Accounting for path and site effects in spatial ground-motion correlation models using Bayesian inference

Lukas Bodenmann[1], Jack W. Baker[2], and Božidar Stojadinović[1]

[1]Department of Civil, Environmental and Geomatic Engineering, ETH Zurich, Zurich, Switzerland
[2]Department of Civil and Environmental Engineering, Stanford University, CA, USA

Correspondence: Lukas Bodenmann (bodenmann@ibk.baug.ethz.ch)

Abstract. Ground-motion correlation models play a crucial role in regional seismic risk modelling of spatially distributed built infrastructure. Such models predict the correlation between ground-motion amplitudes at pairs of sites, typically as a function of their spatial proximity. Data from physics-based simulators and event-to-event variability in empirically derived model parameters suggest that spatial correlation is additionally affected by path and site effects. Yet, identifying these effects has been difficult due to scarce data, and a lack of modelling and assessment approaches to consider more complex correlation predictions. To address this gap, we propose a novel correlation model that accounts for path and site effects via a modified functional form. To quantify the estimation uncertainty, we perform Bayesian inference for model parameter estimation. The derived model outperforms traditional isotropic models in terms of the predictive accuracy for training and testing data sets. We show that the previously found event-to-event variability in model parameters may be explained by the lack of accounting for path and site effects. Finally, we examine implications of the newly proposed model for regional seismic risk simulations.

1 Introduction

Earthquakes can cause widespread damage to the built environment exposing its users to severe and potentially long-lasting societal stress. Analyzing earthquake-induced consequences is key to enhancing efficient and targeted seismic risk management strategies. Empirical ground-motion models (GMMs) are widely used for the prediction of earthquake-induced ground-motion intensity measures (IMs) at individual sites. The assessment of consequences to spatially distributed systems, such as the residential building stock of an urban area or its road network, additionally requires spatial correlation models to characterize the dependency among IMs at different sites (Wesson and Perkins, 2001; Lee and Kiremidjian, 2007).

A predictive spatial correlation model consists of a *functional form*, one or several *dependent variables* and the *model parameters*. Early studies, such as Boore et al. (2003), as well as more recent studies, such as Schiappapietra and Douglas (2020), use an isotropic model where the correlation among sites decays exponentially (the *functional form*) with increasing Euclidean distance between sites (the *dependent variable*). Observations of pairs of IMs from past earthquakes are used to calibrate these models. The isotropic assumption allows grouping of station pairs with similar distance for the estimation of *model parameters* via geo-statistical curve-fitting techniques.

To alleviate the scarcity of data, some researchers pooled data from multiple earthquakes, and assumed that the same cor-
relation model parameters apply to different events (e.g., Goda and Atkinson, 2010; Esposito and Iervolino, 2011). Other
studies estimated separate models for individual earthquakes and reported that the corresponding parameters vary from one
earthquake to another (e.g., Jayaram and Baker, 2009; Goda, 2011). Sokolov et al. (2012) mention that this event-to-event
variability may be caused by site effects (i.e. stronger correlation amongst sites with similar geological conditions) and path
effects (i.e. stronger correlation amongst sites with similar wave propagation paths). While the importance of site and path ef-
fects became apparent in data from physics-based ground-motion simulators (Chen and Baker, 2019), the findings with respect
to recorded ground-motion data differ amongst previous studies. Early work of Jayaram and Baker (2009) and Sokolov et al.
(2012) found that, for example, the heterogeneity of soil conditions may influence spatial correlations. Yet, more recent studies
could not find such evidence and suggested to account for event-to-event variability via random variables for model parameters
in regional seismic risk analyses (Heresi and Miranda, 2019) or to derive region-specific correlation models (Schiappapietra
and Douglas, 2020).

Identifying explanatory factors by estimating correlation models from data of individual earthquakes is challenging. First,
the comparison of a single model parameter estimate per event with a single metric describing a certain aspect of the region the
event was recorded in (such as the heterogeneity of soil conditions) suffers from scarcity of data. Estimation of event-specific
correlation model parameters requires data from particularly well-recorded events, of which there are only a few. Second,
the use of an isotropic model and the condensation to a single parameter estimate per event may hide path and site effects
that are present within the event data. Third, the estimated model parameters are subject to varying degrees of estimation
uncertainty, because the underlying data sets stem from earthquakes that were recorded by a different number and layout of
seismic network stations. Schiappapietra and Douglas (2021) and Baker and Chen (2020) aimed to quantify this uncertainty
via simulating data from an assumed "true" model and comparing the latter to the model estimated from the simulated data
using different estimation techniques. Both studies, however, used the same isotropic model with an exponential functional
form that has only one model parameter, and Baker and Chen reported difficulties in extending the proposed method to models
with more than one parameter.

This study explores novel correlation models that, in addition to spatial proximity, also account for path and site effects.
In contrast to previous studies, we do so by modifying and extending the *functional form* and the *dependent variables* of the
correlation models. The increased complexity of these models calls for a consistent quantification of the inherent estimation
uncertainty, thus complicating the use of conventional geo-statistical curve-fitting techniques. To address this, we use Bayesian
inference to estimate the *model parameters*. While Bayesian inference has been proposed for GMMs in the past (e.g., Moss
and Der Kiureghian, 2006; Stafford, 2019), it has not been applied to study spatial ground-motion correlation.

We present the proposed correlation model in Sect. 2. Section 3 introduces the Bayesian inference scheme used to estimate
the model parameters from the PEER NGA-West2 data set (Ancheta et al., 2014). Using the same data set, Sect. 4 compares
the proposed correlation models to event-specific, isotropic models by employing a novel metric to quantify the predictive
accuracy, and Sect. 5 compares model performance on test data from the 2019 Ridgecrest, California, earthquake sequence

(Rekoske et al., 2020). Finally, Sect. 6 examines implications of the novel correlation model for regional seismic risk simulation studies.

2 Spatial correlation models for ground-motion amplitudes

This study on spatial correlation models builds on empirically derived GMMs that predict a ground-motion IM at site i induced by an earthquake rupture k as

$$\ln IM_{ki} = \mu_{\ln IM}(rup_k, site_i) + \delta B_k + \delta W_{ki} \,, \tag{1}$$

where $\mu_{\ln IM}(\cdot)$ is the predicted mean $\ln IM$ value as a function of rupture (rup) and site ($site$) characteristics (Baker et al., 2021). Amongst others, these typically include earthquake magnitude, rupture mechanism, source-to-site distance and site-specific geological information. The between-event and within-event residuals, δB_k and δW_{ki}, are assumed to be independent, normally distributed variables with standard deviations τ and ϕ, respectively. Empirical GMMs provide the mean function $\mu_{\ln IM}(\cdot)$, as well as the standard deviations τ and ϕ. For a specific event, the between-event residual denotes a common deviation from the predicted mean that is constant for all sites, whereas the within-event residuals vary in space. For this correlation study, we scale the within-event residual by its standard deviation ϕ, and denote the scaled within-event residual as $Z_{ki} = \delta W_{ki}/\phi$.

