviour indicates that epistemic uncertainty cannot be represented adequately by a uniform adjustment to the median ground-motion prediction. Instead, it should be propagated through the hazard calculation using period- and scenario-dependent epistemic branches to obtain a more complete characterization of ground-motion exceedance probabilities in regions with sparse strong-motion data.

1. Introduction

Probabilistic seismic hazard analysis (PSHA) critically depends on two primary components: the seismic source model, which characterizes earthquake occurrence rates and magnitudes, and the ground-motion model (GMM), which predicts the resulting shaking at a site. Among these, the GMM plays a crucial role, as hazard estimates are governed not only by earthquake occurrence but also by how ground motions scale with magnitude, distance, and site conditions. However, GMMs inherently carry significant epistemic uncertainty arising from incomplete knowledge, which stems from limited and heterogeneous datasets, sparse coverage across magnitude–distance–site parameter space, alternative functional forms, parameter uncertainty, and regional ground-motion behavior (Aldama-Bustos et al., 2022). Numerous studies have demonstrated that uncertainty associated with GMMs can dominate the overall variability in PSHA results (Bradley et al., 2012; Porter et al., 2012; Pecker et al., 2017).

In contemporary PSHA practice, epistemic uncertainty is commonly represented using logic trees, where alternative branches correspond to different functional forms or competing GMMs. However, relying on multiple GMMs within a logic-tree framework is often insufficient to capture the full range of uncertainty in the target region. As an alternative, several studies have adopted the "backbone" approach, wherein a single base GMM derived for a host region is systematically modified to the target region and represented through logic-tree branches to account for source, path, and site characteristics (Boore, 2018; Boore et al., 2022). In addition, several approaches such as expert-judgment weighting, model ensembles, and non-ergodic formulations have been proposed to better characterize epistemic uncertainty.

Nevertheless, a critical requirement in GMM development is not only to accurately predict median ground motion but also to robustly quantify epistemic uncertainty, such that the resulting hazard curves and associated fractiles capture the range of epistemically consistent seismic hazard. In this context, Bayesian frameworks provide a rigorous probabilistic alternative by explicitly treating model parameters as random variables and quantifying uncertainty in a systematic manner. The Bayesian approach to neural networks, pioneered by MacKay (1992) and Neal (1996), and later made computationally tractable through approaches such as Monte Carlo dropout (Gal and Ghahramani, 2016), offers a powerful means to incorporate uncertainty directly into data-driven models.

In Bayesian neural network-based GMMs (BNN-GMMs), the network weights are treated as random variables rather than deterministic parameters, resulting in predictive distributions instead of single-point estimates. Consequently, the model yields a full probability distribution of ground motion conditioned on input variables, thereby explicitly capturing epistemic uncertainty associated with model structure and parameters. Sreenath and Raghukanth (2023) developed a Bayesian neural network-based ground-motion model using the NGA-West2 and NGA-Sub databases. Recently, Yedulla et al. (2025) developed a BNN-based GMM for Europe (horizontal

and vertical spectral accelerations) and used it to quantify epistemic uncertainty in the predictions, demonstrating that Bayesian networks can capture a wider and more realistic uncertainty band compared to deterministic ANN models. Similarly, GANs and conditional VAEs have been explored to ensure that the spread of predictions reflects true epistemic variability (e.g., Sriwastav et al., 2025a,2025b). Zhao et al. (2026) developed a BNN-based framework for predicting site nonlinear effects, where the model captures both aleatory and epistemic uncertainties and provides probabilistic estimates of site response using strong-motion datasets.

Recent studies have also demonstrated the suitability of Bayesian neural networks for data-sparse problems. For instance, Moradi et al. (2026) developed a BNN framework for cone penetration test prediction at sparse sites, showing improved predictive stability and effective uncertainty quantification under limited data conditions. Despite these developments, applications of BNNs to ground-motion modelling remain limited, particularly for stable continental regions. Previous BNN-GMMs have primarily considered the NGA-West2, NGA-Subduction, and European databases, while existing GMMs for Peninsular India provide median predictions without directly quantifying posterior epistemic uncertainty or examining its propagation into PSHA. Consequently, the period- and predictor-dependent epistemic uncertainty associated with ground-motion prediction in PI remains insufficiently characterized.

This study develops a BNN-GMM for Peninsular India using the ground-motion database compiled by Meenakshi et al. (2023). The model predicts PGA, PGV, and 5%-damped pseudo-spectral acceleration over periods of 0.01–3.0 s as functions of moment magnitude, Joyner–Boore distance, V_{S30}, and focal depth. Variational Bayesian inference is employed to represent the network weights and biases as posterior probability distributions, enabling probabilistic ground-motion predictions rather than deterministic point estimates. Epistemic uncertainty is quantified through Monte Carlo sampling of the posterior network parameters and examined as a function of oscillator period and different combinations of magnitude, distance, site condition, and focal depth. Independently, logarithmic residuals are decomposed using a mixed-effects framework to estimate the between-event, site-to-site, and remaining single-station components of aleatory variability. Model performance and physical consistency are evaluated using statistical metrics, residual trends, and parametric analyses of magnitude scaling, distance attenuation, site response, depth dependence, and spectral shape. Predictions are further compared with established GMMs and validated against recordings from significant Peninsular Indian earthquakes, including the Bhuj and Koyna events. Finally, a PSHA sensitivity analysis is conducted to quantify the influence of epistemic and aleatory uncertainty on hazard curves and their associated fractiles, thereby assessing the contribution of BNN-derived epistemic uncertainty to seismic-hazard estimates for Peninsular India.

2. Database

In the present study, the ground-motion database compiled by Meenakshi et al. (2023), previously utilized for developing an ANN-based GMM for the Peninsular India (PI) region, is adopted. A brief summary of the composite database is presented here.

The primary data sources include the PEER NGA-East database (Goulet et al., 2021a) and the dataset compiled by Gupta et al. (2017), which contains earthquake records from January 2001 to December 2016. Moreover, additional records from the global NESS database are incorporated to augment the large-magnitude events at near-source distances. Given the limited availability of strong-motion data in PI, region-specific recordings from the Bhuj, Jabalpur, and Koyna earthquakes are also included. These records were obtained from the Indian Meteorological Department (IMD) and processed in the same manner as the rest of dataset to main consistency.

Subsequently, a set of selection criteria is applied to obtain the final database:

i) records with missing metadata were removed;
ii) records with Joyner–Boore distance (R_{jb}) greater than 1500 km were excluded;
iii) records corresponding to events in the Gulf Coast region were omitted because of their distinct attenuation characteristics relative to the NGA-East region;
iv) records with unknown or questionable metadata, including placeholder values such as 999, were removed; and

After applying the selection criteria, the final database comprises 39,835 records from 3,010 events recorded at 3,000 stations. The distribution of data with respect to magnitude and distance, along with their marginal distributions, is presented in **Figure** 1. The dataset is strongly dominated by small-to-moderate magnitude events, with approximately 85% of the records corresponding to M_w 3.0–4.0, while large-magnitude events ($M_w > 6.0$) are sparsely represented. In terms of source-to-site distance, the recordings are primarily concentrated in the 30–300 km range, as evident from both the scatter density and marginal histogram. Notably, the representation of large-magnitude events at larger distances remains limited, reflecting the inherent data sparsity in such regions.

The distribution of input parameters is summarized in **Figure** 2. It should be noted that V_s30 values across the combined database are derived from a mix of direct measurements and proxy-based estimates: NGA-East records are predominantly proxy-based (~88%, largely from topographic-slope and geology correlations; Goulet et al., 2021a), Vs30 in the Gupta et al. (2017) CEUS dataset is entirely proxy-based (topographic slope; Wald and Allen, 2007), and PI/Himalayan records similarly rely predominantly on indirect site classification rather than measured shear-wave velocity profiles. The dataset is predominantly characterized by site conditions with V_s30 values between 180–760 m/s, with the largest contribution from the 360–760 m/s range (~50%), followed by 180–360 m/s (~45%). Stiffer site conditions ($Vs_{30} > 760$ m/s) are minimally represented. The depth distribution indicates that most events occur within the 5–10 km range (~65%), while deeper events (>20 km) are scarce.

In addition, the statistical characteristics of the intensity measures (PGA, PGV, and 5% damped PSA), including mean, median, standard deviation, skewness, and kurtosis, are summarized in **Table 1**. The results indicate that all intensity measures exhibit positive skewness, reflecting asymmetric distributions with extended tails toward higher amplitudes. Furthermore, the kurtosis values exceed that of a Gaussian distribution, particularly at longer periods, indicating heavy-tailed

behavior. These observations confirm that the intensity measures do not strictly follow a normal distribution.

Overall, the dataset exhibits non-uniform distribution across magnitude, distance, and site conditions, along with non-Gaussian statistical characteristics of the intensity measures. Under such conditions, conventional parametric GMPE formulations may be insufficient to capture the underlying complexity of the data. Therefore, data-driven ground motion models are better suited and require to effectively learn the nonlinear relationships between predictor variables and target responses, while providing reliable predictions within the interpolation range of the dataset.

3. Bayesian Neural Network Formulation for Ground-Motion Modeling

The limited availability of strong-motion recordings in Peninsular India (PI), together with incomplete sampling of magnitude, distance, depth, and site conditions, introduces substantial uncertainty into data-driven ground-motion models. Conventional artificial neural networks (ANNs) represent their weights and biases by single optimized values and consequently provide point estimates unless a separate uncertainty model is introduced. Bayesian neural networks (BNNs) address this limitation by assigning probability distributions to selected network parameters, thereby producing posterior predictive distributions rather than single-valued ground-motion estimates.

BNNs can be implemented at different levels of probabilistic depth. In a fully Bayesian network, the weights and biases of all hidden and output layers are represented probabilistically. Although this formulation allows parameter uncertainty to propagate throughout the network, it requires estimation of a large number of posterior parameters. In a partially Bayesian network, probability distributions are assigned only to selected layers, while the remaining layers are deterministic. Another commonly used formulation assigns probability distributions only to the final layer, thereby reducing computational cost but restricting uncertainty representation in the internal features learned by the network. Monte Carlo dropout and ensembles of independently initialized deterministic networks can also approximate predictive uncertainty, although they are methodologically distinct from explicit posterior inference over Bayesian network parameters.

Previous BNN-based GMMs provide important precedents for probabilistic ground-motion prediction. Sreenath and Raghukanth (2023) developed BNN-GMMs for the NGA-West2 and NGA-Sub databases using mean-field variational inference. Their models employed probabilistic hidden and output layers to predict peak ground-motion parameters and response spectra as functions of magnitude, rupture distance, focal depth, V_{S30}, style of faulting, and regional indicators. Yedulla et al. (2025) subsequently developed BNN-GMMs for horizontal response spectra and vertical-to-horizontal spectral ratios using the European ESM2.0 database. Their BNN posterior was approximated through variational inference, and epistemic uncertainty was evaluated by repeatedly sampling the trained network parameters. For each predictor combination, 100 posterior samples were used to produce a distribution of ground-motion predictions. These predictions represent Monte Carlo samples from a single trained BNN.

The present study follows similar methodology and develops a partially Bayesian GMM specifically for PI. The selected formulation assigns probability distributions to the hidden-layer

parameters, allowing uncertainty to propagate through the internal ground-motion representation, while retaining a deterministic output mapping that predicts a conditional mean and variance for each intensity measure. Epistemic uncertainty is estimated from the variability of conditional mean predictions generated by posterior parameter sampling. The selected layer configuration and output structure are described in the following section.

Let $\boldsymbol{w}$ denote the probabilistic network parameters and $\mathcal{D}$ the training dataset containing predictor variables $\boldsymbol{x}$ and logarithmic ground-motion observations $\boldsymbol{y}$. The deterministic output-layer parameters are jointly optimized via gradient descent and are treated as fixed point estimates rather than random variables; for notational clarity, $\boldsymbol{w}$ below refers only to the probabilistic hidden-layer parameters over which the posterior is defined, with the deterministic parameters held implicit in $p(\mathcal{D} \mid \boldsymbol{w})$. The posterior distribution of the network parameters follows Bayes' theorem:

$$p(\boldsymbol{w} \mid \mathcal{D}) = \frac{p(\mathcal{D} \mid \boldsymbol{w})\, p(\boldsymbol{w})}{p(\mathcal{D})}$$

where $p(\boldsymbol{w})$is the prior distribution and $p(\mathcal{D} \mid \boldsymbol{w})$is the likelihood of the observations for a given set of network parameters. A zero-mean Gaussian prior, $p(w_i) = \mathcal{N}\left(0, \sigma_p^2\right)$, is assigned to each Bayesian parameter. The prior standard deviation $\sigma_p$controls the allowable spread of the parameters and is treated as a hyperparameter in the present implementation rather than being fixed at unity. A Gaussian likelihood is adopted in logarithmic ground-motion space, consistent with the conventional representation of ground-motion residuals.

Because the exact posterior cannot be evaluated analytically for a nonlinear neural network, it is approximated using variational inference. The approximate posterior is represented by a mean-field Gaussian distribution,

$$q(\boldsymbol{w} \mid \boldsymbol{\theta}) = \prod_i \mathcal{N}\left(w_i, \mid \mu_i \sigma_i^2\right)$$

where $\boldsymbol{\theta} = \{\mu_i, \sigma_i\}$contains the learnable posterior means and standard deviations. The mean-field assumption treats the Bayesian parameters as statistically independent, substantially reducing the number of posterior quantities that must be estimated compared with a full-covariance representation. During training, parameter samples are generated using the reparameterization procedure, and a softplus transformation is applied to ensure positive posterior standard deviations.

The variational parameters are learned by minimizing the negative evidence lower bound:

$$\mathcal{J}(\boldsymbol{\theta}) = \mathbb{E}_{q(\boldsymbol{w}\mid\boldsymbol{\theta})}[-\log p\left(\mathcal{D} \mid \boldsymbol{w}\right)] + \beta\, D_{\mathrm{KL}}[q(\boldsymbol{w} \mid \boldsymbol{\theta}) \parallel p(\boldsymbol{w})],\ \beta = \frac{1}{N_{\mathrm{train}}}$$

The first term measures agreement between the observations and probabilistic network predictions, whereas the Kullback–Leibler divergence constrains the approximate posterior relative to the prior. Because the model predicts both a conditional mean and a log variance for each intensity measure, the likelihood term is evaluated using a Gaussian negative log-likelihood that accounts for prediction error and the variance assigned to each observation. Therefore, the implemented objective is not equivalent to a conventional mean-squared-error loss. Scaling the Kullback–

Leibler term by $1/N_{\text{train}}$maintains an appropriate balance between data fitting and posterior regularization during mini-batch optimization.

Following training, predictions are obtained through Monte Carlo sampling of the approximate weight posterior. Each sampled parameter set produces a conditional mean and variance for the specified predictor combination. The average of the sampled conditional means provides the final median ground-motion estimate, whereas their variability quantifies parameter-related epistemic uncertainty. Details of the selected architecture, and probabilistic-layer parameterization are provided in the following section.

3.1. Architecture, Hyperparameter Optimization, and Training of the BNN

The predictive performance of a Bayesian neural network is sensitive to both its architectural configuration and its training settings. To identify a suitable configuration, hyperparameter optimization was performed using the Optuna framework, which applies a Bayesian optimization strategy based on the Tree-structured Parzen Estimator (TPE) to efficiently balance exploration and exploitation of the search space. The search space comprised the hidden-layer width (searched between 4 and 16 neurons, applied identically to both Bayesian hidden layers), learning rate, batch size, prior standard deviation, and the minimum log-variance bound; the number of hidden layers was fixed structurally at two, and the number of training epochs was fixed at 150 during the search to ensure a consistent computational budget across trials. Each of the 50 completed trials was evaluated using validation-set RMSE, which served as the objective function minimized by Optuna, with a median pruning strategy used to terminate unpromising trials early.

The relative contribution of each hyperparameter to validation RMSE variance across trials, computed via functional ANOVA, is shown in Supplementary Figure S1. Learning rate was by far the most influential hyperparameter (importance = 0.606), followed by hidden-layer width (0.204) and prior standard deviation (0.109); batch size (0.046) and the minimum log-variance bound (0.034) contributed comparatively little to validation performance, indicating that predictive accuracy was primarily governed by the optimization step size and network capacity, while the prior and variance-clamping settings played a secondary, stabilizing role. The evolution of validation RMSE across all 50 trials, together with the best-so-far RMSE curve, is shown in Supplementary Figure S2; although several trials completed after Trial 1 achieved marginally lower validation RMSE, Trial 1 (RMSE = 0.8903) was adopted as the final configuration because its median predictions more consistently reproduced physically expected earthquake scaling behavior, including magnitude saturation, distance attenuation, and site amplification, while providing comparable predictive accuracy to the lowest-RMSE trials. The position of Trial 1 relative to all completed trials, across each of the five searched hyperparameters, is shown in Figure 3. The selected configuration comprises 9 hidden units, a learning rate of approximately 1.13×10^{-3}, a batch size of 16, a prior standard deviation of approximately 1.72, and a minimum log-variance bound of −10.

Using this configuration, the network was trained by optimizing the variational parameters (posterior means and standard deviations of the Bayesian layer weights) with the Adam optimizer, selected for its adaptive learning-rate behavior and reduced sensitivity to initialization. The training objective is the evidence lower bound (ELBO), comprising a heteroscedastic Gaussian

negative log-likelihood and a Kullback–Leibler (KL) divergence term, rather than a mean-squared-error loss. Each Bayesian layer is parameterized using learnable mean and rho terms, with a softplus transformation applied to rho to ensure a strictly positive standard deviation, and a sigmoid activation applied following each Bayesian hidden layer. During training, stochastic sampling of the Bayesian layer weights is performed at each forward pass using the reparameterization trick, allowing parameter uncertainty to propagate through the network, and the KL-divergence term is scaled by $\beta = 1/N$ to maintain an appropriate balance between data fitting and posterior regularization during mini-batch optimization.

The compiled dataset was divided into training, validation, and testing subsets using a 70/10/20 split; input features were normalized using min–max scaling, while target ground-motion intensity measures, in logarithmic space, were standardized to zero mean and unit variance, with both transformations fit on the training set only. Following training, predictive distributions are obtained through Monte Carlo sampling of the trained weight posterior (100 samples per prediction), rather than from an ensemble of independently trained networks; the resulting sample-to-sample variability in the conditional mean provides the epistemic uncertainty estimate discussed in later sections.

The resulting final architecture comprises five input features, two Bayesian hidden layers with nine neurons each, and a deterministic output layer consisting of 44 nodes (Figure 4). Both hidden layers are modeled within a Bayesian framework to capture parameter uncertainty, while the output layer remains deterministic and maps the final hidden representation to 44 outputs, corresponding to the mean and log-variance of the 22 ground-motion intensity measures (refer to the database section for details); this heteroscedastic output structure allows the network to learn a period-dependent aleatory variance directly from the data.

For the first Bayesian hidden layer (5 inputs, 9 neurons), the number of network parameters includes 45 weights and 9 biases, totaling 54 parameters. The second Bayesian hidden layer (9 inputs, 9 neurons) comprises 81 weights and 9 biases, totaling 90 parameters. A full-covariance Gaussian posterior for these two layers would require estimation of covariance matrices of size 54×54 and 90×90, containing 1,485 and 4,095 unique elements respectively, in addition to the 54 and 90 mean parameters, yielding a combined total of 5,724 unknowns. The deterministic output layer (9 inputs, 44 neurons) comprises 396 weights and 44 biases, totaling 440 fixed parameters estimated via standard backpropagation, with no posterior distribution defined over them. To address the computational burden of full-covariance estimation in the Bayesian layers, a mean-field variational approximation is adopted, wherein all Bayesian parameters are assumed to be statistically independent; under this assumption, each parameter is characterized by its own mean and standard deviation, reducing the number of variational unknowns to twice the number of Bayesian network parameters, i.e., 288 for the present architecture, compared with 5,724 under a full-covariance treatment.

3.2.Performance of BNN

The performance of the developed model is evaluated using standard statistical metrics, including the coefficient of determination (R^2) and the correlation coefficient (R), across the training,

validation, and testing datasets. The detailed performance metrics are summarized in Table 2 for all intensity measures, namely PGA, PGV, and PSA from 0.01 s to 3 s.