The joint distribution of ground-motion IMs at n spatially distributed sites $\mathbf{IM}_k = (IM_{k1}, \ldots, IM_{kn})^\top$ requires characterizing the dependence of the corresponding within-event residuals $\mathbf{Z}_k = (Z_{k1}, \ldots, Z_{kn})^\top$. The joint distribution of the latter is assumed to be multivariate normal (Jayaram and Baker, 2008), e.g. $p(\mathbf{z}_k) = \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$. Note that because each marginal Z_{ki} follows a standard normal distribution, the covariance matrix, $\mathbf{\Sigma}$, is identical to the correlation matrix. To compute the entries of this matrix we employ a model $\rho(\cdot)$ that predicts the correlation between two sites i and j given some *dependent variables* $\mathbf{x}_i$ and $\mathbf{x}_j$ as

$$[\mathbf{\Sigma}]_{ij} = \rho_{\mathrm{M}}(\mathbf{x}_i, \mathbf{x}_j; \boldsymbol{\psi}_{\mathrm{M}}) \,, \tag{2}$$

where subscript M indicates the chosen *functional form* of the model and $\boldsymbol{\psi}_{\mathrm{M}}$ denotes associated *parameters*. The models, which we introduce in the following sections, are defined for a distance (or dissimilarity) metric d between two sites. We often denote the correlation model as $\rho_{\mathrm{M}}(d; \boldsymbol{\psi}_{\mathrm{M}})$.

2.1 Isotropic model based on Euclidean distance metric

A natural first assumption is that the correlation between sites decreases as the Euclidean distance between them increases. If the (projected) Cartesian coordinates of two sites i and j are denoted by $\mathbf{s}_i$ and $\mathbf{s}_j$, then the Euclidean distance between them is defined as $d_{\mathrm{E},ij} = \|\mathbf{s}_i - \mathbf{s}_j\|_2$, as illustrated in Fig. 1a. Because the correlation solely depends on the distance between sites, the resulting correlation model is isotropic. Following Goda and Hong (2008), and more recently Heresi and Miranda (2019), we use a γ-Exponential function (Rasmussen and Williams, 2006, chapter 4) to describe the decay of correlation as a function

of Euclidean distance

$$\rho_{\mathrm{E}}(d_{\mathrm{E}}; \boldsymbol{\psi}_{\mathrm{E}}) = \exp\left(-(d_{\mathrm{E}}/\ell_{\mathrm{E}})^{\gamma_{\mathrm{E}}}\right) , \tag{3}$$

where $\gamma_{\mathrm{E}} \in (0, 2)$ denotes the exponent and $\ell_{\mathrm{E}} \in \mathbb{R}^{+}$ is the lengthscale in km. Both parameters are summarized in the vector $\boldsymbol{\psi}_{\mathrm{E}} = (\ell_{\mathrm{E}}, \gamma_{\mathrm{E}})$, where the subscript E is used to simplify notation for the correlation model comparisons that follow below. Figure 1b illustrates the correlation function for different parameter combinations. If the exponent is one, Eq. (3) simplifies to the exponential function used in many previous ground-motion correlation studies (e.g., Esposito and Iervolino, 2011). As the exponent drops below one, the correlation for distances shorter than the lengthscale are weaker compared to the exponential function, whereas correlations for longer distances are stronger.

2.2 Accounting for path effects using an angular distance metric

Besides the Euclidean distance between two sites, their correlation may also depend on their position relative to the earthquake rupture (due to arriving waves potentially travelling similar propagation paths). In this study we use the epicentral azimuth θ to characterize this relative position and assume that correlation between sites decreases as the difference in their azimuths increases. This difference in epicentral azimuths is herein called the angular distance d_{A} and illustrated in Fig. 1a. The angular distance takes values from 0° to 180°, where 180° indicates two sites that are on opposite sides of the epicenter. To account for path effects, we use the following correlation function that was proposed by Padonou and Roustant (2016) for Gaussian Processes on circular domains

$$\rho_{\mathrm{A}}(d_{\mathrm{A}}; \ell_{\mathrm{A}}) = (1 + d_{\mathrm{A}}/\ell_{\mathrm{A}})(1 - d_{\mathrm{A}}/180)^{180/\ell_{\mathrm{A}}} , \tag{4}$$

where $\ell_{\mathrm{A}} \in (0, 45)$ is the lengthscale in degrees. Figure 1c provides a visual illustration of Eq. (4) for different lengthscales. Model A, as introduced above, assigns strong correlation for sites with similar epicentral azimuths regardless of how close the sites are to each other. To account for spatial proximity in addition to path effects, we introduce model EA:

$$\rho_{\mathrm{EA}}(d_{\mathrm{E}}, d_{\mathrm{A}}; \boldsymbol{\psi}_{\mathrm{EA}}) = \rho_{\mathrm{E}}(d_{\mathrm{E}}) \cdot \rho_{\mathrm{A}}(d_{\mathrm{A}}) , \tag{5}$$

where $\boldsymbol{\psi}_{\mathrm{EA}} = (\gamma_{\mathrm{E}}, \ell_{\mathrm{E}}, \ell_{\mathrm{A}})$ collects all parameters of the individual functions E and A. The multiplicative structure ensures that strong correlations are only present if two sites have similar epicentral azimuths and are close to each other. Figure 1g illustrates the resulting correlation coefficient from a reference site to all other sites in a fictitious region. The comparison with the isotropic model E (Fig. 1f) reveals how model EA assigns weaker correlations to sites with differing paths from the epicenter.

Path effects in the context of correlation models for ergodic GMMs imply stronger correlation between sites that share a similar wave propagation path. This is different from non-ergodic models where one tries to identify systematic and repeatable path effects in areas where multiple events have been recorded by the same seismic network. In the latter context, correlation functions are used to establish probabilistic links between stations in the seismic network in order to estimate the systematic path effects. While these functions are typically defined for Euclidean distance and have a similar functional form as Eq. (3),

Kuehn and Abrahamson (2020) and Liu et al. (2022b) recently proposed to vary the Euclidean lengthscale, ℓ_E, as a function

120 of the epicentral distance. This also induces path effects that vary spatially in a more complex manner than only Euclidean distance. Yet, a quantitative comparison with this approach is non-trivial because of the fundamental differences between ergodic and non-ergodic GMMs.
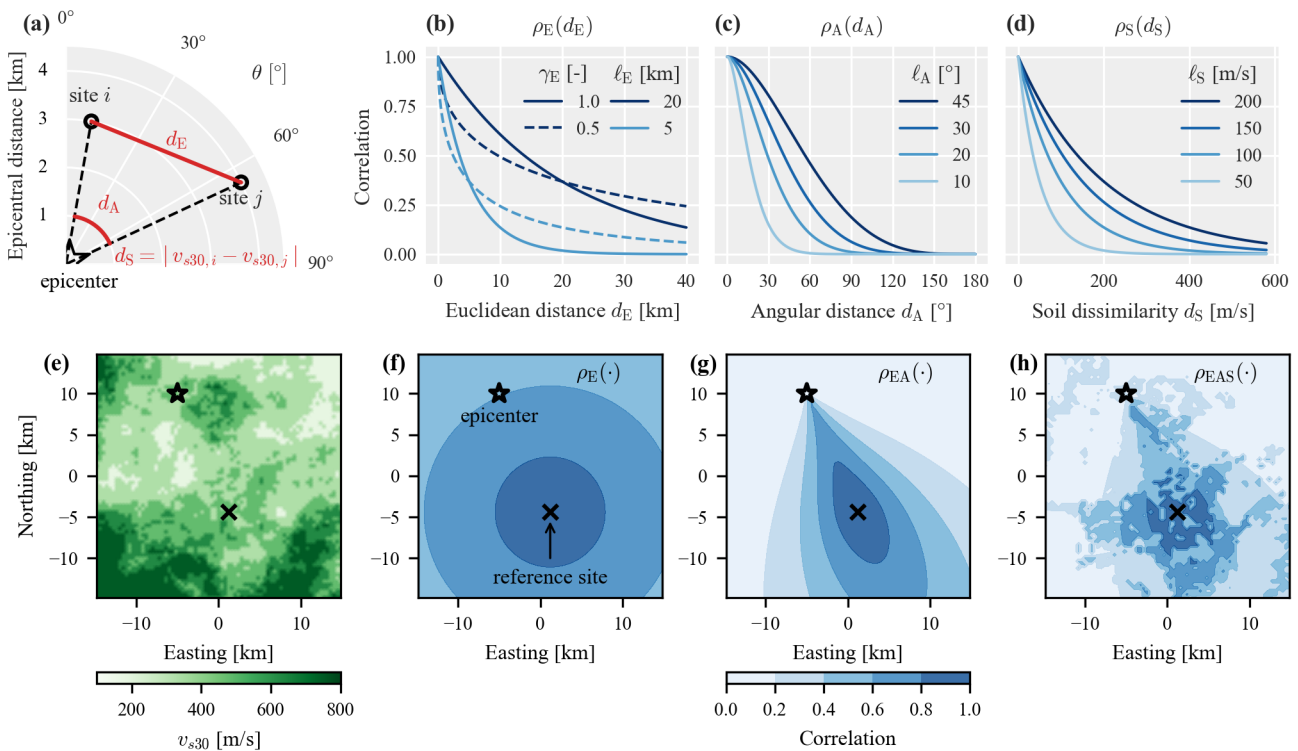
Figure 1. Illustration of the proposed spatial correlation models: Distance and dissimilarity metrics (dependent variables) (a), and correlation coefficients as a function of the corresponding dependent variable for different parameters of individual models E, A and S (b to d). Soil conditions in a fictitious 15x15km region (e) that are used to illustrate correlation coefficients with respect to the indicated reference site for the combined correlation models E, EA and EAS (f to h), where the model parameters are set to their corresponding prior mean values stated in Table 1.

2.3 Accounting for site effects using a soil dissimilarity metric

We account for site effects via measuring dissimilarities in local soil conditions following the premise that sites with similar

125 soil conditions have stronger correlations. We use v_{s30}, the 30m time-averaged shear-wave velocity, as a proxy for the soil conditions and the distance metric d_S is the absolute difference in the two sites' v_{s30} values. We choose an exponential form

of the correlation function

$$\rho_S(d_S; \ell_S) = \exp(-d_S/\ell_S) , \tag{6}$$

where $\ell_S \in \mathbb{R}^+$ is the lengthscale in m/s. The choice of using v_{s30} as a proxy for soil conditions reflects its use in the GMM and its availability in the data sets and case-study regions employed in this study. Other information, such as the depth to bedrock or simply topographic slope (Kotha et al., 2020), could be used in a similar fashion. To simultaneously account for spatial proximity, path, and site effects we introduce correlation model EAS:

$$\rho_{EAS}(d_E, d_A, d_S; \boldsymbol{\psi}_{EAS}) = \rho_E(d_E) \cdot (w\rho_A(d_A) + (1-w)\rho_S(d_S)) , \tag{7}$$

where $w \in (0, 1)$ is a weight parameter and $\boldsymbol{\psi}_{EAS} = (\gamma_E, \ell_E, \ell_A, \ell_S, w)$ collects all parameters. To illustrate model EAS, Fig. 1h plots the predicted correlation coefficient with respect to the indicated reference site using soil conditions as shown in Fig. 1e. The model still predicts path effects, but also higher correlation for sites with similar soil conditions as the reference site.

For model EAS, we explored several combinations of the individual models E, A and S. Compared to the model defined in Eq. (7), a decrease in predictive performance was observed for the case where the individual components are combined as a weighted sum, whereas the decrease was less pronounced for a purely multiplicative structure. We also tried a model with two separate isotropic components E1 and E2 that are multiplied with components A and S, respectively. While this comes at the cost of two additional parameters, this model had a slightly better predictive performance. However, in favor of reduced model complexity, we decided to proceed with the model specified above.

3 Bayesian parameter estimation

We follow a Bayesian approach to estimate the parameters of the correlation models $M \in \{E, EA, EAS\}$, presented in the previous section. Given data $\mathcal{D}$, we aim to derive the posterior distribution of the correlation model parameters $\boldsymbol{\psi}_M$ using Bayes' theorem

$$p(\boldsymbol{\psi}_M|\mathcal{D}) \propto p(\mathcal{D}|\boldsymbol{\psi}_M)p(\boldsymbol{\psi}_M) , \tag{8}$$

where $p(\mathcal{D}|\boldsymbol{\psi}_M)$ denotes the likelihood of jointly observing the data $\mathcal{D}$ conditional on a correlation model M with a specific parameter set $\boldsymbol{\psi}_M$, and $p(\boldsymbol{\psi}_M)$ is the prior distribution of the parameters. The individual components of Eq. (8) are introduced in detail in the following sections.

3.1 Data

We consider ground-motion IM data from the NGA-West2 database (Ancheta et al., 2014). The IMs of interest are elastic, 5%-damped spectral accelerations $Sa(T)$ over a range of periods T, where we consider the median spectral amplitude over all horizontal orientations (SARotD50). To compute within-event residuals we use the GMM of Chiou and Youngs (2014) and mixed-effects regression. We restrict the database to consider only ground motions with closest distance to rupture below 300

km, v_{s30} values between 180 and 760 m/s, and a maximum usable period within the period range of interest. We then consider all earthquake events with more than 40 stations that satisfy the above criteria. Besides the computed residuals, $\mathbf{z}_k$, we extract the geo-coordinates $\mathbf{s}_{ki}$, the epicentral azimuth θ_{ki}, and $v_{s30,ki}$ from all records, which we collectively denote as input vector $\mathbf{x}_{ki}$. The data from event k is denoted as $\mathcal{D}_k = (\mathbf{X}_k, \mathbf{z}_k)$, where the inputs and residuals of the stations that recorded event k

160 are summarized in matrix $\mathbf{X}_k$ and vector $\mathbf{z}_k$, respectively.

The data set obtained by pooling data from all n_k considered events is denoted as $\mathcal{D}_{tot} = \{\mathcal{D}_k \mid k = 1, \ldots, n_k\}$. For $Sa(1s)$, the pooled data set consists of 13,342 records from 128 events. Fig. 2 shows the number of station pairs in bins of Euclidean distance combined with angular distance (a), and with soil dissimilarity (b), respectively. As can be seen in Fig. 2a there are only few station pairs available at short Euclidean distances (smaller than 5km) and angular distances greater than 60°, mainly

165 because this combination requires data points that have been recorded at short epicentral distances, which are in general scarce. The number of station pairs available at different combinations of Euclidean distance and soil dissimilarity, shown in Fig. 2b, are more evenly distributed, with relatively few data points from close-by stations with strongly differing soil conditions (e.g. $d_S > 350$ m/s).
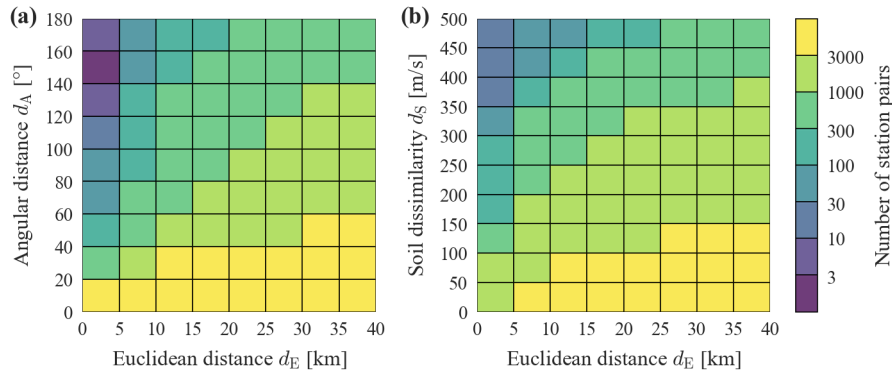
Figure 2. For the pooled training data set: Number of station pairs in joint bins of Euclidean and angular distance in (a), and number of station pairs in joint bins of Euclidean distance and soil dissimilarity in (b). Note the logarithmic color scale.