Overall, the model demonstrates good predictive capability across all datasets, with consistent performance between the training, validation, and testing sets. Across the individual intensity measures, R^2 ranges from 0.775 to 0.870 and the correlation coefficient (R) ranges from 0.881 to 0.933 (Table 2), with all correlations statistically significant ($p < 0.001$). Relative to shorter periods, PSA at periods greater than 0.3 s generally shows higher R^2 and lower MSE, rising from $R^2 \approx 0.79$–0.81 in the 0.15–0.30 s range to $R^2 = 0.870$ at 3.0 s, which can be attributed to the relatively smoother behavior and reduced variability at longer periods. The lowest R^2 (0.775) and highest MSE occur near 0.075–0.10 s, corresponding to the period range most sensitive to local site resonance and near-surface amplification effects.

The target versus predicted values are shown in Figure 5 for PGA, PGV, and PSA at 0.1 s and 1 s, with R^2 values of 0.797, 0.835, 0.775, and 0.850, respectively. The results indicate that the data points are closely clustered around the 1:1 reference line, demonstrating good agreement between predictions and observations. This trend is consistent across the training, validation, and testing datasets, further confirming the robustness of the model. Some dispersion is observed, particularly at lower amplitude levels; however, no significant systematic bias is evident, as the predictions are symmetrically distributed around the ideal line.

In addition, the mean squared error (MSE) values for the training, validation, and testing datasets are presented in Table S1 for all output intensity measures. The MSE values are comparable across the three datasets, indicating consistent model performance; for the selected configuration (Figure 4), the pooled training, validation, and test MSE are 0.823, 0.793, and 0.824, respectively, with corresponding R^2 of 0.813, 0.815, and 0.817. The absence of significant differences between training and testing errors suggests that the model does not suffer from overfitting and generalizes well to unseen data. This consistency further supports the robustness and stability of the developed BNN framework. Overall, the results from Table 2, Table S1, and Figure 5 demonstrate that the BNN model provides reliable and consistent predictions of ground-motion intensity measures.

4. Uncertainty Quantification

Predictive uncertainty is characterized using two complementary procedures. Within-model epistemic uncertainty is quantified through Monte Carlo sampling of the variational posterior distributions of the Bayesian network parameters. Aleatory variability is estimated independently by decomposing the logarithmic ground-motion residuals using a mixed-effects model. The former represents uncertainty in the conditional median prediction, whereas the latter represents the remaining variability of individual recordings about that median. These components are subsequently combined to evaluate total predictive uncertainty and are propagated separately in the PSHA calculations.

The predictive model is expressed in logarithmic space as:

$$ln(IM) = f_{BNN}(M_W, R_{jb}, V_s30, depth) + \epsilon_{alea}$$

where $f_{\mathrm{BNN}}(\cdot)$represents the median prediction from the ensemble BNN model, and $\varepsilon_{\mathrm{alea}}$denotes the aleatory residual.

$$\sigma_{epi} = std(f_{BNN}^{(k)}(\cdot))$$

$$\sigma_{alea} = std(y_{obs} - y_{pred})$$

Here, $f_{\mathrm{BNN}}^{(k)}$represents the predictions from multiple Bayesian neural network realizations (where $k$denotes different weight samples or ensemble members), whereas y_{obs}corresponds to the observed ground-motion values and y_{pred} denotes the predicted ground-motion values.

4.1. Aleatory variability

Aleatory variability is quantified using the residuals between the predicted and observed ground-motion intensity measures. Following standard practice in ground-motion modeling, these residuals are decomposed into inter-event, site-to-site, and within-event components using a likelihood-based mixed-effects approach. The estimation is carried out using restricted maximum likelihood (REML), iterating until the likelihood converges.

The total residual is expressed as:

$$\mathrm{r_{ij}} = \eta_{\mathrm{i}} + \delta_{\mathrm{j}} + \epsilon_{\mathrm{ij}}$$

where η_i represents the inter-event variability for the i-th earthquake, δ_j denotes the site-to-site variability for the j-th station, and ε_{ij} corresponds to the within-event (record-to-record) variability. The decomposition is implemented using a linear mixed-effects model, where earthquake and site identifiers are treated as random effects, fitted separately for each spectral ordinate using REML, from which the inter-event, site-to-site, and within-event components are obtained (Khosravikia and Clayton, 2021).

The aleatory standard deviations are provided in Table 3 (ln units). The total sigma (σ_{total}) ranges from approximately 0.76 to 0.89, the inter-event term (τ) ranges from approximately 0.42 to 0.63, the site-to-site term (φ_{S2S}) ranges from approximately 0.43 to 0.64, and the single-station term (φ_{SS}) is comparatively narrow, ranging from approximately 0.40 to 0.42 across all periods. Notably, φ_{S2S} exceeds φ_{SS} at every period.

The finding that φ_{S2S} exceeds φ_{SS} at every period (Table 3) departs from the pattern typically reported in ground motion studies, where the single-station term is usually the larger of the two components. The possible contributing factor is the site characterization underlying the combined database. V_s30 in the NGA-East component is predominantly proxy-based, with only approximately 12% of contributing stations having measurement-based Vs30, the remainder relying on topographic-slope and geology-based correlations (Goulet et al., 2021a). V_s30 in the Gupta et al. (2017) CEUS dataset is entirely proxy-based, derived from topographic slope (Wald and Allen, 2007), with no directly measured values at any recording site. Similarly, V_s30 for the

Peninsular India records relies predominantly on indirect site classification rather than measured shear-wave velocity profiles. The larger φ_{S2S} observed in the present study is therefore interpreted as physically grounded limitation of V_s30 as a sole site descriptor in this tectonic setting, consistent with prior findings in comparably characterized stable continental and hard-rock regions (Al Atik, 2015), rather than as an artifact of proxy-based site characterization or limited station repeat-occupancy in the combined database.

Further, to assess the model performance, residuals are plotted against magnitude (M_w), rupture distance (R_{jb}), and site condition (V_s30), as shown in Figure 6 for PGA and PGV, and in Figure 7 for PSA (0.1 s) and PSA (1 s). The residuals are well distributed around zero, with no visible systematic trend with respect to any of the input variables. A similar behavior is observed with respect to distance, where the residuals show no consistent increase or decrease from near-field to far-field ranges. Across site conditions, the residuals remain centered around zero over the full range of V_s30, suggesting that the site amplification effects are reasonably well represented.

The scatter of residuals is broadly uniform, without pronounced systematic trends. Some dispersion is observed at sparsely sampled ranges (e.g., very large magnitudes or near-field distances), attributable to limited data availability at these extremes. The binned mean residuals (red markers) remain close to zero across the full range of each input variable, further indicating the absence of systematic bias. Overall, the residuals exhibit a broadly random and symmetric scatter around zero, supporting that the model predictions are unbiased across the range of input parameters considered.

4.2. Epistemic uncertainty

In this study, epistemic uncertainty is quantified using Monte Carlo (MC) sampling from the variational posterior distribution of the trained Bayesian neural network's weights, rather than from an ensemble of independently trained networks. For a given input scenario, the trained model is evaluated S times (e.g., S = 100 forward passes), with a new set of weights sampled from the posterior at each pass via the reparameterization trick. This results in a distribution of predictions reflecting the uncertainty associated with the learned network parameters.

The mean of these S sampled predictions provides the representative median ground-motion estimate for the given scenario:

$$\mu(x) = \frac{1}{S}\sum_{s=1}^{S} y_s\,(x)$$

Epistemic uncertainty is then estimated from the variability across these S sampled predictions, i.e., their standard deviation about $\mu(x)$. A larger spread indicates higher epistemic uncertainty, reflecting limited data support or reduced model confidence in that region of the input space, while a smaller spread indicates more consistent predictions across posterior samples.

To further illustrate epistemic uncertainty in comparison with the observed data recordings from four earthquakes that span a range of magnitudes, source-to-site distances, and site conditions, is considered and it is shown in Figure 8. For each event, the individual Monte Carlo posterior samples are shown in gray, the corresponding mean prediction in black, and the recorded response spectrum in red. In contrast to a deterministic model, which produces only a single point estimate, the distribution of posterior samples provides a direct representation of the model's epistemic uncertainty at each oscillator period. These samples were generated by repeatedly drawing from the posterior distribution of the trained model weights. Across the four events, the mean predictions capture the overall shape of the recorded response spectra, while the observations generally fall within or near the posterior sample band over the entire range of periods.

Further, the total predictive uncertainty is evaluated and shown together with its epistemic and aleatory components in Figure 9 as a function of period. The aleatory component (σ_{total}, panel a) ranges from approximately 0.82 to 0.89, following a two-peak pattern with elevated values near 0.04–0.10 s and again near 0.75–1.5 s, separated by a relative minimum near 0.15–0.40 s, consistent with Table 3. The epistemic component (Figure 9, panel b) is substantially smaller, ranging from approximately 0.04 to 0.10 across the evaluated scenarios. For the baseline, stiff-site, and deep-event scenarios, epistemic uncertainty increases fairly steadily with period, reaching its highest values (≈0.09–0.10) beyond 1 s. The long-distance scenario (M_w = 6.0, R_{jb} = 800 km) shows a marked increase in epistemic uncertainty relative to the baseline over the 0.05–0.3 s range, peaking near 0.15 s (≈0.082), consistent with the sparse representation of far-field recordings (R_{jb} > 500 km) in the database, which is dominated by records within approximately 250 km (Figure 1). Beyond this period range, epistemic uncertainty for the long-distance scenario declines to the lowest values among all scenarios evaluated (≈0.04–0.045), suggesting the elevated uncertainty is localized to short-to-intermediate periods rather than persisting uniformly across the spectrum. This behavior is broadly consistent with the general finding that Bayesian neural network ground-motion models return larger epistemic deviations in comparatively underrepresented regions of the predictor space, with no corresponding effect on the aleatory component, as similarly reported by Sreenath and Raghukanth (2023).

As a result, total predictive uncertainty (panel c) is dominated by the aleatory component and closely follows its period-dependent shape, ranging from approximately 0.82 at short periods to approximately 0.89 near 1.0 s, with the long-distance and low-magnitude scenarios producing marginally lower total uncertainty than the baseline at longer periods due to their comparatively lower epistemic contribution at those periods.

4.3.Influence of Predictor Variables on Ground Motion

The scaling behavior of the developed BNN-based ground-motion model is evaluated with respect to key predictor variables, namely magnitude (M_w), Joyner–Boore distance (R_{jb}), shear-wave velocity (V_{s30}), and focal depth. The model is assessed for its ability to reproduce observed ground-motion trends while maintaining consistency with the fundamental physical principles governing earthquake behavior.

The variation of ground-motion intensity measures (PGA, PGV, and PSA at 0.1 s and 1 s) with respect to distance is shown in Figure 10 for different magnitudes. The attenuation of ground

motion with increasing distance is observed across all intensity measures. It is also clearly noticed that the decay is steeper at short distances and gradually reduces at larger distances, capturing the effects of geometric spreading and anelastic attenuation. Ground-motion saturation at larger magnitudes is observed for PGA and PSA at short periods (e.g., 0.1 s), whereas it is less evident for PGV and PSA at longer periods (e.g., 1 s). This behavior can be attributed to differences in the frequency content and scaling characteristics of ground motion, as PGA and PSA at short periods are dominated by high-frequency components, which tend to saturate at large magnitudes, particularly at short distances. Consequently, ground-motion amplitudes become relatively similar for higher magnitudes.

In addition, event-specific comparisons between observed and predicted ground motions are presented in Figure 11 for selected earthquakes. The results show that the predicted median trends follow the observed data across all intensity measures and distance ranges, showing that the model is also able to generalize at an event level. The shaded regions represent the total predictive uncertainty (95% CI) derived from the Bayesian framework, and largely cover the observed data, indicating that the model effectively captures the variability in ground-motion behavior. The uncertainty bands also appear to widen in regions where observational data are comparatively sparse, consistent with reduced data support in this range (Figure 1).

Further, the scaling of ground motion with magnitude for different source-to-site distances is shown in Figure 12. Intensity measures increase with magnitude across all distances, reflecting the expected increase with magnitude. The rate of increase in amplitude is not linear, depicting stronger amplification of ground motion at shorter distances, whereas at larger distances, the influence of magnitude decreases as attenuation effects dominate. These trends are consistent with established GMMs, indicating that the model is able to capture ground-motion scaling with magnitude and distance.

Furthermore, the influence of site conditions is examined by varying V_s30, as shown in Figure 13. Amplification of ground motion is observed at lower V_s30 sites relative to rock/stiff strata. It is also noted that the peak amplitude occurs near 0.25 s for rock strata, whereas for softer sites the peak shifts toward 0.3–0.35 s, capturing the expected nonlinear site-response behavior. However, since this amplification is slight, an additional comparison was made between predicted ground motions and observed data grouped by V_s30 range. Figure 14 and Figure S3 present the variation of predicted ground motions with Joyner–Boore distance (R_{jb}) for magnitude bins centered at M_w = 5.0 and M_w = 7.5, respectively. The predictions correspond to a reference condition of V_s30= 760 m/s and a fixed focal depth of 5 km, while the observations are grouped according to two V_s30 ranges. In the present dataset, the distinction between the two V_s30 groups (360–760 m/s and 760–1500 m/s) is not strongly pronounced. At larger distances, particularly for higher-magnitude events (M_w = 7.5, Figure S3), some deviation between predictions and observations is observed. Those observations correspond to recordings from the Bhuj earthquake. For these data points, precise V_s30 measurements are not available, and a representative value of V_s30 = 1500 m/s has been assumed to characterize rock conditions; this approximation may contribute to the observed discrepancies in this distance range. From these predictions, the model is able to capture the general attenuation behavior with distance across the intensity measures considered, with predicted median trends generally lying within the spread of the observed data.

The effect of focal depth on ground-motion prediction is shown in Figure 15. Across the depth and distance combinations evaluated, the predicted amplitudes show little sensitivity to focal depth, with predictions for shallower and deeper events remaining closely comparable at similar magnitude and distance.

4.4.Comparison with Existing Global Ground Motion Models

Several ground motion models (GMMs) have been developed for stable continental regions (SCR), particularly for the Central and Eastern North America (CENA) region, using a range of methodologies including stochastic simulations, hybrid empirical methods (HEM), and empirical regressions. In order to evaluate the performance of the proposed model, a comparison is carried out with representative and widely used GMMs, namely Boore et al. (2018) – AB10 and BCA10d, Pezeshk et al. (2018) – PZCT-M1ES and PZCT-M2ES, Graizer (2016) – G16, and the NGA-East model (Goulet et al., 2021b). The model proposed in this study is referred to as the "Proposed GMM."

Figure 16 presents the comparison of predicted ground-motion amplitudes (PGA, PSA at 0.1 s, and PSA at 1 s) as a function of magnitude for different source-to-site distances (5, 30, and 100 km). Overall, the predictions from the proposed GMM follow trends similar to those of the existing models, showing a consistent increase in amplitude with magnitude. At shorter distances (R = 5 km), the proposed model predicts noticeably lower amplitudes than several of the existing models, particularly AB10 and NGA-East, across the full magnitude range considered. As distance increases, the predictions from the proposed GMM align closely with the existing models. Across all panels, the proposed GMM is flows the same trend as with the M23 model, which is developed using a comparable combined database including region-specific recordings; this agreement suggests that the differences relative to the other comparison models reflect, at least in part, differences in the underlying calibration datasets rather than the specific choice of modeling framework.

Figure 17 shows the comparison of response spectra across periods for different magnitudes (M_w= 5.5, 6.5, and 7.5) and distances (R = 5, 30, and 100 km). The spectral shapes obtained from the proposed model are smooth and consistent with those from other models, exhibiting a clear peak at short-to-intermediate periods followed by a gradual decay. At R = 5 km, the proposed model predicts systematically lower amplitudes than several existing models (e.g., AB10, PZCT) across the magnitude range evaluated, most noticeably at short-to-intermediate periods; this gap narrows at longer periods and at larger distances, where predictions from all models, including the proposed GMM, show closer agreement. Notably, the proposed GMM and the M23 model remain closely aligned across nearly all magnitude–distance combinations evaluated, including at M_w = 7.5, consistent with their shared underlying database. This suggests that the amplitude differences observed among the compared models, particularly at short distances, are driven at least in part by differences in near-field, large-magnitude data coverage across the underlying databases rather than by the choice of modeling framework alone.

In addition, the comparison of total standard deviation (σ) with contributions from epistemic and aleatory components, along with values from selected existing studies are shown in Figure 18. It can be observed that the epistemic uncertainty estimated in the present study remains relatively

small (approximately 0.1–0.15) across all periods, indicating stable model predictions. In contrast, the aleatory variability dominates, with values around 0.7–0.8, leading to a total standard deviation close to ~0.9.

When compared with other models, the total uncertainty from the proposed GMM is found to be consistent with published values, lying within the overall range reported in previous studies. While some models show slightly higher or lower variability at certain periods, the general trend remains comparable. The relatively small epistemic component further suggests that the BNN framework provides stable predictions with reduced model uncertainty. Overall, the comparisons indicate that the proposed GMM produces predictions that are consistent with established models in terms of magnitude scaling, distance attenuation, spectral shape, and associated uncertainty. The slight differences observed in certain ranges are within acceptable limits and comparable to the variability among existing models. These results demonstrate that the proposed methodology is capable of providing reliable ground-motion estimates and can be effectively applied to data-sparse regions with similar tectonic characteristics.

4.5.Comparison of Model Predictions with Peninsular India Ground-Motion Data

A comparison between recorded PI earthquake data and the corresponding predictions from the developed BNN-GMM is presented in this section. The PI region is characterized by a low-to-moderate seismicity rate, and it was the occurrence of the 1967 M_w 6.3 Koyna earthquake that first drew seismologists' attention to the region (Chandra, 1977; Guha and Basu, 1993). Only 12 earthquakes with $M_w > 6$ have occurred in PI since the 1900s, among which the 1993 Killari (Latur) earthquake (M_w 6.1) and the 2001 Bhuj earthquake (M_w 7.7) are among the most significant and damaging. For the Bhuj and Jabalpur earthquakes, only far-field recordings are available, while a limited set of near-field recordings exists for the Koyna region. This scarcity of well-distributed recordings across the magnitude–distance range needed for direct GMPE development is a central motivation for the data-driven, uncertainty-aware framework adopted in this study, and is also the reason regional GMPEs for PI have historically relied heavily on simulated rather than recorded ground motion (Raghukanth and Iyengar, 2007; Boore, 1983; Toro et al., 1997). Given this data scarcity, the developed BNN-GMM is validated against the two PI earthquakes with the most complete available recordings for the 2001 Bhuj earthquake (Cramer and Kumar, 2003) and the 1967 Koyna earthquake to assess whether the model captures observed ground-motion characteristics in the PI region.

For the Bhuj earthquake, recorded PGA and PSA at 0.4 s, 0.75 s, and 1.0 s are compared against the model's predicted attenuation with distance across a representative range of site conditions (V_s30 = 180, 360, and 760 m/s), with the 95% confidence interval, and the results are shown in Figure 19 and Figure S4. The model reproduces the expected attenuation of ground motion with distance, and the majority of observed ground motion data fall within the predicted 95% CI despite the natural scatter typical of a single recorded event. The model is further evaluated against the recorded response spectrum of the 1967 M_w 6.5 Koyna earthquake, compared with both the observed spectrum and the predictions of the regional ANN-based GMM of Meenakshi et al. (2023) (Figure 20). The developed model reproduces the overall spectral shape, including the rise, peak, and decay with period, and the recorded spectrum falls within the model's total predictive uncertainty band throughout.

Moreover, since the developed model is purely data-driven and does not require seismological source or path parameters, unlike the stochastic simulation-based models, generally relied upon for this region, it can be suggested that the developed model provides reasonable estimates for predicting ground motion in the PI region. Furthermore, since the model explicitly separates and quantifies epistemic and aleatory uncertainty, an additional effort has been done in the following section to study the influence of these uncertainty components on the resulting seismic hazard estimates.

4.6.Seismic Hazard analysis

To demonstrate the applicability of the proposed ground-motion model, a Probabilistic Seismic Hazard Analysis (PSHA) is performed for Peninsular India (PI). The analysis follows the classical PSHA framework, wherein seismic source characterization, earthquake recurrence, and ground-motion prediction are integrated using the total probability theorem. The mean annual rate of exceedance of a ground-motion level $y^*$is expressed as:

$$\mu_{y^*} = \sum_{i=1}^{n_{\text{sources}}} \lambda_i \left(M_{\min}\right) \int_{M_{\min}}^{M_{\max}} \int_{0}^{r_{\max}} P(Y > y^* \mid m, r)\, f_R(r)\, f_M(m)\, dr\, dm$$

where μ_{y^*}denotes the annual exceedance rate, $\lambda_i$represents the occurrence rate of earthquakes exceeding the minimum magnitude, and $f_M$and $f_R$denote the probability density functions of magnitude and distance, respectively. The probability of exceedance over an exposure period $T$is given by:

$$P(Y > y^*; T) = 1 - \exp(-\mu_{y^*} T)$$

In India, due to the limited availability of fault slip rate data, a distributed seismicity approach is generally adopted for hazard analysis. Accordingly, the PSHA framework developed by Sreejaya et al. (2022) is followed in this study, and only the essential steps are summarized here.

Peninsular India is divided into 33 seismogenic source zones by NDMA (2010) based on past seismicity, seismotectonic characteristics, and regional geology. Among these, PI comprises seven zones, including three zones in South India (SZ–17, 20, and 29), three zones in Central India (SZ–16, 18, and 19), and one zone corresponding to the Kutch region (SZ–27). Seismicity characterization is performed using the updated earthquake catalogue compiled by Sreejaya et al. (2022), which includes additional data from multiple seismic networks. The catalogue is declustered using the Gardner–Knopoff algorithm to remove dependent events, and completeness is evaluated using the Stepp method.

The seismic activity of each source zone is modeled using the Gutenberg–Richter recurrence relationship:

$$\log(\lambda) = a - bM$$

where a and b are seismicity parameters describing the occurrence rate and relative distribution of magnitudes. The parameters, along with the maximum magnitude ($M_{\max}$), are estimated using maximum likelihood approaches, including the Kijko and Weichert methods. To account for spatial variability in seismicity, an elliptical smoothing approach is employed, which incorporates the distribution of past earthquakes and regional tectonic features. The study region is discretized

into grid cells, and earthquake occurrence rates (for $M \geq 4.0$) are normalized using appropriate completeness periods before applying smoothing along fault orientations. For detailed implementation of the PSHA framework and region-specific seismicity parameters for Peninsular India, the reader is referred to Sreejaya et al. (2020).

Ground-motion estimates are obtained using the proposed ground-motion model, and seismic hazard calculations are carried out using the CRISIS2003 software, which facilitates the incorporation of user-defined seismicity parameters and attenuation relationships within a probabilistic framework. Epistemic uncertainty in seismicity parameters is accounted for through a logic-tree approach by considering alternative estimates of the b-value and maximum magnitude $M_{\max}$derived from the Kijko and Weichert methods, with equal weights assigned to each branch.

To assess the influence of uncertainty in ground-motion modelling on hazard estimates, three scenarios are considered: (i) total variability, combining aleatory variability with the BNN's epistemic uncertainty propagated through continuous Monte Carlo sampling of the posterior network parameters; (ii) aleatory variability only, using the median BNN-GMM prediction together with the total aleatory standard deviation (decomposition from the residuals); and (iii) epistemic uncertainty represented through a three-branch logic tree, applied in addition to the aleatory variability retained in Scenario (ii).

For Scenario (iii), the logic tree is constructed as follows. For each site and spectral period, the controlling earthquake scenario defined by moment magnitude (M_w) and Joyner-Boore distance (R_{JB}) is identified from hazard deaggregation of the corresponding Scenario (i) hazard curve. The deaggregation results are shown in Figure S5. The controlling M_w is 5.3 with R_{JB} as 52 km. At this controlling (M_w, R_{JB}) combination, together with the appropriate V_s30=760 m/s and focal depth of 10 km, the BNN posterior is sampled to obtain the period-dependent epistemic standard deviation, $\sigma_{epistemic}(T)$, of the median ground-motion prediction. This continuous epistemic distribution is then discretized into three branches: a central branch at the median prediction, and low and high branches obtained by scaling the median by corresponding to the 5^{th} and 95^{th} percentiles of the epistemic distribution. Following the extended Pearson–Tukey three-point discrete approximation (Pearson and Tukey, 1965; Keefer and Bodily, 1983), weights of 0.63, 0.185, and 0.185 are assigned to the central, low, and high branches, respectively; this scheme reproduces the mean and variance of the underlying continuous epistemic distribution using only three discrete realizations and is widely adopted to represent epistemic uncertainty in ground-motion logic trees. Each branch is propagated independently through the full PSHA calculation in CRISIS2003, with the aleatory standard deviation applied within each branch as in Scenario (ii), and the three resulting branch-specific hazard curves are combined using the assigned weights to obtain the final hazard curve for Scenario (iii).

The hazard curves are generated for four important cities in Peninsular India, namely Mumbai, Chennai, Bangalore, and Hyderabad, where the first two cities lie in seismic Zone III as per IS 1893:2002, whereas the other two lie in seismic Zone II. The annual rate of exceedance curves for the 4 cities for three scenarios are shown in Figure 21 and 22 corresponding to PGA and PSA at 1s. The comparison demonstrates that the influence of ground-motion uncertainty depends on the site, spectral period, and exceedance level. At relatively high annual exceedance rates, the aleatory-only scenario produces comparable or, in some cases, higher hazard estimates, particularly for

Mumbai and Hyderabad. As the annual exceedance rate decreases, however, the curves generally intersect, and the scenarios incorporating epistemic uncertainty produce higher exceedance rates in the high-ground-motion range. Consequently, for a specified low annual exceedance rate, the inclusion of epistemic uncertainty generally results in a larger ground-motion estimate. This behaviour reflects the increasing contribution of the upper GMM branch to rare, high-intensity ground motions, whereas the lower branch can reduce the logic-tree-averaged exceedance rate at more frequent, low-intensity levels. The effect is site- and period-dependent: relatively small differences are observed for Chennai, whereas more pronounced differences occur for Mumbai at low-to-moderate intensities and for Bangalore and Hyderabad in the low-exceedance-rate region. The results therefore indicate that epistemic uncertainty is particularly important for hazard estimates associated with long return periods and should not be represented by a constant adjustment to the hazard curve.

In addition to the hazard curves, uniform hazard response spectra (UHRS) were calculated for return periods of 475 and 2,475 years and are presented for the four cities in Figures 23 and 24, respectively. Each spectrum represents the spectral acceleration associated with the same annual exceedance rate at every period. It is observed that the spectral ordinates for the 2,475-year return period are consistently higher than those for the 475-year return period because the longer return period represents a lower probability of exceedance and a more severe ground-motion level. For all four cities, the spectra exhibit a broad maximum approximately between 0.08 and 0.20 s, followed by a progressive reduction toward longer periods. This indicates that the hazard is dominated by short-period ground motions, which are particularly relevant to low- and medium-rise structures. The magnitude and shape of the spectra nevertheless vary among the cities, demonstrating the influence of differences in regional seismicity, source-to-site distance, and the period dependence of the adopted GMMs.

The influence of the uncertainty treatment is strongly site-dependent. For Mumbai, the aleatory-only scenario produces consistently higher spectral accelerations across almost the entire period range at both return periods. Hyderabad shows a similar pattern, particularly at short and long periods. Conversely, the three uncertainty scenarios yield closely spaced spectra for Chennai and Bangalore, indicating limited sensitivity to the uncertainty treatment. The differences are generally more pronounced at the 2,475-year return period, especially for Mumbai and Hyderabad, demonstrating that the effect of ground-motion uncertainty becomes more significant at lower annual exceedance rates. Overall, Mumbai exhibits the highest spectral accelerations among the four cities, particularly over the short- and intermediate-period ranges. Chennai generally shows the next-highest spectral demand, while Bangalore and Hyderabad exhibit comparatively lower values. The results demonstrate that the influence of ground-motion uncertainty cannot be represented by a single scaling factor because its effect varies with the site, vibration period, and return period

5. **Summary and Conclusions**

This study developed a Bayesian neural network ground-motion model (BNN-GMM) for Peninsular India to address the limited characterization of epistemic uncertainty in existing regional GMMs. The model predicts PGA, PGV, and 5%-damped pseudo-spectral acceleration at 20 periods using magnitude (M_w), distance (R_{JB}), V_s30, and focal depth. Trained through

variational inference on 39,835 records from 3,010 earthquakes, the model achieved a test (R^2) of 0.817 and reproduced physically consistent magnitude scaling, distance attenuation, site amplification, and spectral behaviour. Comparisons with established stable-continental-region GMMs and recordings from the Bhuj and Koyna earthquakes further supported its predictive capability within the range of the available data.

The BNN framework separated within-model epistemic uncertainty from aleatory variability. Epistemic uncertainty ranged from approximately 0.04 to 0.11 natural-log units and generally increased at longer periods and in sparsely sampled regions of the predictor space. Aleatory variability ranged from approximately 0.76 to 0.89 natural-log units and dominated the total predictive uncertainty. In addition, PSHA analysis has been done to assess the influence of epistemic uncertainty on the ground motion for four major cities in PI- Mumbai, Chennai, Bangalore, and Hyderabad. The results showed that the influence of epistemic uncertainty varies with site, spectral period, and exceedance level and cannot be represented by a constant adjustment. The close agreement between continuous posterior propagation and the three-branch logic-tree approximation also demonstrates that the latter provides an efficient representation of the BNN-derived epistemic distribution for the cases considered.

The principal limitations arise from the scarcity of large-magnitude and near-source regional recordings and the extensive use of proxy-based (V_s30) estimates, which likely contributes to the relatively large site-to-site variability. Moreover, the calculated epistemic uncertainty represents uncertainty in the fitted network parameters and does not encompass all sources of model-form or regional uncertainty. Future work should incorporate additional Peninsular Indian recordings, measured site parameters, fully Bayesian or alternative probabilistic architectures, and broader PSHA applications covering additional cities and critical infrastructure.

Author Contributions

Yellapragada Meenakshi: Conceptualization, methodology, software, validation, formal analysis, investigation, data curation, visualization, and writing—original draft. S. T. G. Raghukanth: Conceptualization, methodology, supervision, resources, project administration. Both authors reviewed and approved the final manuscript.

Declaration of Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this manuscript.

Funding Declaration

The authors declare that there is no specific funding was received for this work.

Data and Resources

The ground-motion records used in this study were compiled from the Pacific Earthquake Engineering Research Center Next Generation Attenuation-East (PEER NGA-East) database

(https://peer.berkeley.edu/research/nga-east), the central and eastern United States dataset described by Gupta et al. (2017), and version 2.0 of the NEar-Source Strong-motion (NESS) flat file (https://ness.mi.ingv.it/). Region-specific recordings from the Bhuj, Jabalpur, and Koyna earthquakes were obtained from the India Meteorological Department (IMD).

References

1. Al Atik, L. (2015). NGA-East: Ground-motion standard deviation models for Central and Eastern North America. PEER Report No. 2015/07. Pacific Earthquake Engineering Research Center, University of California, Berkeley.
2. Aldama-Bustos, G., Douglas, J., Strasser, F. O., Daví, M., & MacGregor, A. (2022). Methods for assessing the epistemic uncertainty captured in ground-motion models. Bulletin of Earthquake Engineering. Advance online publication. https://doi.org/10.1007/s10518-022-01515-8
3. Boore, D. M. (1983). Stochastic simulation of high-frequency ground motions based on seismological models of the radiated spectra. Bulletin of the Seismological Society of America, 73(6A), 1865–1894. https://doi.org/10.1785/BSSA07306A1865
4. Boore, D. M. (2018). Ground-motion models for very-hard-rock sites in eastern North America: An update. Seismological Research Letters, 89(3), 1172–1184. https://doi.org/10.1785/0220170218
5. Boore, D. M., Youngs, R. R., Kottke, A. R., Bommer, J. J., Darragh, R., Silva, W. J., Stafford, P. J., Al Atik, L., Rodriguez-Marek, A., & Kaklamanos, J. (2022). Construction of a ground-motion logic tree through host-to-target adjustments applied to an adaptable ground-motion prediction model. Bulletin of the Seismological Society of America, 112(6), 3063–3082. https://doi.org/10.1785/0120220056
6. Bradley, B. A., Stirling, M. W., McVerry, G. H., & Gerstenberger, M. (2012). Consideration and propagation of epistemic uncertainties in New Zealand probabilistic seismic-hazard analysis. Bulletin of the Seismological Society of America, 102(4), 1554–1568. https://doi.org/10.1785/0120110257
7. Bureau of Indian Standards. (2002). Criteria for earthquake resistant design of structures—Part 1: General provisions and buildings (IS 1893 [Part 1]: 2002, 5th rev.). New Delhi, India.
8. Chandra, U. (1977). Earthquakes of Peninsular India—A seismotectonic study. Bulletin of the Seismological Society of America, 67(5), 1387–1413.
9. Cramer, C. H., & Kumar, A. (2003). 2001 Bhuj, India, earthquake engineering seismoscope recordings and Eastern North America ground-motion attenuation relations. Bulletin of the Seismological Society of America, 93(3), 1390–1394. https://doi.org/10.1785/0120020194
10. Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. In M. F. Balcan & K. Q. Weinberger (Eds.), Proceedings of the 33rd International Conference on Machine Learning (Vol. 48, pp. 1050–1059). Proceedings of Machine Learning Research.
11. Goulet, C. A., Bozorgnia, Y., Kuehn, N., Al Atik, L., Youngs, R. R., Graves, R. W., & Atkinson, G. M. (2021a). NGA-East ground-motion characterization model part I: Summary of products and model development. Earthquake Spectra, 37(1_suppl), 1231–1282. https://doi.org/10.1177/87552930211018723

12. Goulet, C. A., Kishida, T., Ancheta, T. D., Cramer, C. H., Darragh, R. B., Silva, W. J., Hashash, Y. M. A., Harmon, J., Parker, G. A., Stewart, J. P., & Youngs, R. R. (2021b). PEER NGA-East database. Earthquake Spectra, 37(1_suppl), 1331–1353. https://doi.org/10.1177/87552930211015695
13. Graizer, V. (2016). Ground-motion prediction equations for Central and Eastern North America. Bulletin of the Seismological Society of America, 106(4), 1600–1612. https://doi.org/10.1785/0120150374
14. Guha, S. K., & Basu, P. C. (1993). Catalogue of earthquakes (magnitude ≥ 3.0) in Peninsular India. Atomic Energy Regulatory Board Technical Document No. TD/CSE-1, Mumbai, India.
15. Gupta, A., Baker, J. W., & Ellsworth, W. L. (2017). Assessing ground-motion amplitudes and attenuation for small-to-moderate induced and tectonic earthquakes in the central and eastern United States. Seismological Research Letters, 88(5), 1379–1389. https://doi.org/10.1785/0220160199
16. Keefer, D. L., & Bodily, S. E. (1983). Three-point approximations for continuous random variables. Management Science, 29(5), 595–609. https://doi.org/10.1287/mnsc.29.5.595
17. Khosravikia, F., & Clayton, P. (2021). Machine learning in ground motion prediction. Computers & Geosciences, 148, 104700. https://doi.org/10.1016/j.cageo.2021.104700
18. MacKay, D. J. C. (1992). A practical Bayesian framework for backpropagation networks. Neural Computation, 4(3), 448–472. https://doi.org/10.1162/neco.1992.4.3.448
19. Meenakshi, Y., Vemula, S., Alne, A., & Raghukanth, S. T. G. (2023). Ground motion model for Peninsular India using an artificial neural network. Earthquake Spectra, 39(1), 596–633. https://doi.org/10.1177/87552930221144330
20. Moradi, P., Al-Harahsheh, Y., Nobahar, M., & Abu-Farsakh, M. Y. (2026). A Bayesian neural network framework for CPT prediction at sparse sites. Transportation Geotechnics, 57, 101837. https://doi.org/10.1016/j.trgeo.2025.101837
21. National Disaster Management Authority. (2010). Development of probabilistic seismic hazard map of India. Technical report, Government of India, New Delhi.
22. Neal, R. M. (1996). Bayesian learning for neural networks. Lecture Notes in Statistics (Vol. 118). Springer. https://doi.org/10.1007/978-1-4612-0745-0
23. Pearson, E. S., & Tukey, J. W. (1965). Approximate means and standard deviations based on distances between percentage points of frequency curves. Biometrika, 52(3/4), 533–546.
24. Pecker, A., Faccioli, E., Gurpinar, A., Martin, C., & Renault, P. (2017). An overview of the SIGMA research project: A European approach to seismic hazard analysis. Springer International Publishing. https://doi.org/10.1007/978-3-319-58154-5
25. Pezeshk, S., Zandieh, A., Campbell, K. W., & Tavakoli, B. (2018). Ground-motion prediction equations for Central and Eastern North America using the hybrid empirical method and NGA-West2 empirical ground-motion models. Bulletin of the Seismological Society of America, 108(4), 2278–2304. https://doi.org/10.1785/0120170179
26. Porter, K. A., Field, E. H., & Milner, K. (2012). Trimming the UCERF2 hazard logic tree. Seismological Research Letters, 83(5), 815–828. https://doi.org/10.1785/0220120012
27. Raghu Kanth, S. T. G., & Iyengar, R. N. (2007). Estimation of seismic spectral acceleration in Peninsular India. Journal of Earth System Science, 116(3), 199–214. https://doi.org/10.1007/s12040-007-0020-8
28. Sreejaya, K. P., Raghukanth, S. T. G., Gupta, I. D., Murty, C. V. R., & Srinagesh, D. (2022). Seismic hazard map of India and neighbouring regions. Soil Dynamics and Earthquake Engineering, 163, 107505. https://doi.org/10.1016/j.soildyn.2022.107505

29. Sreenath, V., & Raghukanth, S. T. G. (2023). Stochastic ground motion models to NGA-West2 and NGA-Sub databases using Bayesian neural network. Earthquake Engineering & Structural Dynamics, 52(1), 248–267. https://doi.org/10.1002/eqe.3759
30. Sriwastav, R. K., Kammula, S. T., Yedulla, J., & Raghukanth, S. T. G. (2025a). Ground motion modeling and epistemic uncertainty using conditional variational autoencoders. Machine Learning: Earth, 1(1), 015005. https://doi.org/10.1088/3049-4753/adcfa5
31. Sriwastav, R. K., Yedulla, J., Burde, P., & Raghukanth, S. T. G. (2025b). Application of conditional generative adversarial network in ground motion modelling encompassing epistemic uncertainty. Journal of Seismology, 29, 1637–1671. https://doi.org/10.1007/s10950-025-10338-4
32. Toro, G. R., Abrahamson, N. A., & Schneider, J. F. (1997). Model of strong ground motions from earthquakes in Central and Eastern North America: Best estimates and uncertainties. Seismological Research Letters, 68(1), 41–57. https://doi.org/10.1785/gssrl.68.1.41
33. Wald, D. J., & Allen, T. I. (2007). Topographic slope as a proxy for seismic site conditions and amplification. Bulletin of the Seismological Society of America, 97(5), 1379–1395. https://doi.org/10.1785/0120060267
34. Yedulla, J., Sriwastav, R. K., Choubey, A., Tulasi, K. S., & Raghukanth, S. T. G. (2025). Bayesian neural network-based ground motion model for horizontal and V/H spectral ordinate with epistemic uncertainty for the European region. Journal of Earthquake Engineering, 29(12), 2457–2484. https://doi.org/10.1080/13632469.2025.2528913
35. Zhao, J.-W., Meng, S.-B., & Liu, Z.-X. (2026). Bayesian neural network for uncertainty quantification and prediction of site nonlinear effect. Journal of Applied Geophysics, 247, 106156. https://doi.org/10.1016/j.jappgeo.2026.106156