3.2 Likelihood

170 For event k, the likelihood, $p(\mathcal{D}_k|\boldsymbol{\psi}_{\mathrm{M}})$, denotes the joint probability that correlation model M with parameters $\boldsymbol{\psi}_{\mathrm{M}}$ and inputs $\mathbf{X}_k$ assigns to scaled within-event residuals $\mathbf{z}_k$ computed from the recorded ground-motion IMs as discussed above. This distribution is multivariate normal, so $p(\mathcal{D}_k|\boldsymbol{\psi}_{\mathrm{M}}) = \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$. For the pooled data set $\mathcal{D}_{tot}$, we consider the residuals from distinct events to be independent, so the likelihood factorizes: $p(\mathcal{D}_{tot}|\boldsymbol{\psi}_{\mathrm{M}}) = \prod_k p(\mathcal{D}_k|\boldsymbol{\psi}_{\mathrm{M}})$. The joint probability of interest is thus analytically tractable. The correlation model and its parameters $\boldsymbol{\psi}_{\mathrm{M}}$ define the entries of the correlation matrix $\boldsymbol{\Sigma}$.

Table 1. Prior distributions for the parameters ψ and corresponding means and the 5% and 95% quantiles ($q_{0.05}$ and $q_{0.95}$).

Parameter	Domain	Distribution	Mean	$q_{0.05}$	$q_{0.95}$
γ_E [-]	(0,2)	$p(\gamma_E/2) = \mathrm{Beta}(\alpha = 2, \beta = 2)$	1.0	0.3	1.7
ℓ_E [km]	$\mathbb{R}^+$	$p(\ell_E) = \mathrm{InvGamma}(\alpha = 2, \beta = 30)$	30.0	6.3	84.0
ℓ_A [°]	(0,45)	$p(180/\ell_A - 4) = \mathrm{Gamma}(\alpha = 2, \beta = 0.25)$	18.2	7.8	33.3
ℓ_S [m/s]	$\mathbb{R}^+$	$p(\ell_S) = \mathrm{InvGamma}(\alpha = 2, \beta = 100)$	100.0	21.0	284.6
w [-]	(0,1)	$p(w) = \mathrm{Beta}(\alpha = 2, \beta = 2)$	0.5	0.1	0.9

3.3 Prior distributions

We chose weakly informative prior distributions for the parameters based on guidance provided in Kuehn and Stafford (2021) and Liu et al. (2022b), and assume that their joint distribution, $p(\psi_M)$ in Eq. (8), is factorizing. The prior distributions and the corresponding prior mean, as well as the 5% and 95% quantiles are stated in Table 1. Note that some priors are defined on a scaled version of the parameter to account for its corresponding domain. For instance the exponent, γ_E, is defined for a range from zero to two, which means to get samples of this parameter we first sample from a Beta distribution, defined from zero to one, and multiply the obtained samples by two.

3.4 Posterior distribution

The posterior distribution, $p(\psi_M|\mathcal{D})$ in Eq. (8), is not analytically tractable, thus we use Markov Chain Monte Carlo (MCMC) sampling for Bayesian inference (Neal, 1993). Specifically, we employ the No-U-Turn Sampler (Hoffman and Gelman, 2014) as implemented in the software library NumPyro (Phan et al., 2019). Instead of a closed-form posterior distribution, MCMC provides n_r sampled parameter sets $\psi_{M,r}$ from the posterior.

In the following we discuss the estimated parameters by first focusing on isotropic models E that are derived from data $\mathcal{D}_k$ of a single event k, so-called event-specific models. Then we expand the discussion to the case where parameters are estimated from the pooled training data $\mathcal{D}_{tot}$, so-called pooled models, where we also include the combined models EA and EAS. Both upcoming paragraphs focus on the spectral acceleration at a period of one second, $Sa(1s)$.

3.4.1 Event-specific models

Event-specific correlation models have parameters estimated from data of an individual event. Figure 3 illustrates the joint posterior distributions of the parameters ψ_E estimated separately for three events with increasing number of records. The top row shows the individual samples $\psi_{E,r}$ obtained via MCMC and an estimated kernel density for illustrative purposes. The bottom row shows the correlation model evaluated using all sampled parameter sets $\psi_{E,r}$.

The results shown in Fig. 3 illustrate slight differences in the three event-specific models but they also highlight the uncertainty involved in estimating correlation model parameters from data of individual earthquakes. As pointed out in the introduction, this estimation uncertainty complicates the identification of factors that could explain the variability in event-specific

correlation models. In Sect. 4 we will compare the predictive accuracy of such event-specific models to the predictive accuracy

200   of pooled models as introduced in the following.
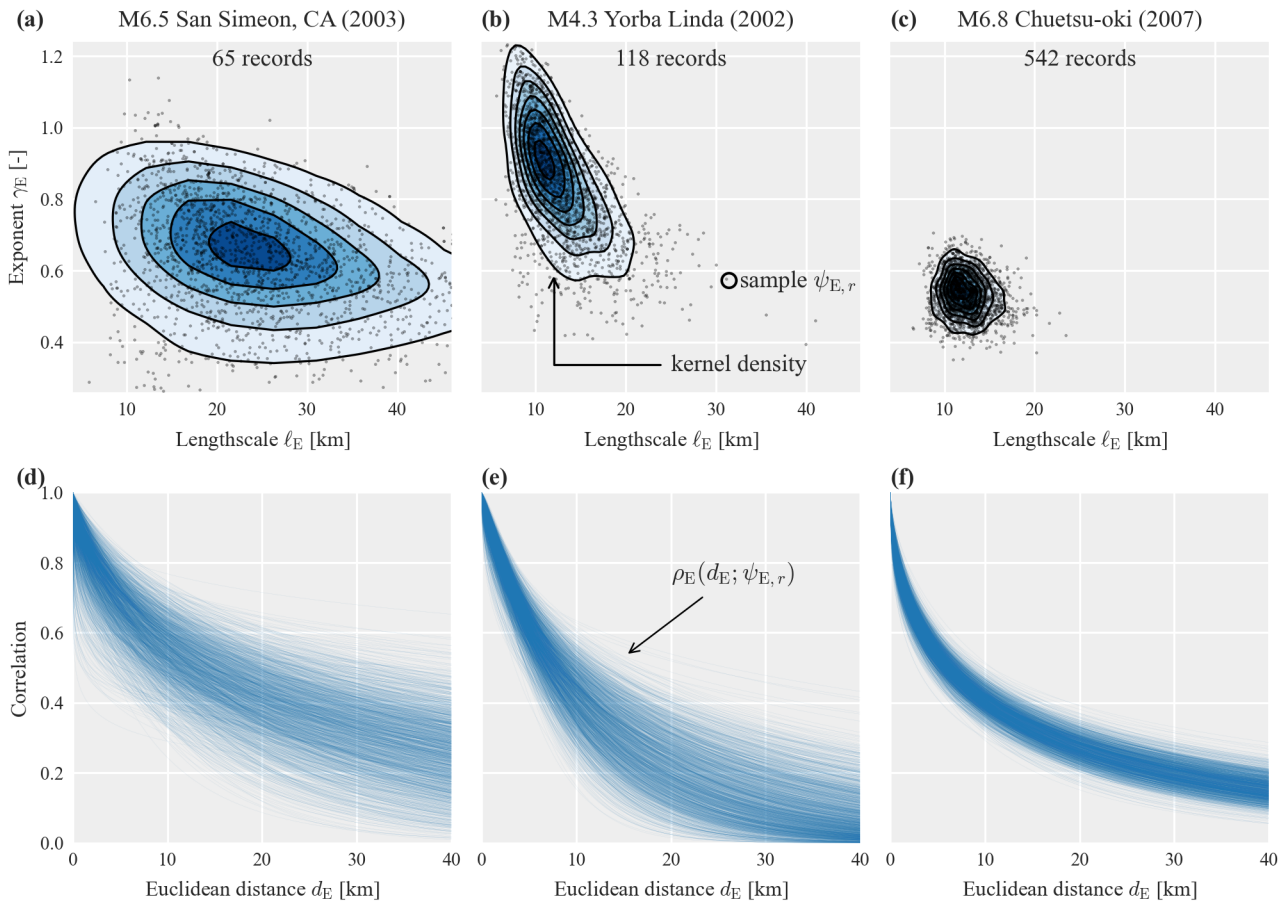
Figure 3. Isotropic model E estimated separately for three events with increasing number of records: (a-c) Posterior distributions of the parameters, and (d-f) Predicted correlation coefficient as a function of Euclidean distance for all posterior parameter samples.