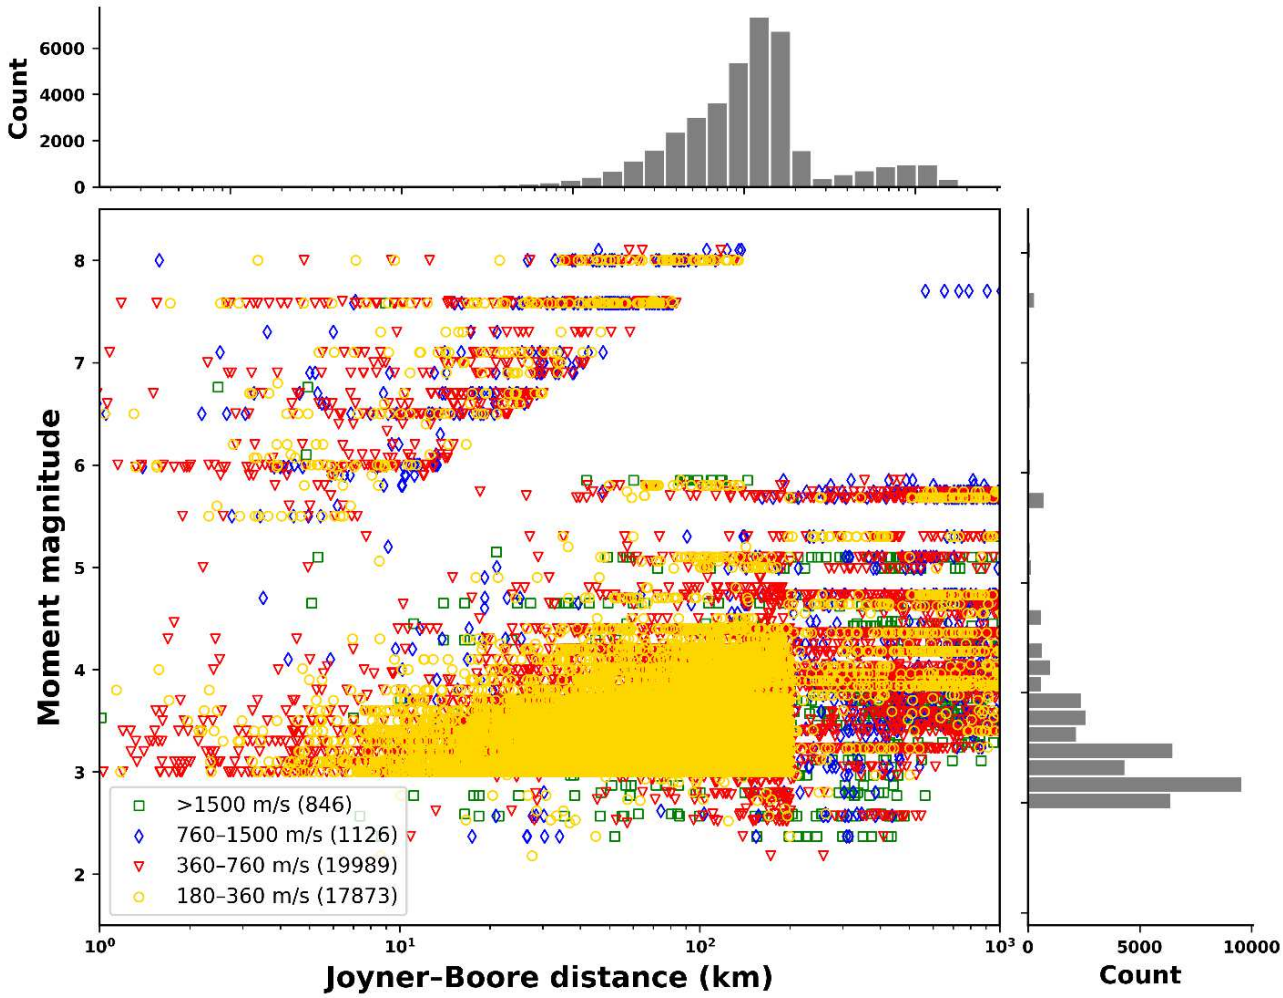


Figure 1. Distribution of moment magnitude and Joyner-Boore distance in the final ground-motion database, with marginal histograms.



Figure 2. Distributions of moment magnitude, Joyner-Boore distance, V_s30, and focal depth in the final database.

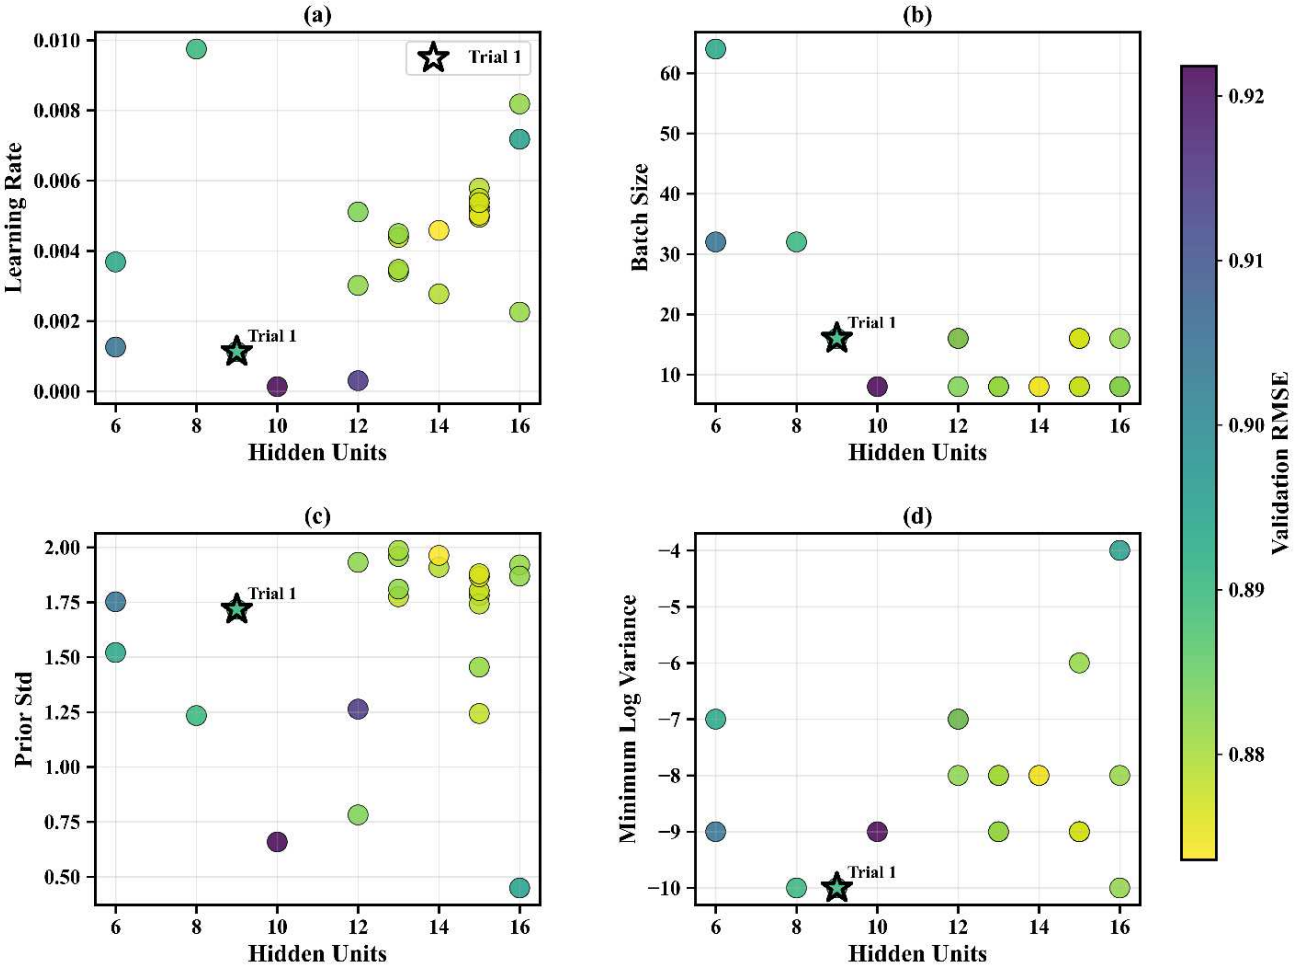


Figure 3. Hyperparameter values and validation RMSE across completed optimization trials; Trial 1 is marked.

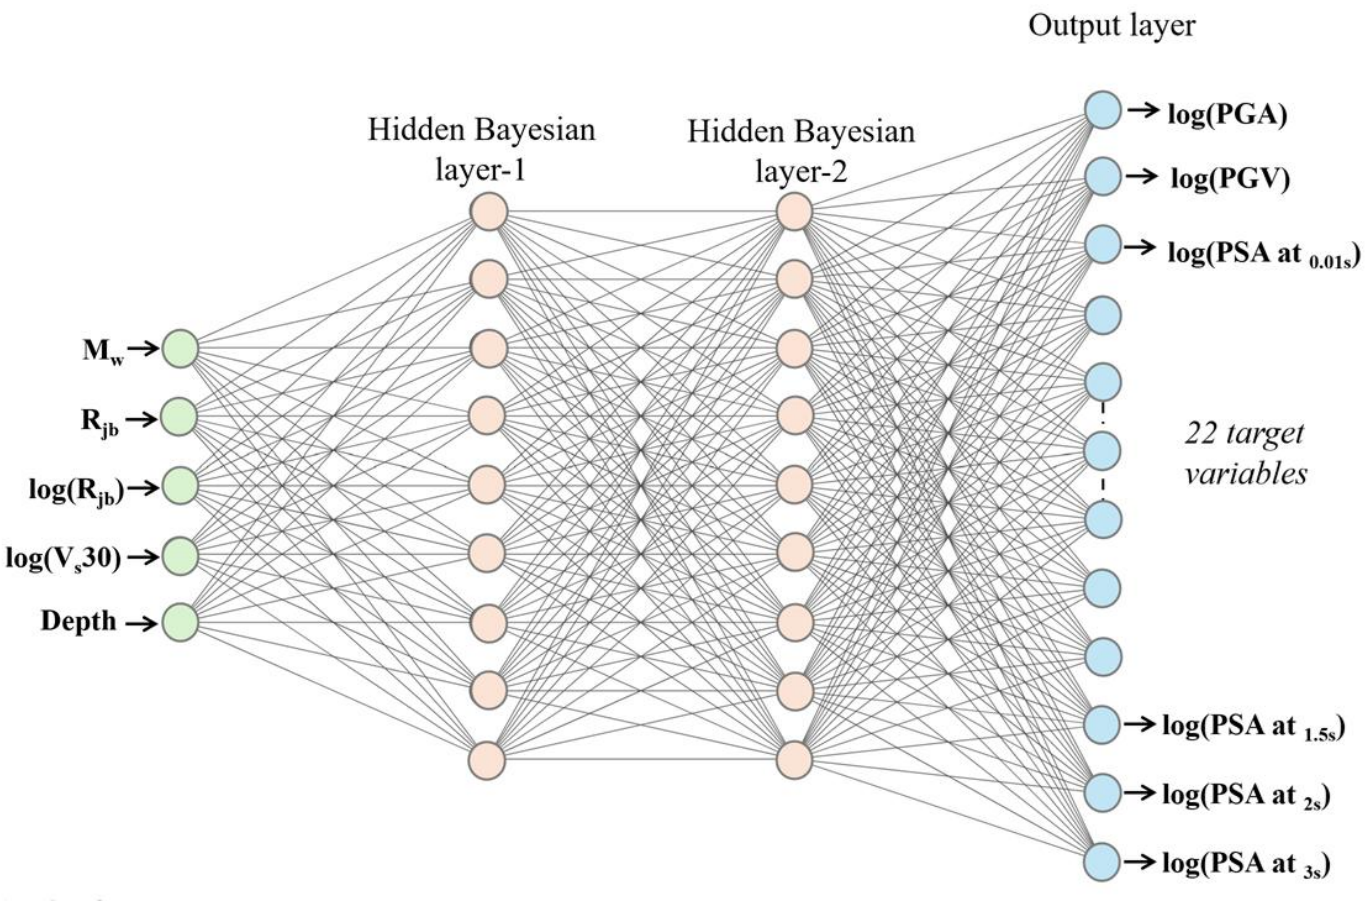


Figure 4. Architecture of the proposed Bayesian neural network ground-motion model.

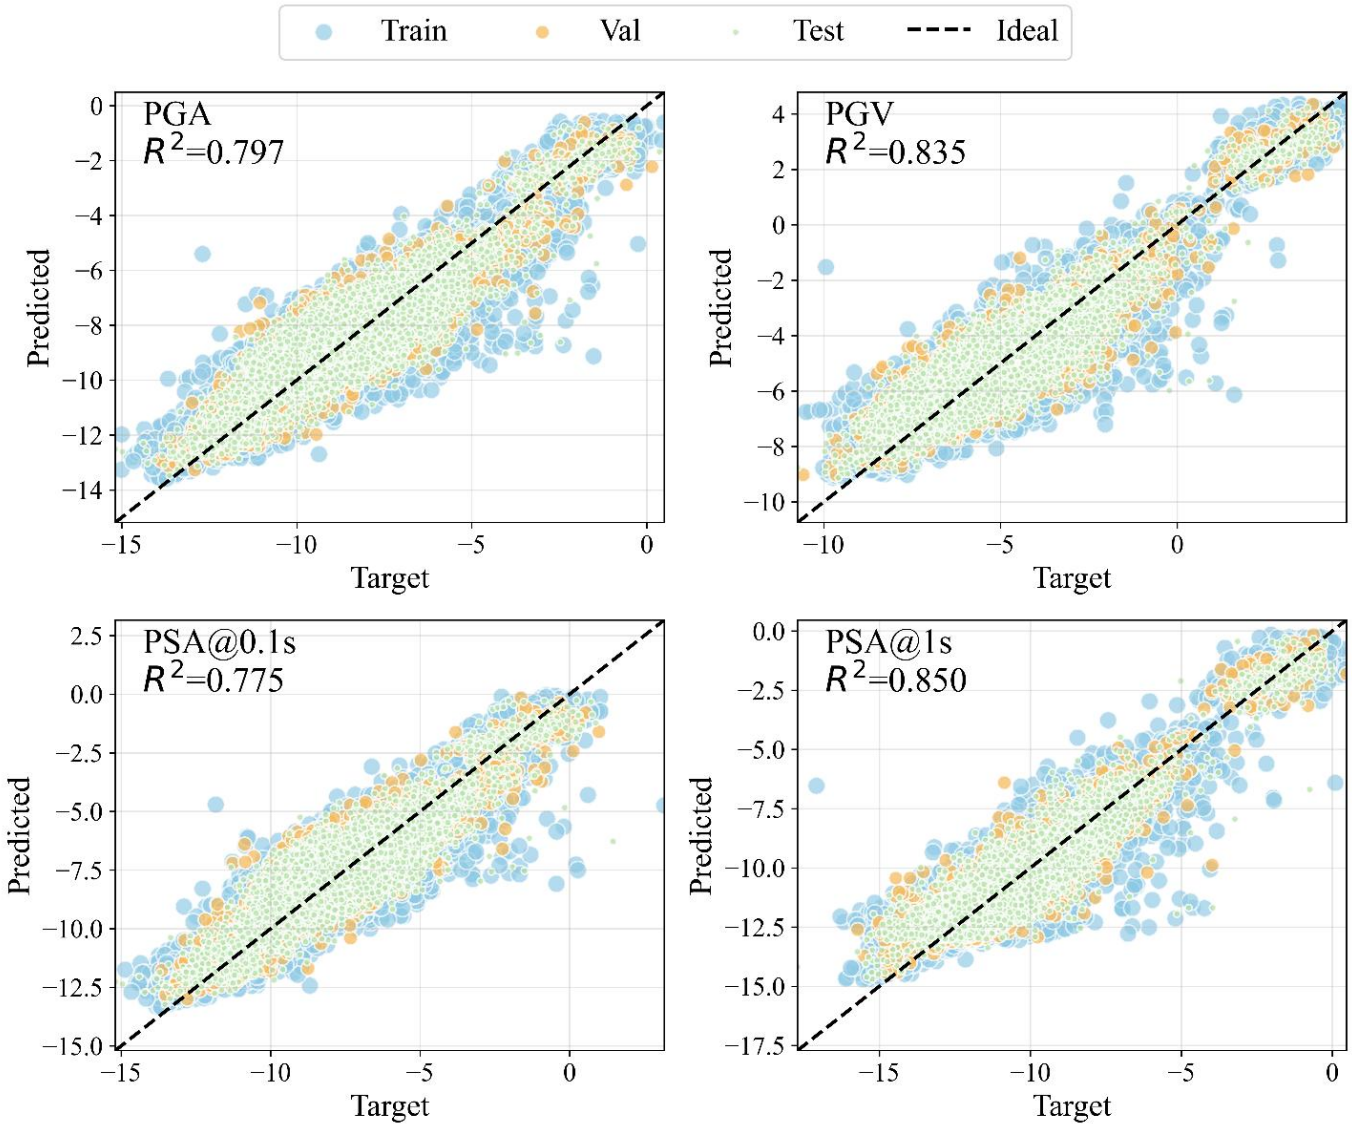


Figure 5. Target versus predicted logarithmic ground-motion values for PGA, PGV, PSA at 0.1 s, and PSA at 1 s across the training, validation, and test datasets.

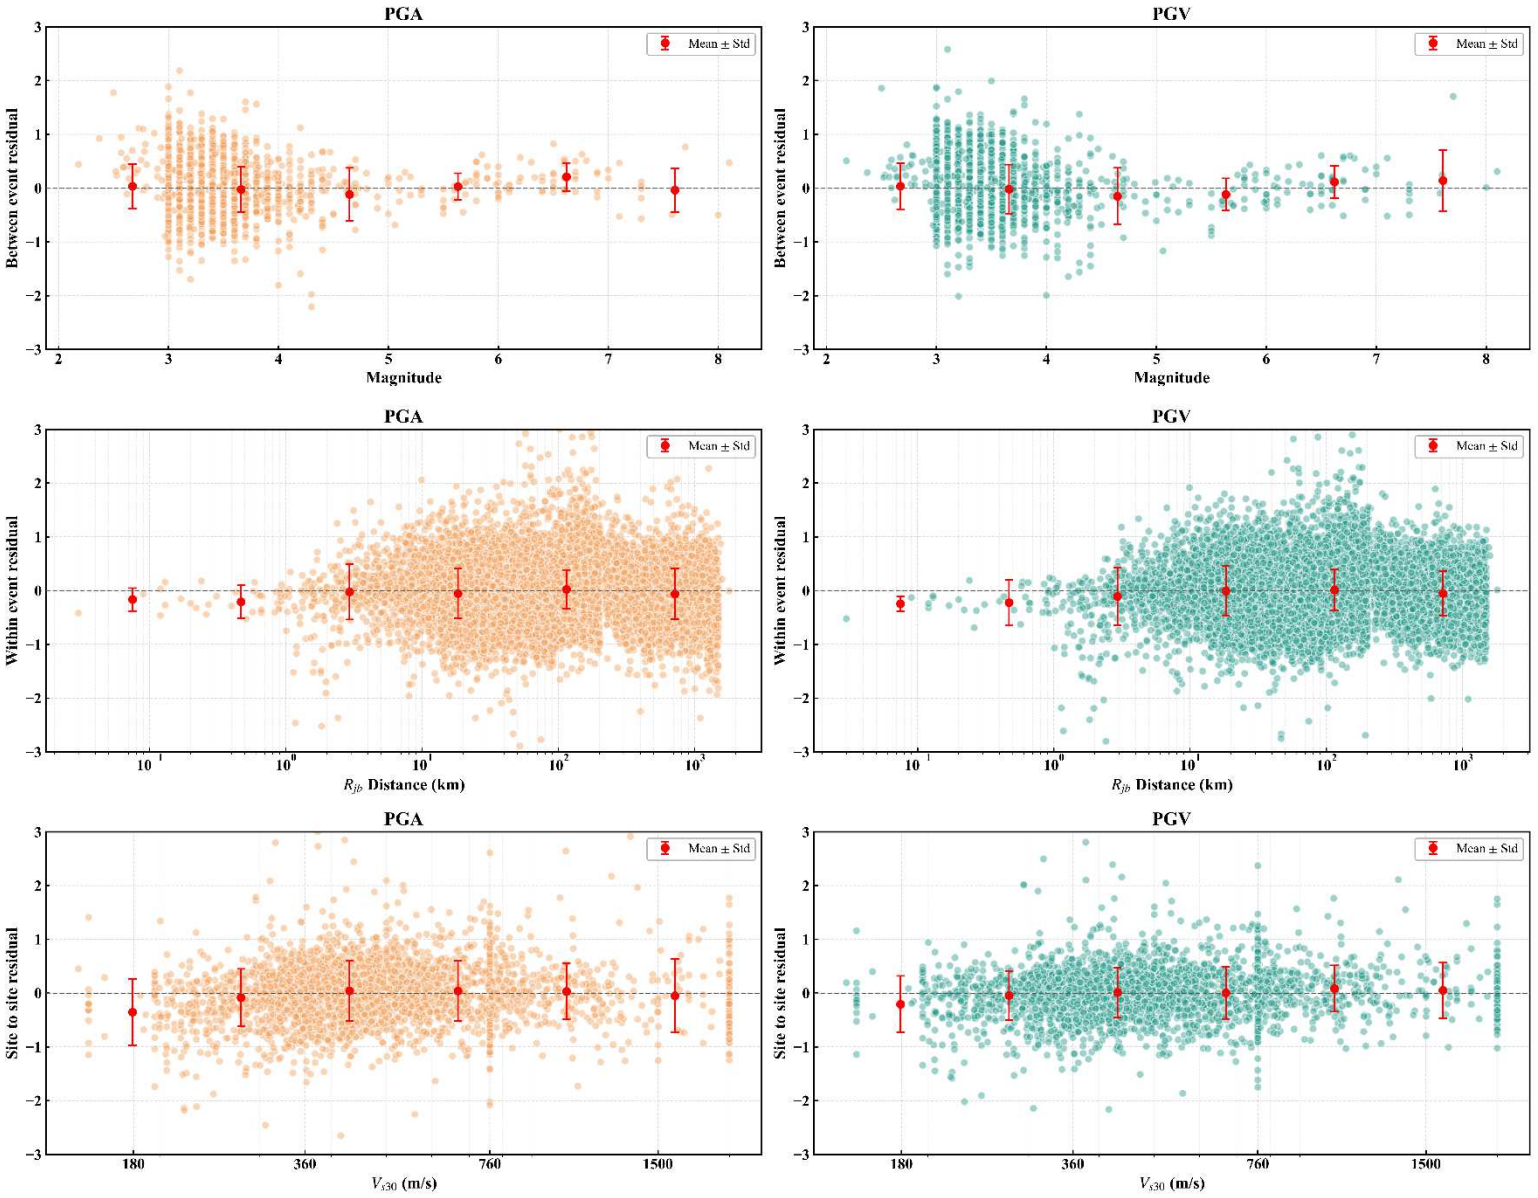


Figure 6. Residual trends with magnitude, Joyner-Boore distance, and Vs30 for PGA and PGV.

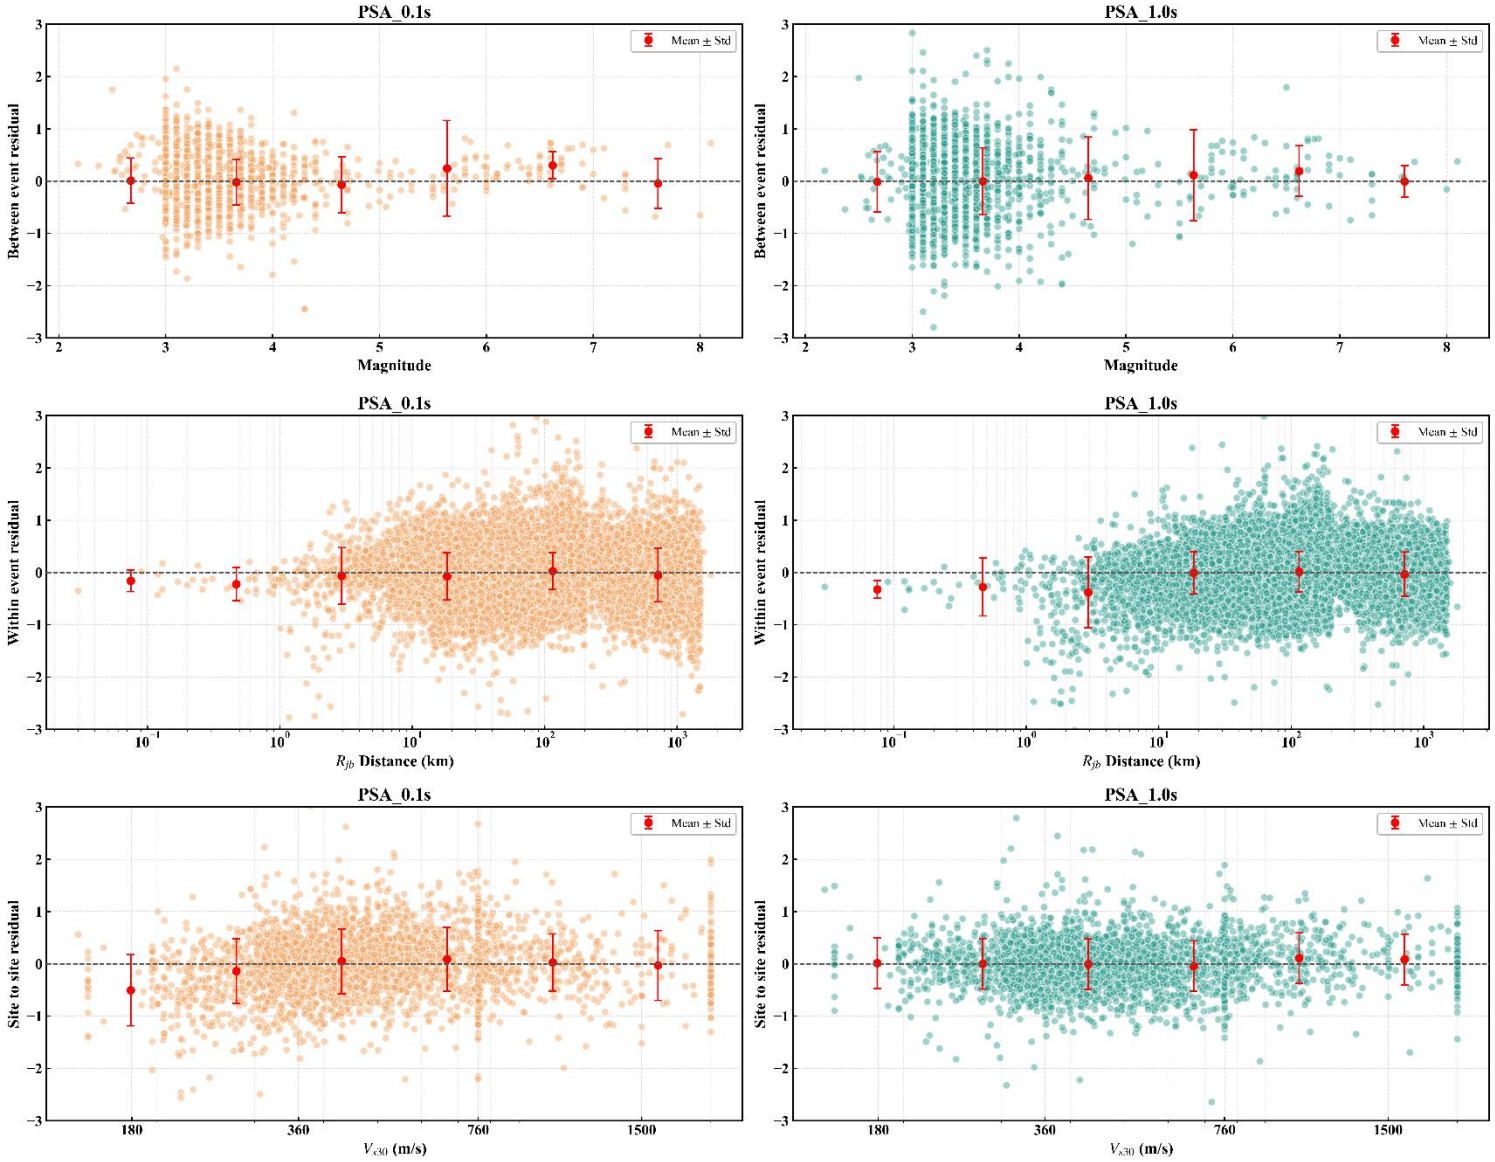


Figure 7. Residual trends with magnitude, Joyner-Boore distance, and Vs30 for PSA at 0.1 s and PSA at 1 s.

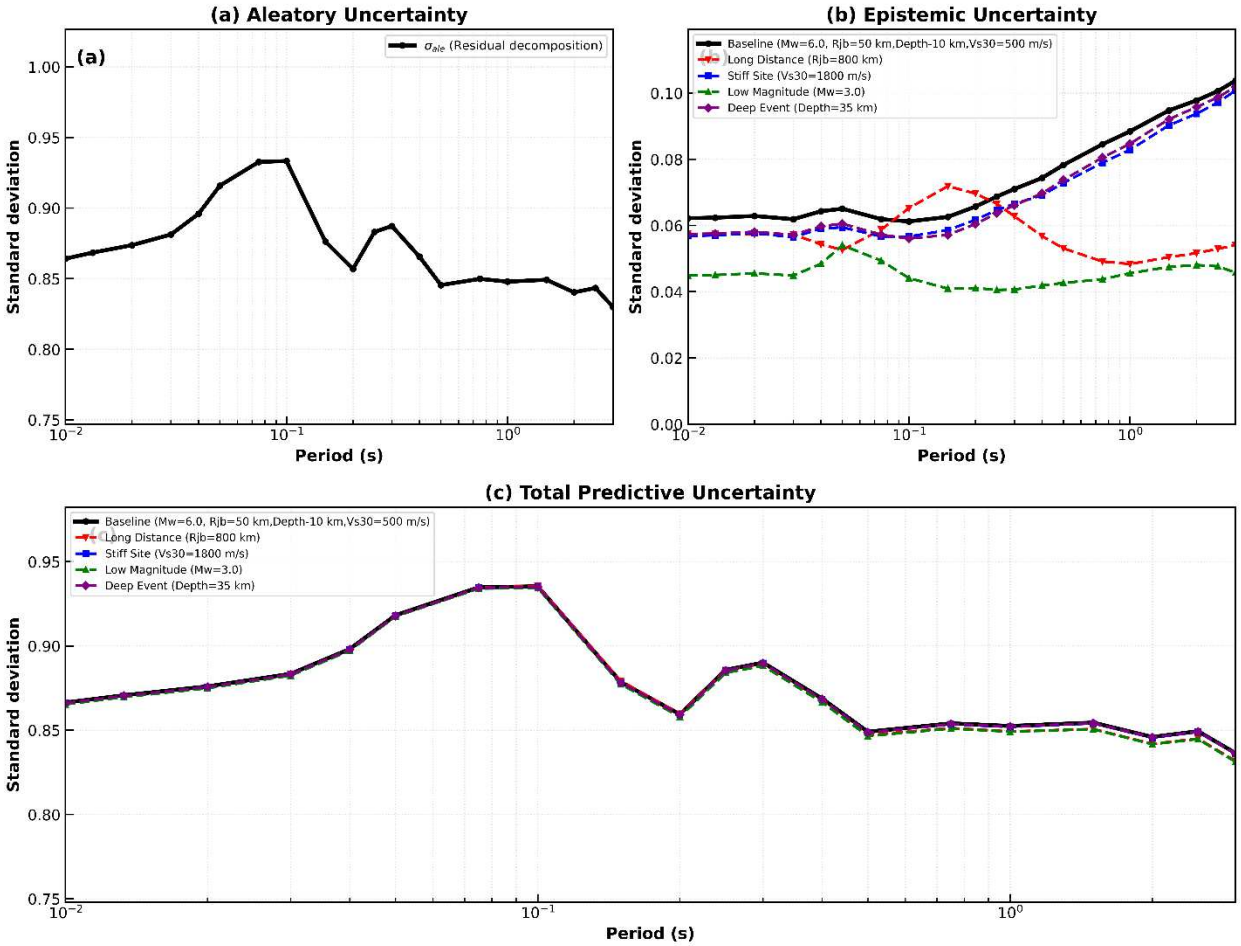


Figure 9. Aleatory, epistemic, and total predictive uncertainty as functions of oscillator period.

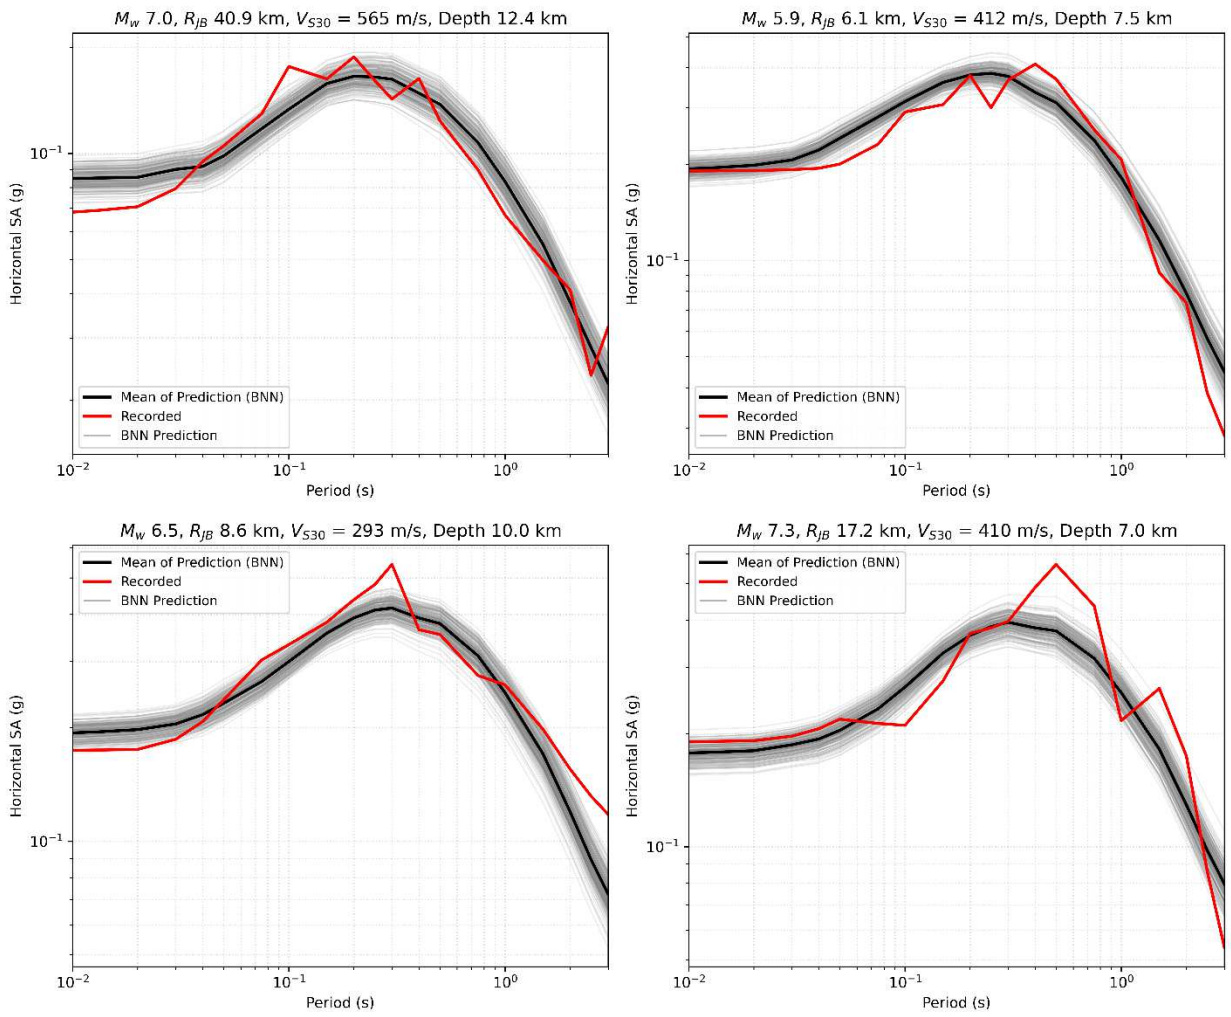


Figure 8. Recorded response spectra, posterior samples, and mean BNN predictions for four earthquake scenarios.

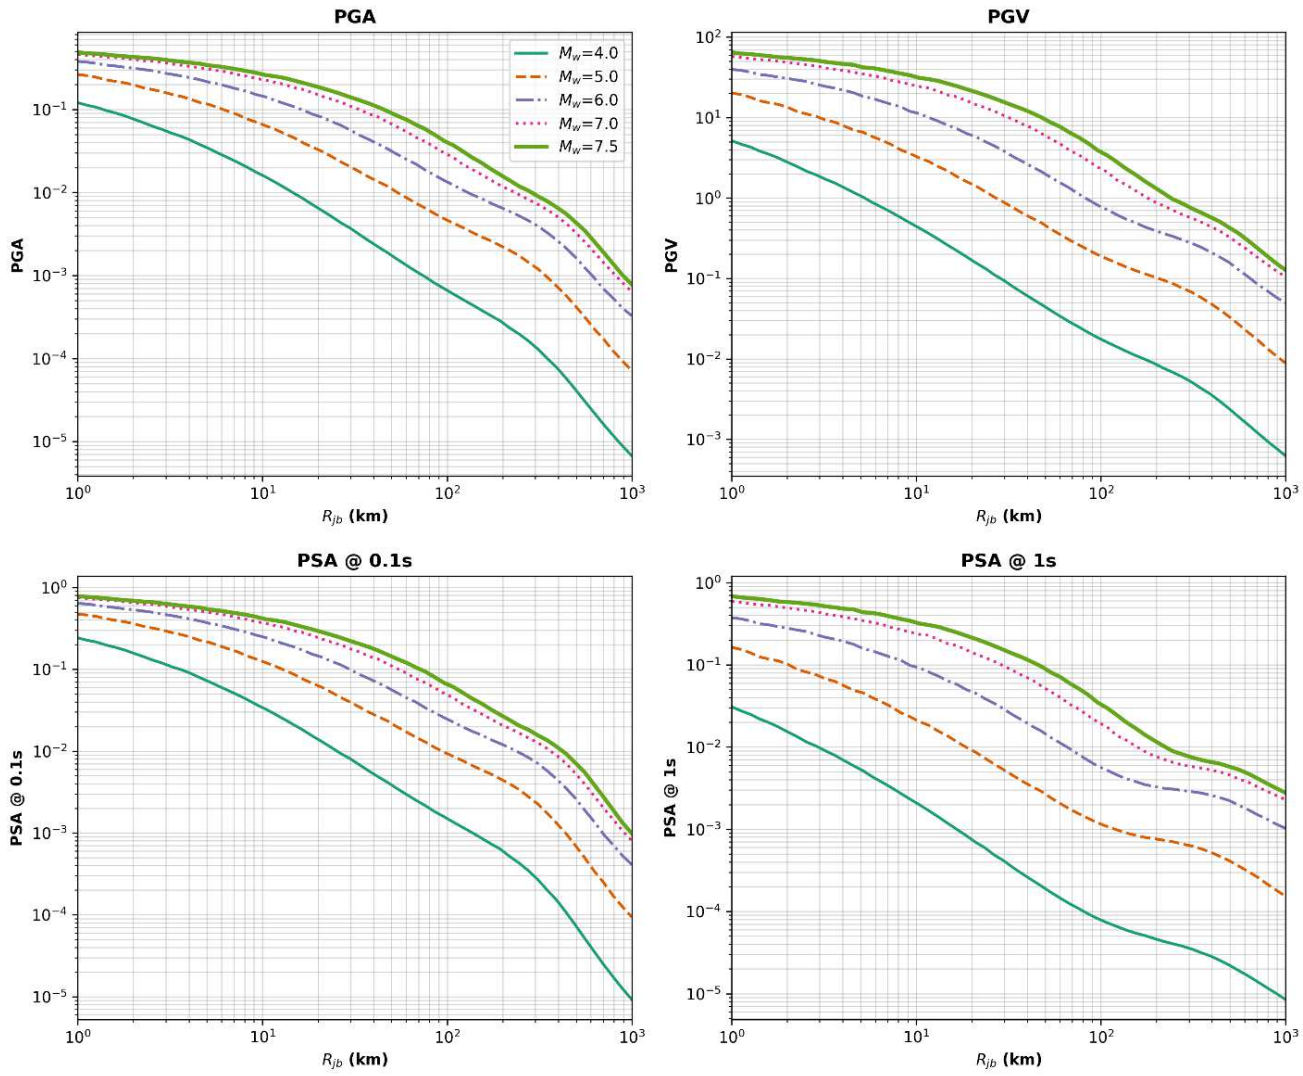


Figure 10. Predicted PGA, PGV, PSA at 0.1 s, and PSA at 1 s versus Joyner-Boore distance for different magnitudes.