3.4.2   Pooled models

Pooled models are derived by combining data from multiple individual earthquakes. In contrast to event-specific models, we use the same parameters to describe correlations of data from all events.

Table 2 summarizes the posterior parameters of correlation models E, EA and EAS, all inferred from the pooled training data

205    $\mathcal{D}_{tot}$. First, we note that the angular lengthscale, ℓ_A, of models EA and EAS is around 20°. As can be seen from Fig. 1c, this means that two sites whose epicentral azimuths differ by more than 90° are essentially uncorrelated. Second, the lengthscale applied to the Euclidean distance, ℓ_E, is longer for models EAS and EA compared to the isotropic model E, while the exponent

Table 2. Parameters for the different correlation models estimated from the pooled training data set. Stated quantities are the mean, and the 5% and 95% quantiles from the posterior samples.

Model		γ_e [-]	ℓ_e [km]	ℓ_a [°]	ℓ_s [m/s]	w [-]
E	Mean	0.40	16.0	-	-	-
	5%, 95% quantiles	0.38, 0.42	14.7, 17.4	-	-	-
EA	Mean	0.35	21.3	23.5	-	-
	5%, 95% quantiles	0.33, 0.36	19.2, 23.5	20.8, 26.7	-	-
EAS	Mean	0.41	29.8	20.4	169.2	0.70
	5%, 95% quantiles	0.38, 0.43	27.0, 32.8	17.8, 23.6	106.1, 250.8	0.62, 0.77

γ_{E} is similar for all three models. Thus, sites with similar azimuthal difference and located on similar soil are correlated over longer Euclidean distances.

For models EAS and E, Fig. 4 illustrates the predicted correlation as a function of Euclidean distance, d_{E}, for the three soil dissimilarities, d_{S}, of 10, 100 and 400 m/s, and at three increasing angular distances, d_{A}, of 5, 30 and 60°. The solid line indicates the function evaluated with the mean posterior parameters and the shaded area indicates the interval between the 95% and 5% quantile of all sampled functions from the posterior. The isotropic model E depends solely on Euclidean distance, so the predicted correlation is identical in all three panels (a-c). The comparison of this pooled model E to the event-specific isotropic models shown in Fig. 3 reveals the reduction in estimation uncertainty obtained by pooling data from multiple events. For model EAS, on the other hand, we observe increased uncertainty for larger angular distances and short Euclidean distances, which reflects the low amount of data available at such combinations (see Fig. 2). By comparing model EAS in panels (a) and (b), we observe that increasing the angular distance from 5° to 30° has roughly the same effect as increasing the soil dissimilarity from 10 m/s to 400 m/s. For the plotted range of Euclidean distance, the isotropic model E is similar to model EAS at small angular distances. This is because most data at short distances have small azimuthal differences.

4 In-sample model performance

We next evaluate the three models in terms of their predictive accuracy on test data from either an individual event or multiple events. Given the posterior parameters of model M inferred from training data, $\mathcal{D}_{train}$, we compute the Logarithmic Posterior Predictive Density (LPPD) of test data, $\mathcal{D}_{test}$, as

$$\mathrm{LPPD}_{\mathrm{M}}(\mathcal{D}_{test}; \mathcal{D}_{train}) = \ln \int p(\mathcal{D}_{test}|\boldsymbol{\psi}_{\mathrm{M}}) p(\boldsymbol{\psi}_{\mathrm{M}}|\mathcal{D}_{train}) \, \mathrm{d}\boldsymbol{\psi}_{\mathrm{M}} \approx \ln \frac{1}{n_r} \sum_{r=1}^{n_r} p(\mathcal{D}_{test}|\boldsymbol{\psi}_{\mathrm{M},r}) \,, \tag{9}$$

where $\boldsymbol{\psi}_{\mathrm{M},r}$ is the r-th sample from the posterior $p(\boldsymbol{\psi}_{\mathrm{M}}|\mathcal{D}_{train})$ obtained via MCMC (Gelman et al., 2014). A higher value of this metric indicates a model that is more compatible with the data, in the sense of predicting higher probabilities of observing the given test data. By marginalizing over the posterior distribution of the parameters, the LPPD takes into account the uncertainty associated with parameter estimation. To compare the models we use the relative difference between the LPPD
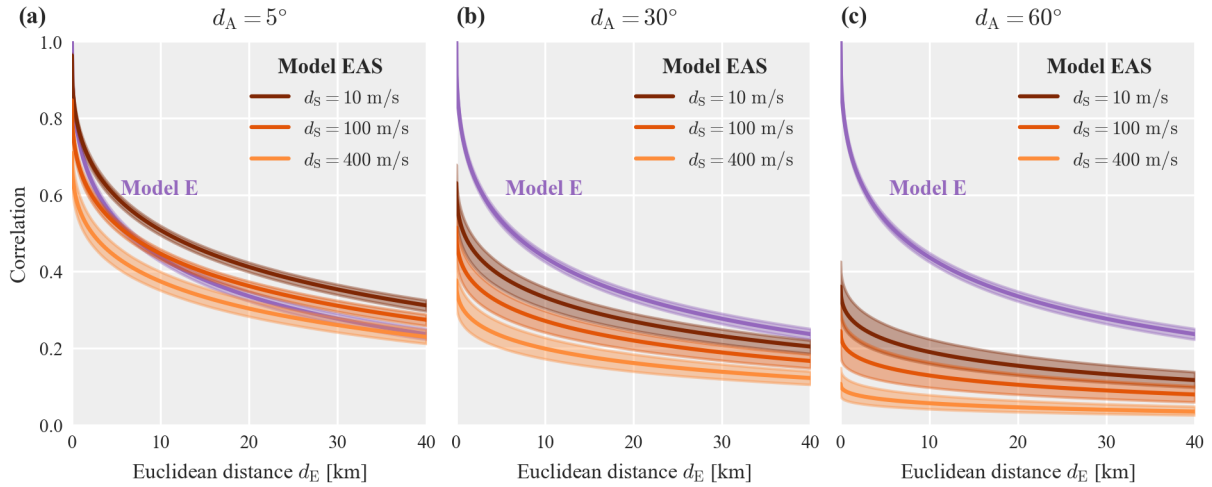
Figure 4. Posterior correlation models EAS and E for $Sa(1s)$ as a function of Euclidean distance and soil dissimilarity plotted at three angular distances: (a) 5°, (b) 30° and (c) 60°. The shaded area indicates the interval between the 95% and 5% quantile of all sampled functions from the posterior.

of a model M and the LPPD of a certain baseline model BL:

$$\text{Rel. difference} = \frac{\text{LPPD}_{\text{BL}} - \text{LPPD}_{\text{M}}}{\text{LPPD}_{\text{BL}}} \ . \tag{10}$$

Note that the LPPD is negative (log-scale), thus a positive relative difference indicates that model M has higher predictive accuracy (higher LPPD) than the baseline model BL.