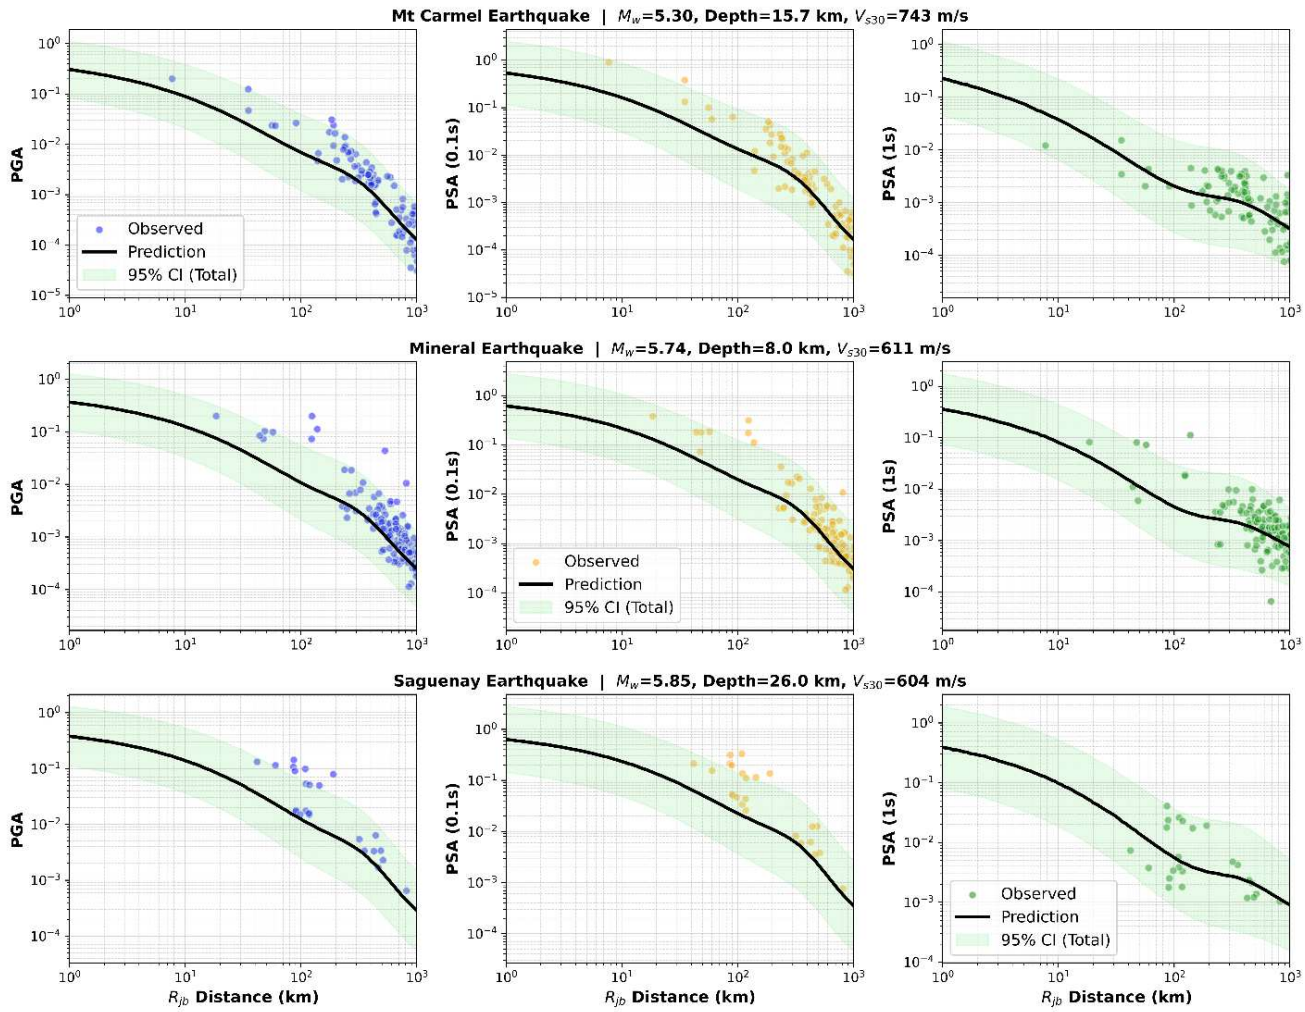


Figure 11. Event-specific observed and predicted ground motions with 95% predictive intervals.

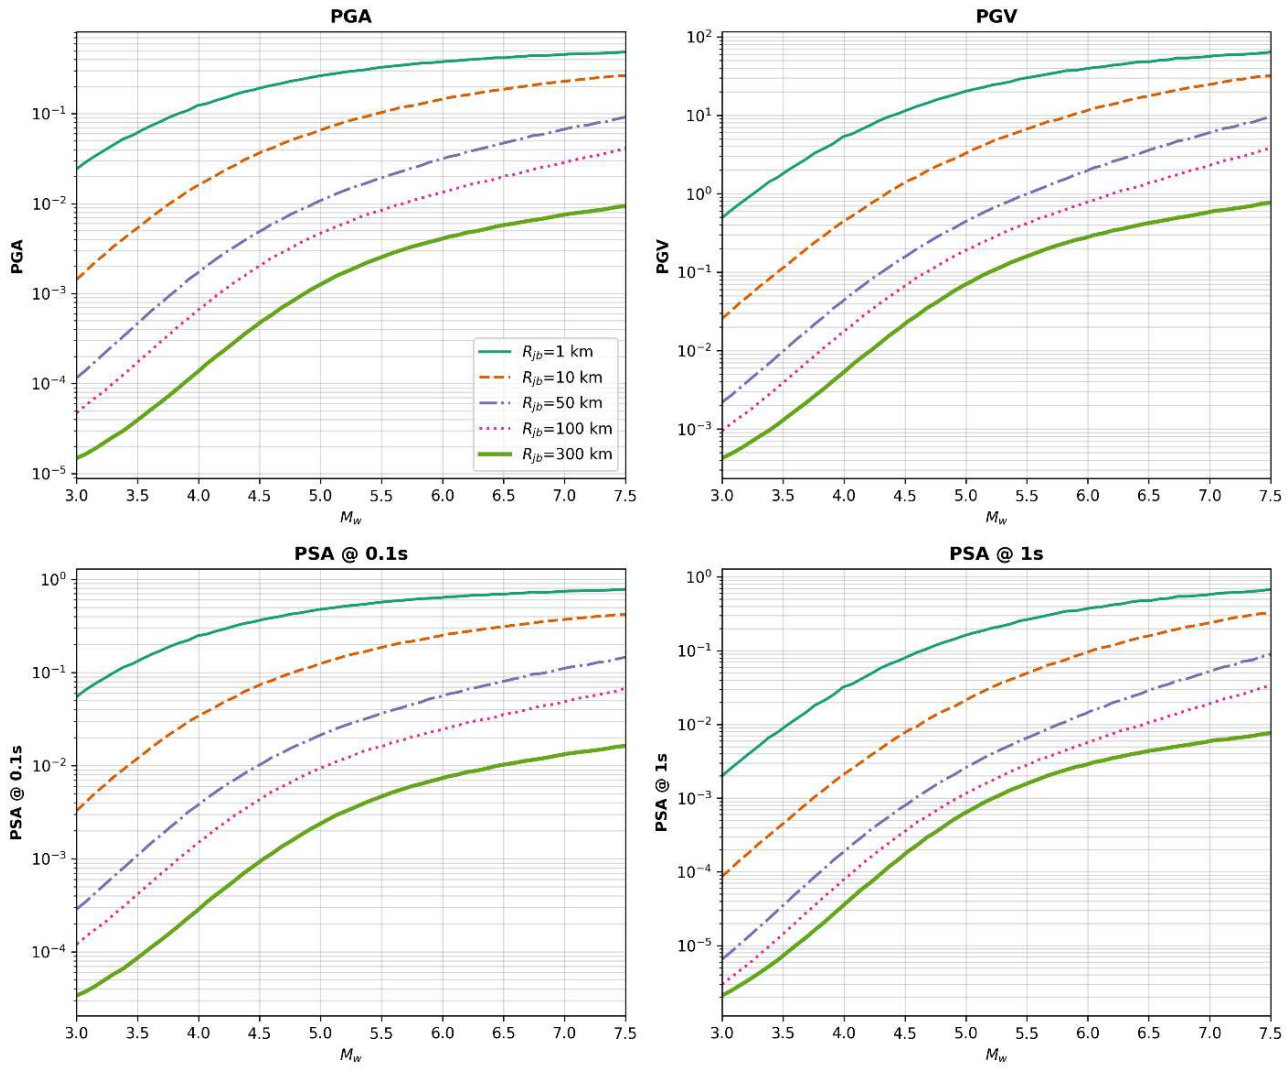


Figure 12. Magnitude scaling of PGA, PGV, PSA at 0.1 s, and PSA at 1 s for different source-to-site distances.

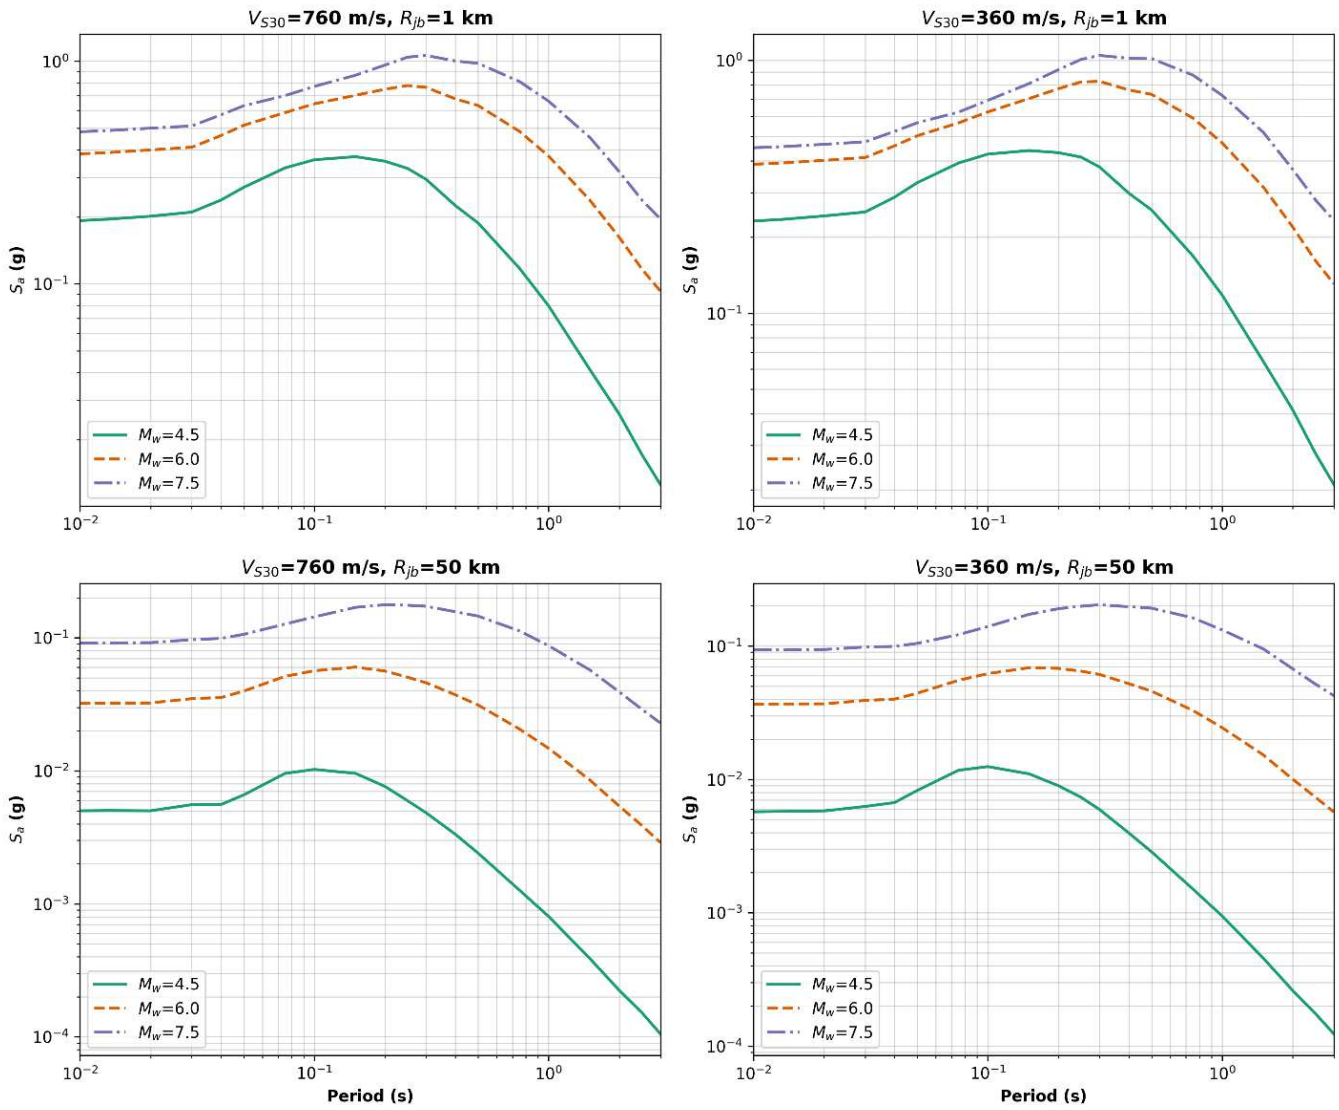


Figure 13. Influence of Vs30 on predicted response spectra for different magnitudes and distances.

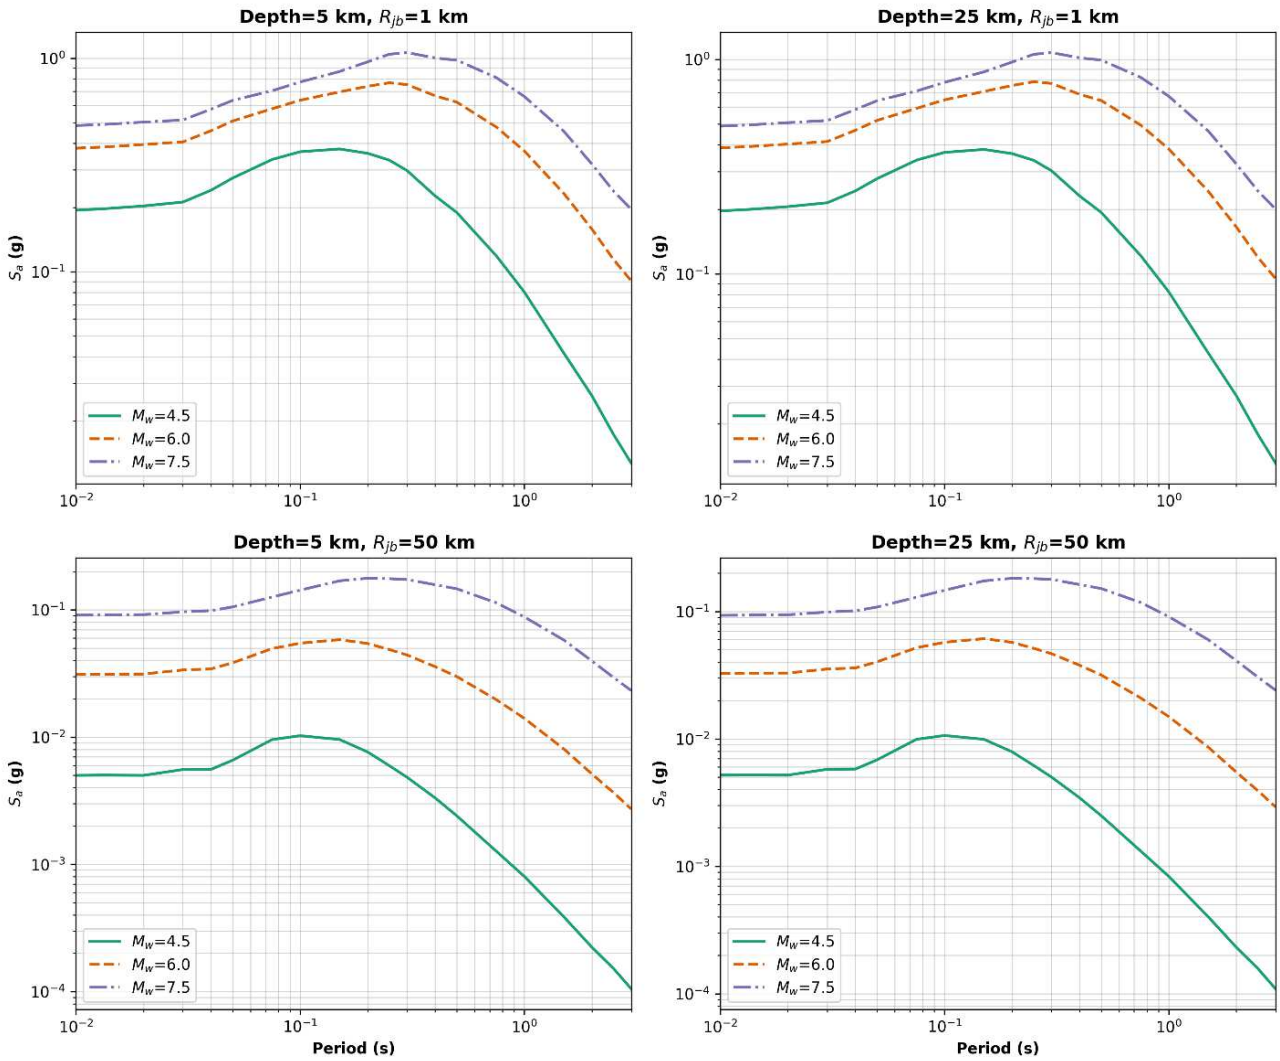


Figure 15. Influence of focal depth on predicted response spectra for different magnitudes and distances.

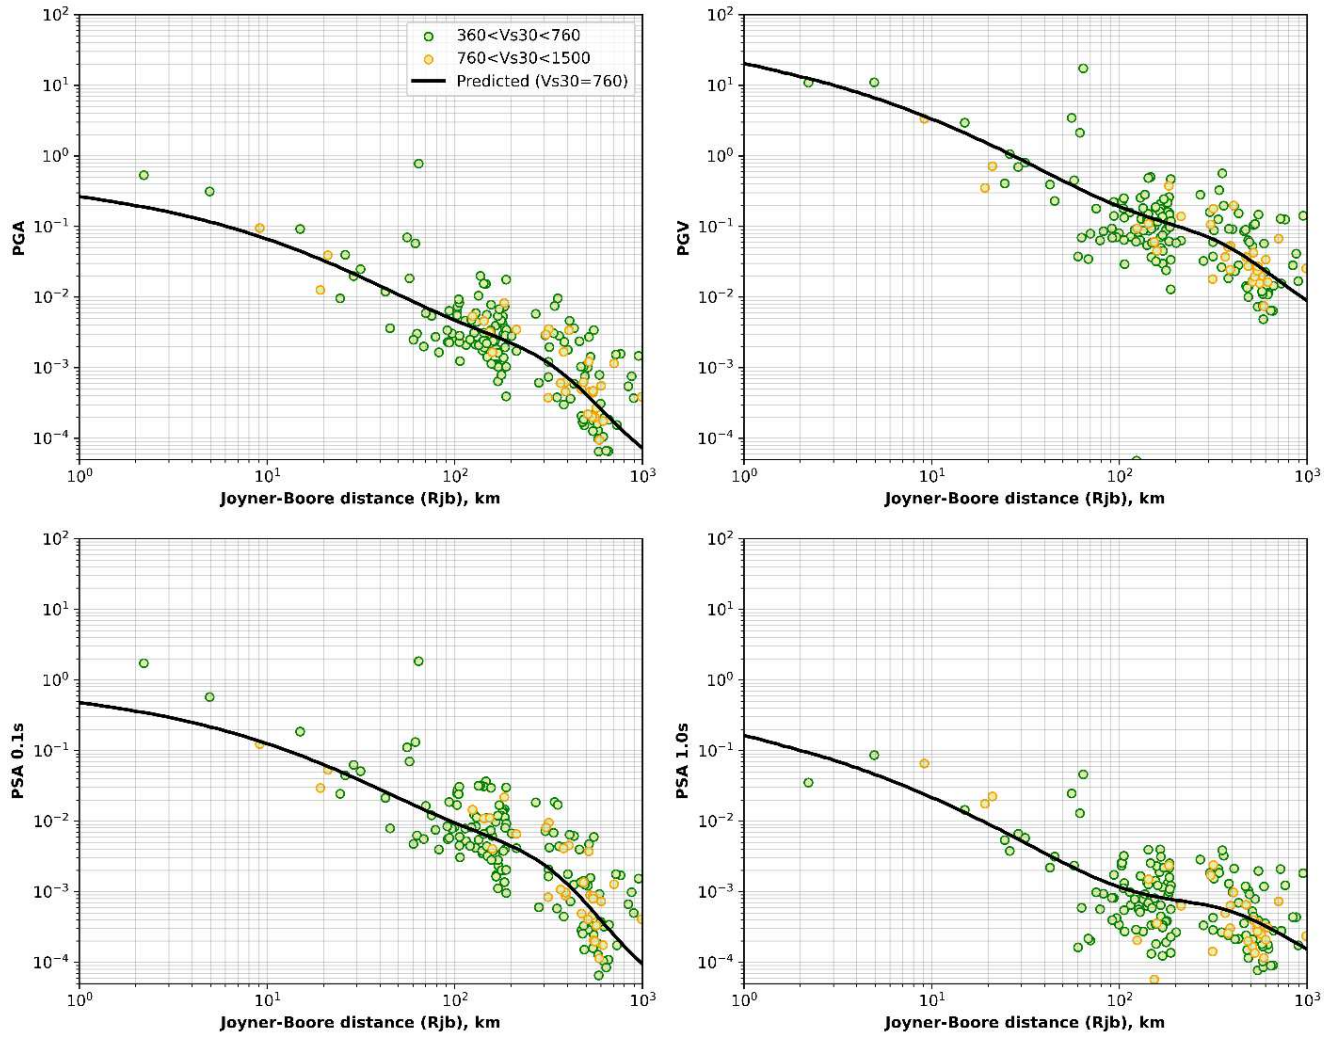


Figure 14. Observed and predicted ground motions grouped by Vs30 range as functions of Joyner-Boore distance.