This section focuses on in-sample performance, which means that the test data is a subset of the training data, i.e., $\mathcal{D}_{test} \subseteq \mathcal{D}_{train}$. The training data stems from the NGA-West2 database as described in Sect. 3.1. Specifically, we first focus on $Sa(1s)$ and compare the performance of pooled models E and EAS (estimated in Sect. 3.4.2) to the performance of event-specific isotropic models E (estimated in Sect. 3.4.1) on data $\mathcal{D}_k \subset \mathcal{D}_{tot}$ from each event k. Then we analyze the performance of the pooled models on the entire training data set $\mathcal{D}_{tot}$ to examine the relative importance of site and path effects in correlation of $Sa(T)$ at different periods T.

4.1 Comparing event-specific and pooled models

As stated in the introduction, previous studies found that the parameters of an isotropic correlation model estimated from data of different events vary from earthquake to earthquake. To examine this event-to-event variability, we compute for each event k: (i) the predictive accuracy $\text{LPPD}_{\text{M}}(\mathcal{D}_k; \mathcal{D}_{tot})$ of pooled models $\text{M} \in \{\text{E}, \text{EAS}\}$, and (ii) the predictive accuracy of the corresponding isotropic event-specific model $\text{LPPD}_{\text{E}}(\mathcal{D}_k; \mathcal{D}_k)$.

To illustrate the aforementioned LPPD metrics, we first use 125 records from the Hector-Mine earthquake in Fig. 5. The frequency histograms indicate the log-likelihoods computed for individual samples from the posterior distribution of the model

parameters, $\ln p(\mathcal{D}_k | \boldsymbol{\psi}_{\mathrm{M},r})$. The X values indicate the final LPPD metric which is computed from the histogram values as specified in Eq. (9).

We see that the event-specific model E has the largest variance in log-likelihood values, as the model parameters are more
250 uncertain due to limited data and some sampled parameter values give low probability of observing the data (i.e., low log-likelihoods). For most realizations, however, the event-specific model E outperforms the pooled model E for the given event, as would be expected. By computing the LPPD metric using Eq. (9) we marginalize over the posterior parameters and the resulting LPPDs of both models are similar, with the event-specific model performing slightly better. We also see that the pooled model EAS has a higher predictive accuracy than both the pooled and event-specific isotropic models E.
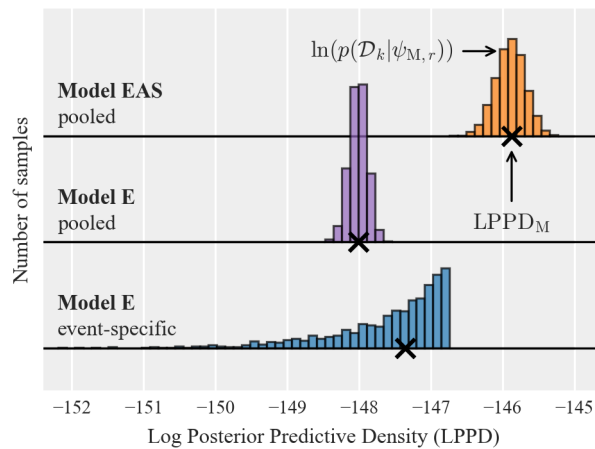
Figure 5. Log Posterior Predictive Density (LPPD) for data from the Hector-Mine earthquake event of models E and EAS estimated from the pooled data set, as well as model E with parameters estimated only from data of that event. The histograms show the log-likelihood of the data conditional on samples from the posterior parameters, $\ln p(\mathcal{D}_k | \boldsymbol{\psi}_{\mathrm{M},r})$.

255 Figure 6 extends the above comparison to all events in the training data by comparing the LPPDs of pooled model E (panel a) and pooled model EAS (panel b) for each event to the LPPD of the corresponding event-specific isotropic model E using the relative difference metric specified in Eq. (10). The event-specific model E serves as a baseline in computing the relative difference. Thus a positive relative difference indicates that the pooled model predicts the data from that event with higher accuracy than the corresponding event-specific isotropic model.

260 Figure 6a reveals that the pooled model E has worse predictive accuracy than the event-specific model for most events, with the exception of some events with fewer than 200 records. The exceptions are explained by the increased estimation uncertainty when using sparse data from solely one event, as discussed above. The better predictive accuracy of event-specific models points to event-to-event variability in model parameters when using isotropic correlation models that are only based on Euclidean distance.

On the other hand, Fig. 6b shows that the predictive accuracy of the pooled model EAS is, for most events, higher than the event-specific isotropic model. This is despite the increased number of parameters (and associated estimation uncertainty) and highlights the benefit of accounting for path and site effects in correlation models. It also indicates that the event-to-event variability in correlation model parameters is at least partially explained by the lack of accounting for such effects. Finally, we note that the earthquake magnitude does not seem to affect the relative comparison in LPPDs, neither for model E nor for model EAS.
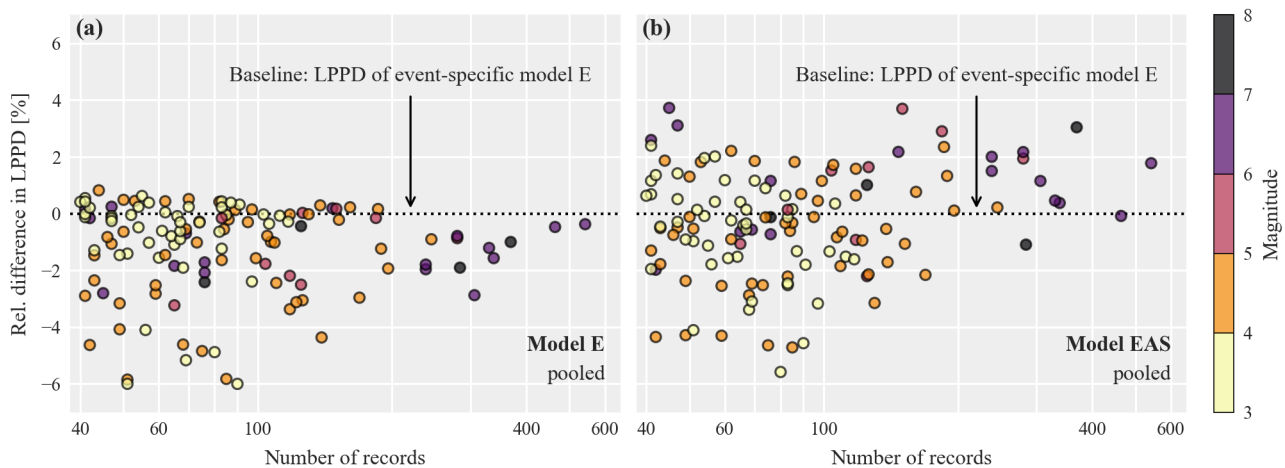
Figure 6. Relative difference in Log Posterior Predictive Density (LPPD) of pooled models E (a) and EAS (b) to the LPPD of event-specific models E as a function of the available number of records and magnitude. Each point represents one event in the considered data set, and the coloring indicates the event's magnitude. The relative difference is computed using Eq. (10) and positive values indicate higher predictive accuracy (higher LPPD) of the pooled model for this event compared to the corresponding event-specific model E.