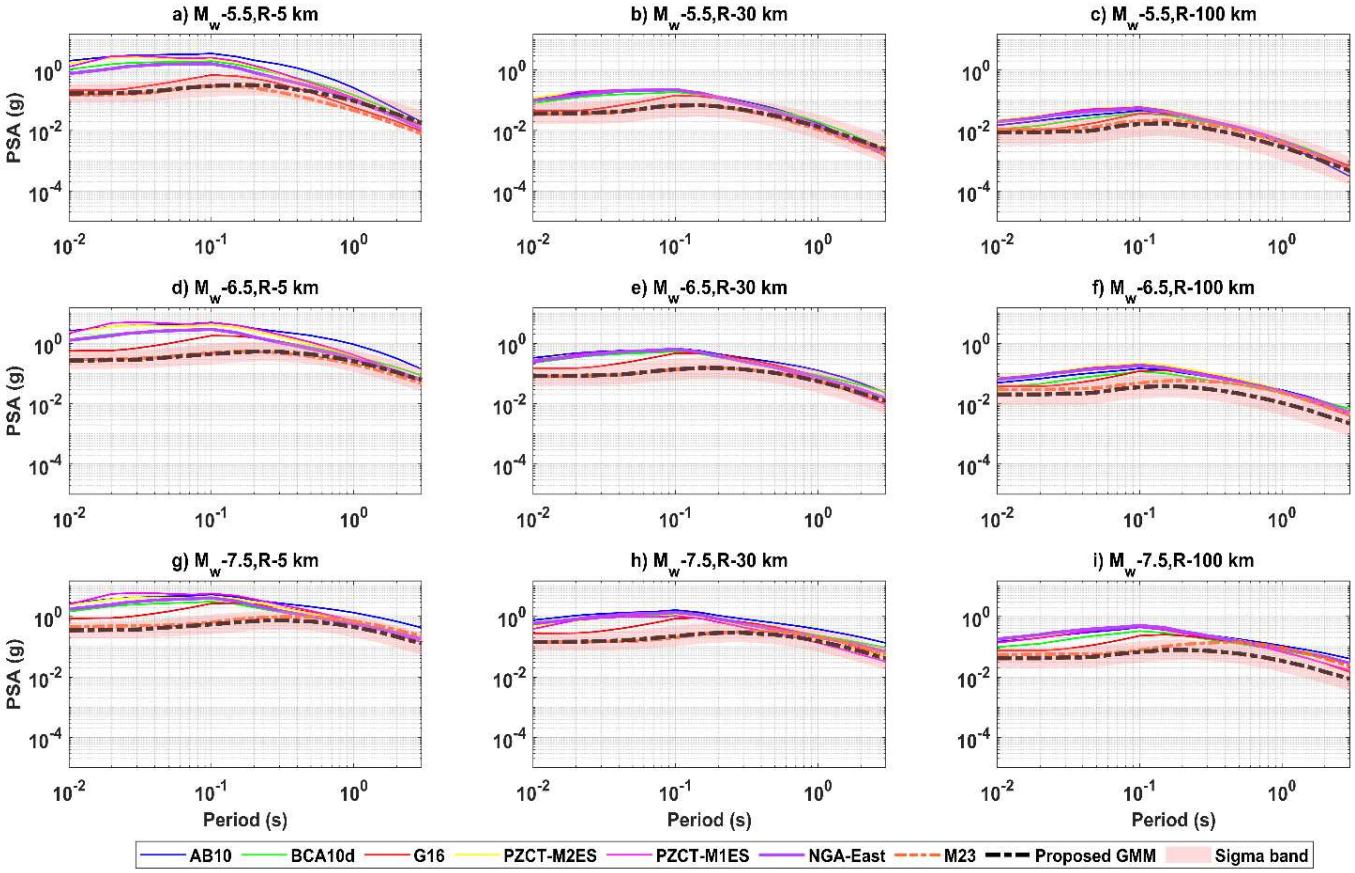


Figure 17. Comparison of predicted response spectra with published GMMs for different magnitudes and distances.

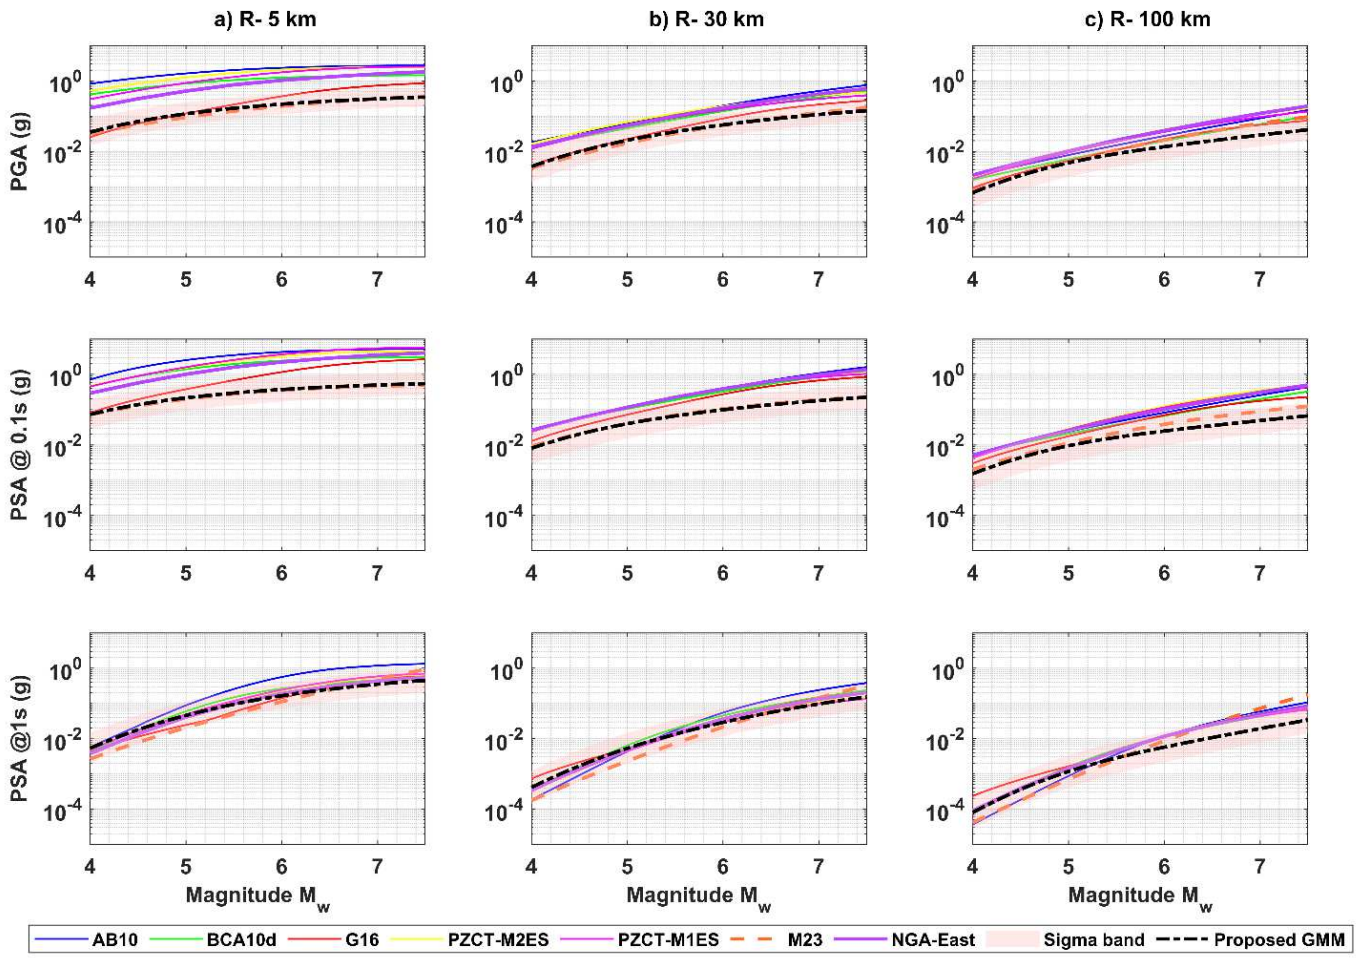


Figure 16. Comparison of predicted ground-motion amplitudes with published GMMs as functions of magnitude at different distances.

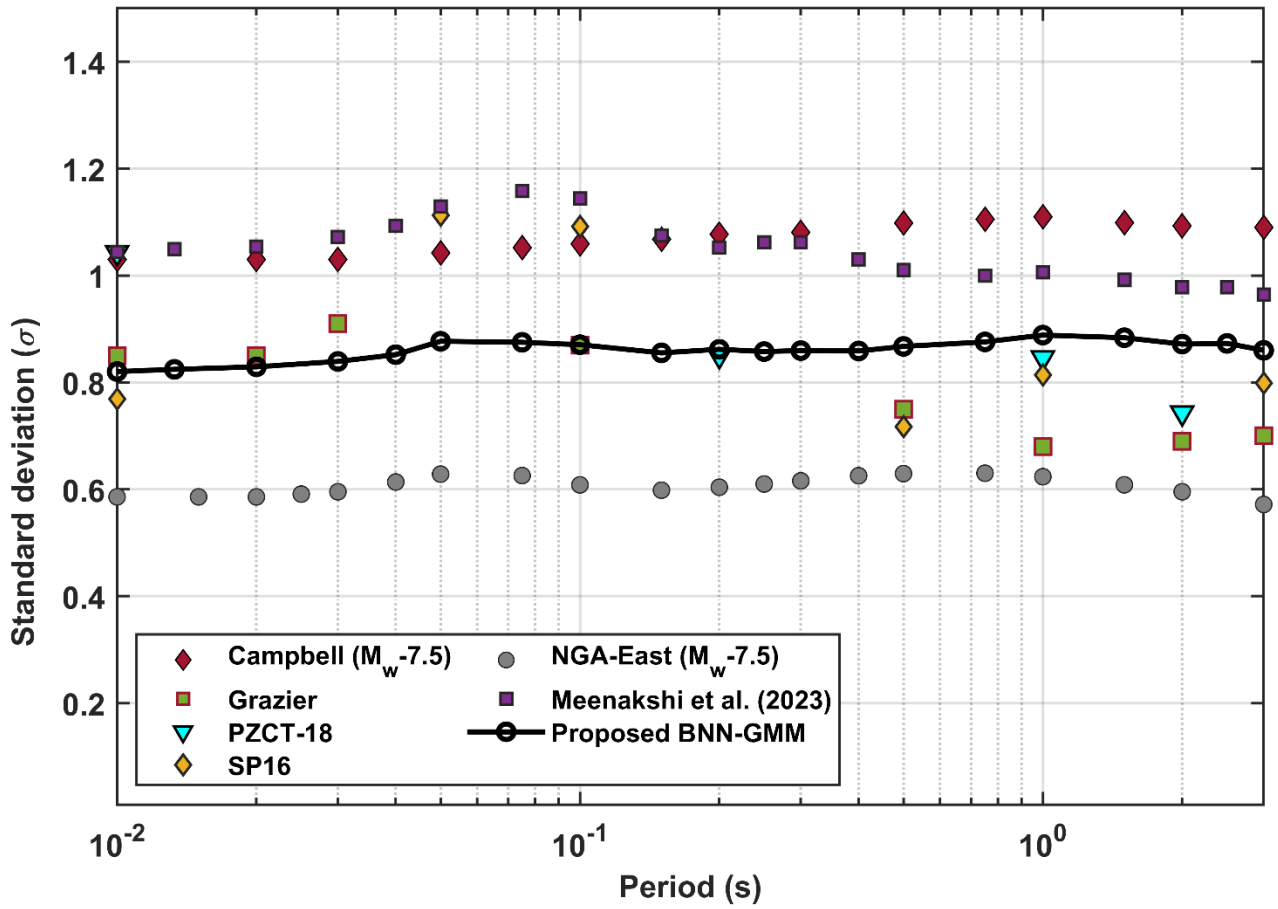


Figure 18. Comparison of total standard deviation and its aleatory and epistemic components with values from published GMMs.

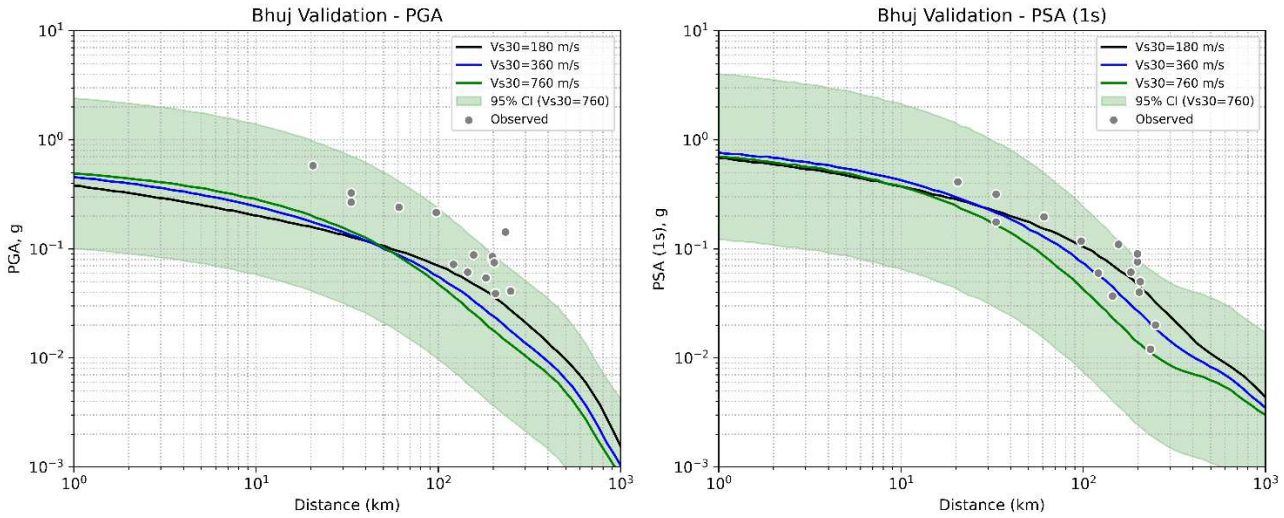


Figure 19. Comparison of recorded Bhuj earthquake ground motions with BNN-GMM predictions and 95% predictive intervals.

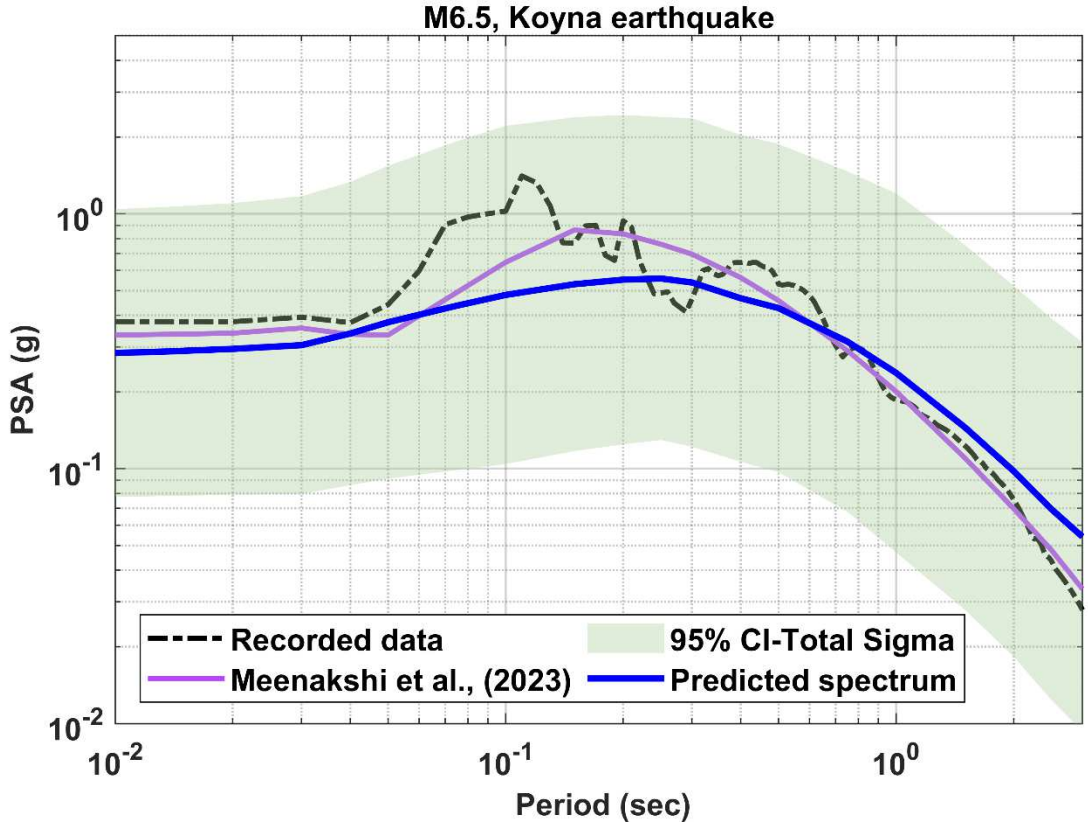


Figure 20. Recorded response spectrum of the Koyna earthquake compared with the BNN-GMM and Meenakshi et al. (2023).

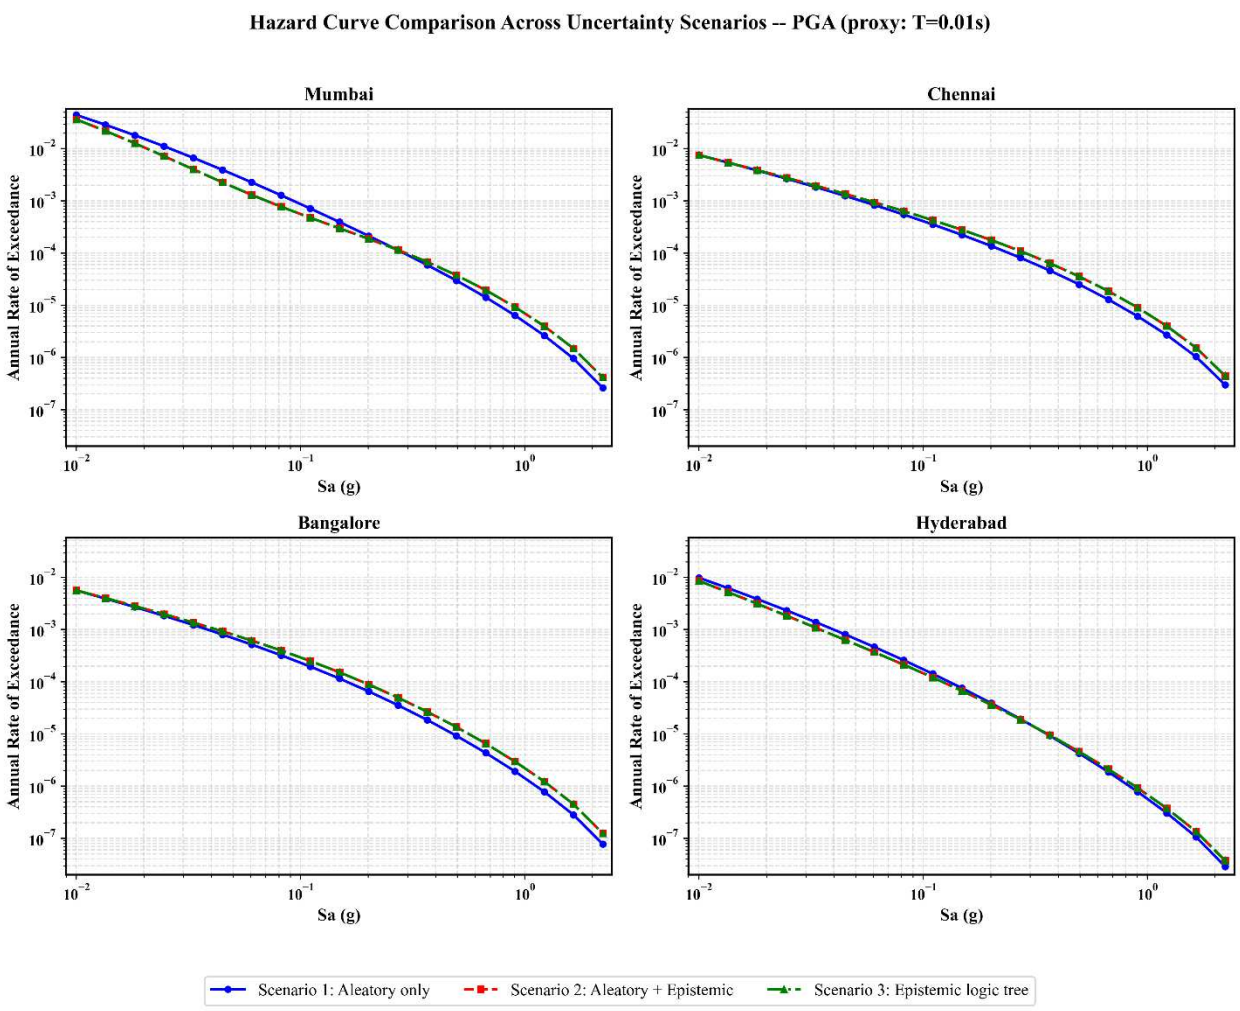


Figure 21. Annual exceedance hazard curves for PGA at Mumbai, Chennai, Bangalore, and Hyderabad under the three uncertainty scenarios.

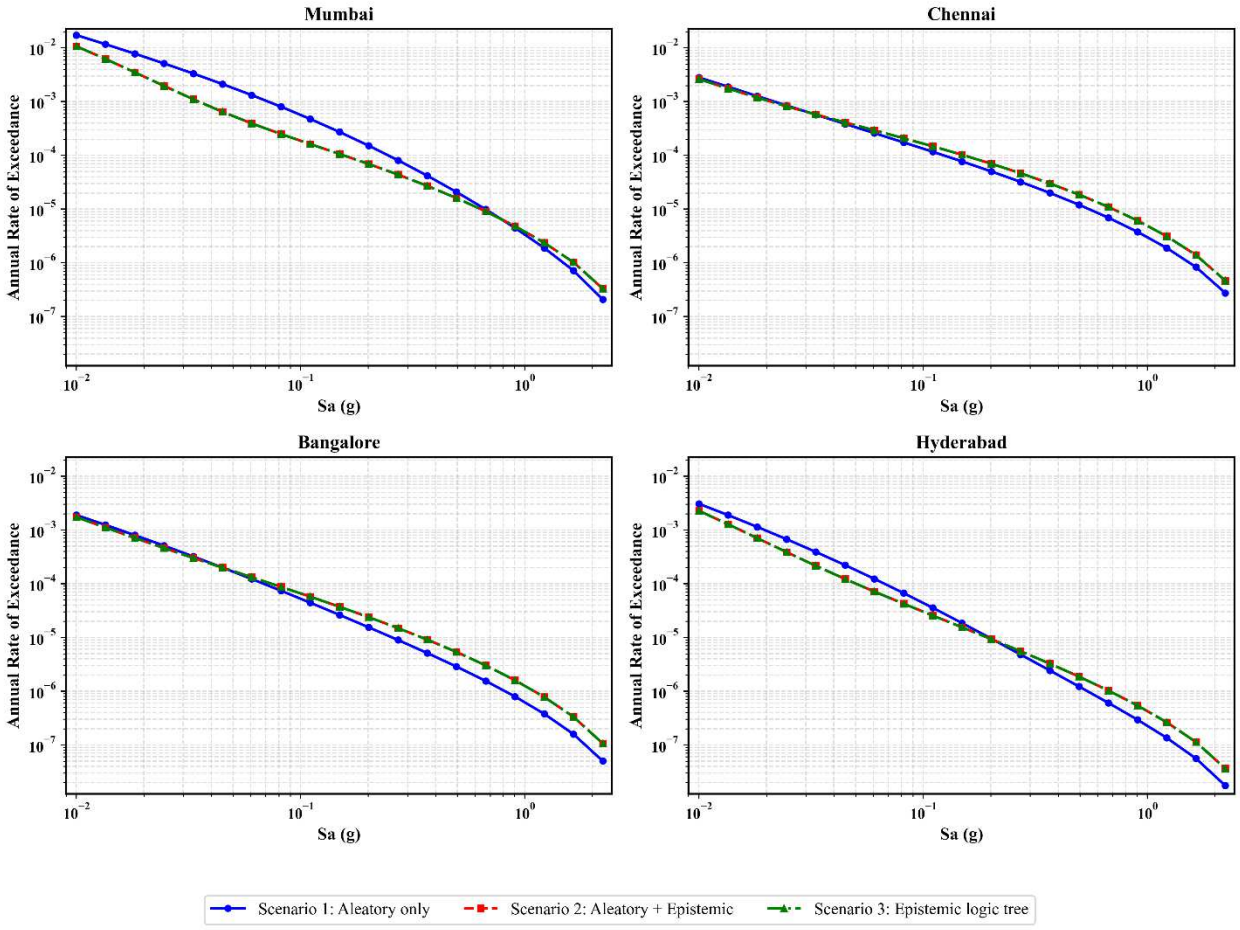


Figure 22. Annual exceedance hazard curves for PSA at 1 s at Mumbai, Chennai, Bangalore, and Hyderabad under the three uncertainty scenarios.

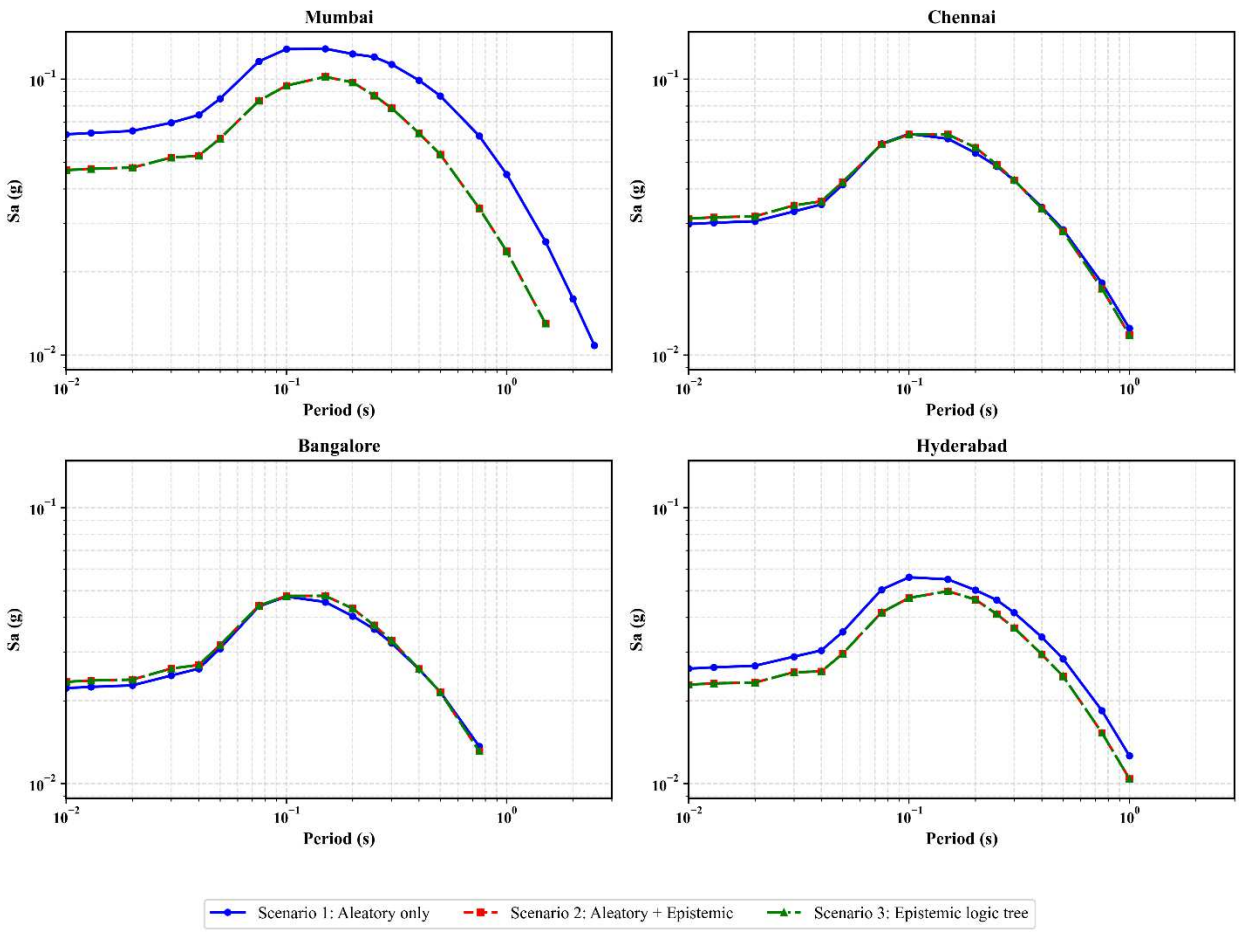


Figure 23. Uniform hazard response spectra for Mumbai, Chennai, Bangalore, and Hyderabad at a 475-year return period under the three uncertainty scenarios.

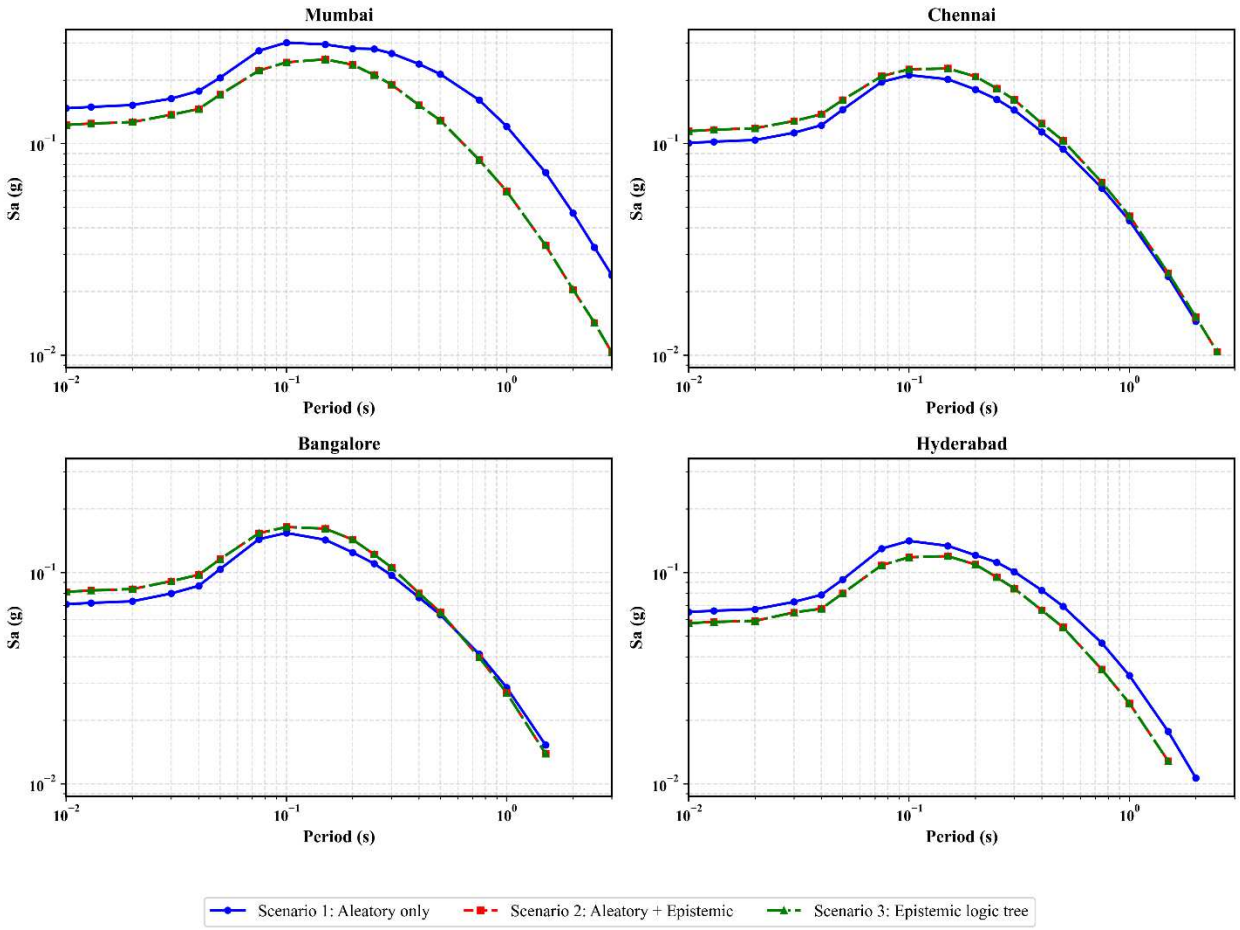


Figure 24. Uniform hazard response spectra for Mumbai, Chennai, Bangalore, and Hyderabad at a 2,475-year return period under the three uncertainty scenarios.

Table 1. Statistics of the input and target variables

Input variables							
Variable	**Mean**	**Median**	**Std**	**Min**	**Max**	**Skewness**	**Kurtosis**
M_w	3.567	3.300	0.793	2.180	8.100	3.015	10.840
R_{JB}	182.450	109.970	259.535	0.030	1794.07	3.048	9.127
log (R_{JB})	4.636	4.700	1.048	-3.507	7.492	-0.157	1.945
log (V_{s30})	6.030	5.936	0.455	4.970	7.601	1.016	1.265
Depth	6.789	5.895	9.608	-999.000	36.000	-86.211	9042.395

Target variables							
Variable	**Mean**	**Median**	**Std**	**Min**	**Max**	**Skewness**	**Kurtosis**
PGA	-8.394	-8.534	2.018	-15.186	0.500	0.855	1.910
PGV	-5.100	-5.407	2.075	-10.748	4.790	1.582	4.182
PSA at 0.01s	-8.389	-8.530	2.021	-15.187	2.756	0.861	1.946
0.0133s	-8.386	-8.525	2.025	-15.188	2.924	0.860	1.932
0.02s	-8.382	-8.518	2.033	-15.187	3.197	0.856	1.911
0.03s	-8.305	-8.442	2.035	-15.182	3.271	0.839	1.888
0.04s	-8.243	-8.349	2.077	-15.181	3.470	0.781	1.641
0.05s	-7.995	-8.041	2.126	-15.189	3.492	0.604	1.223
0.075s	-7.544	-7.572	2.117	-15.165	3.157	0.460	1.066
0.10s	-7.470	-7.521	2.084	-15.196	3.163	0.485	1.272
0.15s	-7.693	-7.868	2.012	-15.171	2.760	0.807	2.060
0.20s	-7.962	-8.219	1.993	-15.116	2.234	1.082	2.638
0.25s	-8.184	-8.461	2.005	-15.001	1.821	1.229	2.918
0.30s	-8.447	-8.756	2.016	-14.871	1.507	1.375	3.328
0.40s	-8.950	-9.312	2.044	-15.660	1.207	1.598	4.081
0.50s	-9.370	-9.757	2.080	-16.241	0.961	1.728	4.596
0.75s	-10.155	-10.581	2.171	-17.119	0.924	1.836	5.083
1.00s	-10.729	-11.170	2.236	-17.714	0.465	1.876	5.247
1.50s	-11.570	-12.051	2.305	-18.543	0.639	1.972	5.683
2.00s	-12.202	-12.720	2.341	-19.122	0.213	2.061	6.002
2.50s	-12.625	-13.168	2.352	-19.569	0.285	2.120	6.208
3.00s	-12.990	-13.553	2.358	-19.935	-0.350	2.197	6.514

Table 2. Performance metrics for PGA, PGV, and PSA across oscillator periods.

Period	MSE	RMSE	MAE	R2	R
PGA	0.826	0.909	0.704	0.797	0.893
PGV	0.709	0.842	0.644	0.835	0.914
0.01	0.831	0.912	0.704	0.796	0.893
0.0133	0.839	0.916	0.707	0.795	0.892
0.02	0.847	0.92	0.711	0.795	0.892
0.03	0.867	0.931	0.716	0.791	0.889
0.04	0.896	0.946	0.728	0.792	0.89
0.05	0.934	0.966	0.745	0.793	0.891
0.075	0.987	0.993	0.773	0.78	0.883
0.10	0.975	0.987	0.775	0.775	0.881
0.15	0.851	0.922	0.716	0.79	0.889
0.20	0.803	0.896	0.689	0.798	0.893
0.25	0.841	0.917	0.707	0.791	0.889
0.30	0.847	0.92	0.713	0.792	0.89
0.40	0.795	0.891	0.688	0.81	0.9
0.50	0.753	0.868	0.67	0.826	0.909
0.75	0.749	0.866	0.67	0.841	0.917
1.00	0.748	0.865	0.668	0.85	0.922
1.50	0.748	0.865	0.664	0.859	0.927
2.00	0.733	0.856	0.658	0.866	0.931
2.50	0.742	0.861	0.662	0.866	0.931
3.00	0.721	0.849	0.648	0.87	0.933

Table 3. Aleatory, epistemic, and total standard-deviation components across oscillator periods.

Period	Between event sigma	Site-to-site sigma	Within event-single station sigma	Aleatory sigma	Epistemic sigma	Total sigma
PGA	0.42	0.566	0.399	0.809	0.042	0.811
PGV	0.445	0.472	0.401	0.763	0.042	0.764
0.01	0.44	0.565	0.398	0.819	0.042	0.82
0.0133	0.442	0.568	0.4	0.823	0.042	0.824
0.02	0.447	0.57	0.4	0.828	0.042	0.829
0.03	0.447	0.582	0.404	0.838	0.041	0.839
0.04	0.444	0.6	0.408	0.851	0.045	0.852
0.05	0.437	0.64	0.407	0.876	0.05	0.877
0.075	0.445	0.639	0.396	0.874	0.048	0.875
0.10	0.45	0.628	0.399	0.869	0.044	0.871
0.15	0.46	0.599	0.399	0.854	0.043	0.855
0.20	0.475	0.592	0.406	0.861	0.043	0.862
0.25	0.481	0.578	0.41	0.857	0.042	0.858
0.30	0.499	0.563	0.413	0.858	0.042	0.859
0.40	0.53	0.531	0.415	0.858	0.042	0.859
0.50	0.557	0.513	0.421	0.866	0.042	0.867
0.75	0.598	0.482	0.418	0.875	0.042	0.876
1.00	0.623	0.486	0.404	0.887	0.043	0.888
1.50	0.631	0.458	0.412	0.882	0.045	0.884
2.00	0.624	0.45	0.408	0.87	0.046	0.872
2.50	0.634	0.44	0.405	0.872	0.046	0.873
3.00	0.621	0.432	0.405	0.859	0.045	0.86

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supplementarydocumentdraft1.docx