4.2 Pooled model performance for Sa at other periods

Whereas the previous results considered within-event residuals of $Sa(1s)$, this section expands the discussion to different periods T. We again use the NGA-West2 data set to estimate parameters of pooled models for residuals at nine periods between 0.01s and 6s, and evaluate the corresponding LPPD for the entire training data set, $\mathrm{LPPD_M}(\mathcal{D}_{tot}; \mathcal{D}_{tot})$. To enable a comparison across periods we additionally compute the predictive accuracy of an independent model (i.e., no correlation between within-event residuals). For the training data set with n records the latter is computed as $\mathrm{LPPD_I} = \sum_i^n \ln p(z_i)$ where each individual scaled within-event residual z_i follows a standard normal distribution.

Figure 7a compares the resulting LPPD values in terms of their relative difference to the baseline LPPD of the independent model using Eq. (10). For the considered periods, all three models have increased predictive accuracy compared to the independent model. The major part of the benefit in LPPD stems from the isotropic model E, although accounting for path and site effects (models EA and EAS) leads to further increases in LPPD. Figure 7b focuses on the relative difference of model EA and

EAS compared to the baseline of the isotropic model E. Interestingly, for $T \leq 0.3$s the LPPD of model EAS is very similar to the LPPD of model EA. This indicates that, for short periods, additionally accounting for site effects (as measured with model EAS) adds only a minor benefit compared to model EA which only accounts for path effects. For longer periods, however, the
285 site effects become more important.
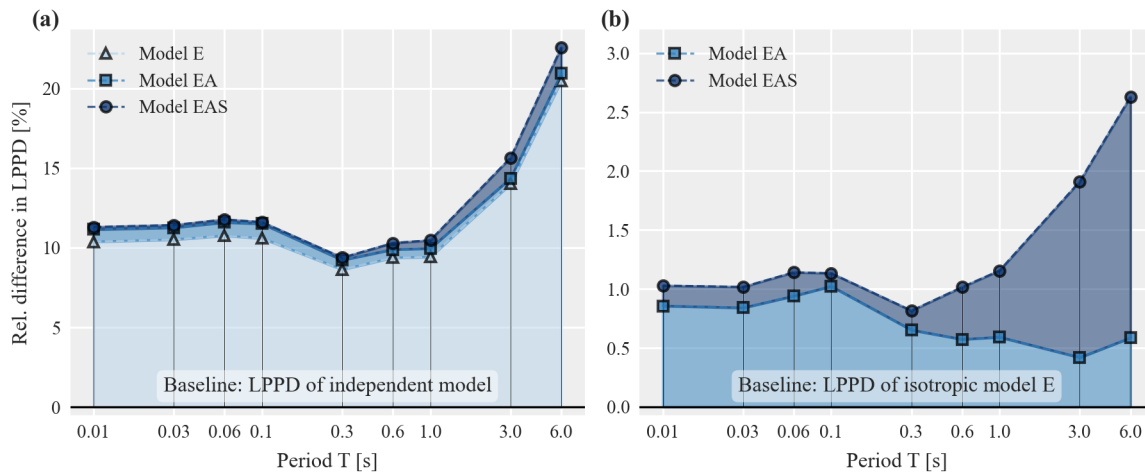
Figure 7. For within-event residuals of $Sa(T)$ at different periods T from the NGA-West2 data set: Relative difference (Eq. 10) in Log Posterior Predictive Density (LPPD) of pooled models E, EA and EAS to (a) the LPPD of the independent model, and (b) the LPPD of pooled model E.

5 Out-of-sample model performance

We use recorded ground-motion data from the 2019 Ridgecrest, California, earthquake sequence (Rekoske et al., 2020) to compare the out-of-sample performance of the herein proposed correlation models. Specifically, we compute the within-event residuals of $Sa(1s)$ using the regionalized ergodic GMM developed by Liu et al. (2022a). This GMM has been derived using
290 the NGA-West2 database and a subset of 9,554 records from 81 earthquakes of the Ridgecrest data set. The latter subset serves here as the first test data set denoted as $\tilde{\mathcal{D}}_1$. The second test data set, $\tilde{\mathcal{D}}_2$, consists of data from three events with magnitudes 5.4, 6.4 and 7.1 that was not used to fit the GMM. For more details on the ground-motion data and model the reader is referred to Rekoske et al. (2020) and Liu et al. (2022a).

The out-of-sample performance is first assessed via the LPPD of the test set $\tilde{\mathcal{D}}_1$ obtained with the pooled correlation models
295 E, EA and EAS from Sect. 3.4.2. Table 3 shows the resulting relative difference in the LPPDs to the LPPD of an independent model and compares the values to the in-sample performance on the training set. The in-sample values are identical to the results shown in Fig. 7 at period $T = 1s$. For the test set $\tilde{\mathcal{D}}_1$ the increase in LPPD of the model EAS is similar to the in-sample LPPD, whereas the increase of models E and models EA is more than 1% point smaller. The comparison of models EA and

EAS reveals that the additional gain in predictive accuracy obtained from accounting for site effects is larger for the test data
300 (Ridgecrest) than for the training data (NGA-West2).

Table 3. In- and out-of-sample performance in terms of relative difference (Eq. 10) in LPPD of pooled models E, EA and EAS to the baseline
LPPD of an independent model for the training and test data sets.

	Relative difference in LPPD [%]	
Model	In-sample (NGA-West2 $\mathcal{D}_{tot}$)	Out-of-sample (Ridgecrest $\tilde{\mathcal{D}}_1$)
E	9.42	8.32
EA	9.96	8.73
EAS	10.47	10.32

For the three events in the second test data set, $\tilde{\mathcal{D}}_2$, Fig. 8 shows the computed within-event residuals in the top row and
the relative difference in LPPD compared to the corresponding independent model in the bottom row. Similar to Fig. 5, the
histograms show the log-likelihood of the different models conditional on samples from the posterior distributions of parameter
values, whereas the X mark the marginalized LPPD. For all three events, model EAS has a higher predictive accuracy than the
305 isotropic model E.
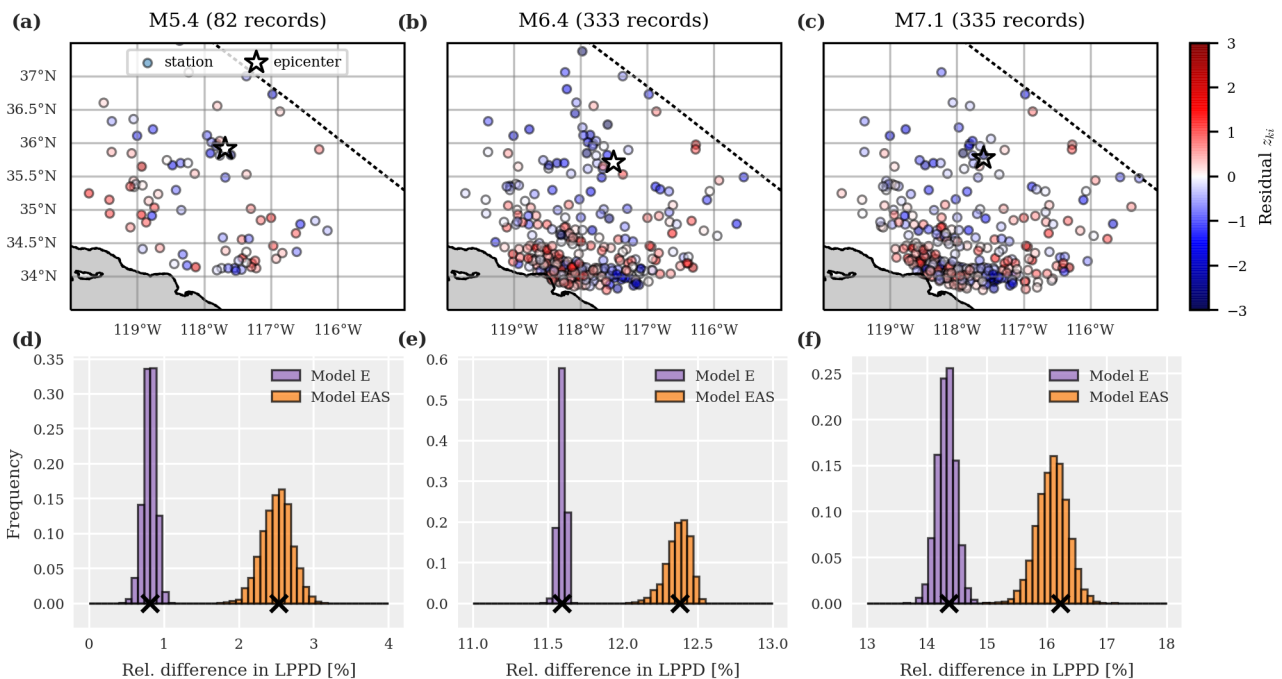
Figure 8. For the three earthquakes in the second test set from the 2019 Ridgecrest sequence: The spatial distribution of scaled within-event
residuals $\mathbf{z}_k$ of $Sa(1s)$ (top row) and the relative difference (Eq. 10) in Log Posterior Predictive Density (LPPD) of pooled models E and
EAS to the baseline LPPD of the independent model (bottom row).

6 Implications for regional seismic risk

Figure 9 shows the configuration of a simplified case study to assess the implications of the different correlation models for regional seismic risk simulations. Four 5x5km subregions are located in the San Francisco Bay Area, California, subjected to a M6.25 rupture on the Hayward fault. The four subregions are gridded into individual sites with a spacing of 3 arc-seconds

310 (approximately 90m). The quantity of interest is the proportion of sites, a, within one or several subregions, where the $Sa(1s)$ jointly exceeds a certain threshold value sa, which we denote as random variable $A_{Sa(1s)>sa}$. While regional seismic risk is typically quantified in terms of fatalities, financial losses or downtime, the joint distribution of Sa values is a major component in the underlying workflow and an important driver of such regional seismic risk metrics (Weatherill et al., 2015).

The rupture scenario is taken from the UCERF2 earthquake rupture forecast (Field et al., 2008). We use OpenSHA (Field

315 et al., 2003) to compute the mean $\ln Sa$ value at all sites, as well as the between- and within-event standard deviations, τ and ϕ, from the GMM of Chiou and Youngs (2014). The v_{s30} values, illustrated in Fig. 9, are obtained from Thompson (2022). Then we sample scaled within-event residuals $\mathbf{z}$ from the posterior predictive distribution $p(\mathbf{z}|\mathcal{D}_{tot})$ and multiply them with the within-event standard deviation ϕ to get samples of δW. For δB we sample from a zero-mean normal distribution with standard deviation τ. Finally, we obtain sampled ground-motion fields by summing the mean $\ln Sa$ value and the sampled residuals using

320 Eq. (1), and compute a sample of $A_{Sa(1s)>sa}$ by counting the number of sites where the simulated ground-motion field exceeds a certain threshold value sa.
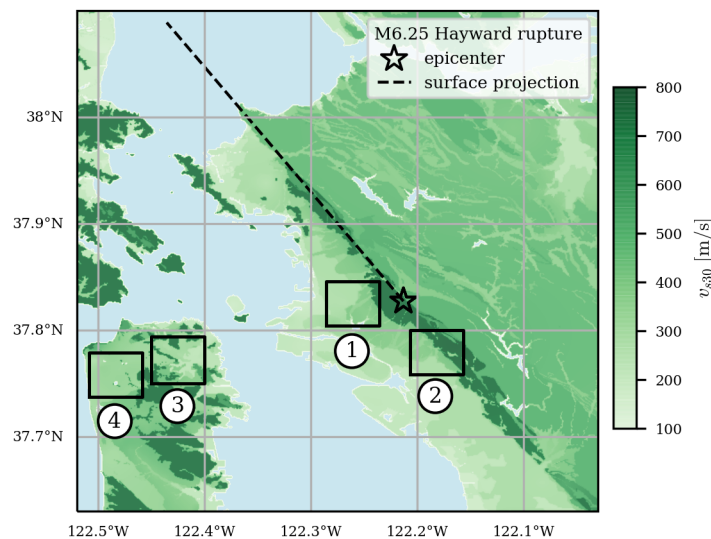
Figure 9. Map of the case study area used for regional risk assessment and the considered M6.25 rupture on the Hayward fault. Numbered boxes indicate the subregions of interest. Colors indicate soil conditions, quantified via v_{s30} (blue areas are water).

Note that the above process requires realizations of $\mathbf{z}$ from each sampled parameter set from the posterior $\boldsymbol{\psi}_{\mathrm{M},r}$ obtained via MCMC. This sampling from the posterior predictive distribution is computationally expensive because, for each $\boldsymbol{\psi}_{\mathrm{M},r}$, it

requires the evaluation of a novel covariance matrix, and a subsequent sampling from the corresponding multivariate normal.
325 Thus, we first explore whether it is sufficient to use the mean values from the posterior distribution of each parameter to build
one covariance matrix and generate all samples from that corresponding multivariate normal. Figure 10 compares the resulting
exceedance probability curves for the proportion of sites in subregion two where $Sa(1s)$, induced by the considered scenario
rupture, jointly exceeds a threshold value of 0.75g. To calculate the exceedance probability curves, we used the posterior
parameters from model EAS, because this model has more parameters and increased estimation uncertainty compared to
330 model E (as shown in Fig. 4). We observe that the exceedance probability curves obtained via both approaches are visually
identical. The same comparison for the other subregions produced the same findings. This indicates that the less expensive
approach of using the mean posterior parameters is sufficient for these regional risk estimates.
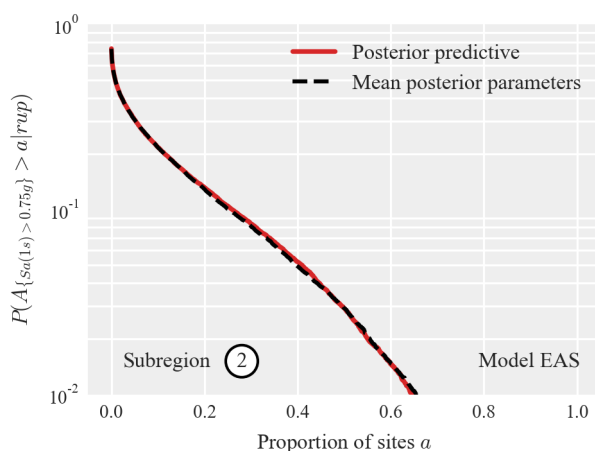
Figure 10. Exceedance probability curves for the proportion of sites a in subregion two where $Sa(1s)$ jointly exceeds a threshold value of
0.75g, obtained via: sampling from the posterior predictive distribution of model EAS (dashed line) and sampling conditional on the mean
values of the posterior parameter distributions (solid line).

We used the posterior mean parameters of model E and model EAS to compute exceedance probability curves for the
proportion of sites within the different subregions where $Sa(1s)$ jointly exceeds a certain threshold values. Figure 11 shows
335 results for subregions one and two (top row) and subregions three and four (bottom row). Note that for the former we use a
threshold value of 0.75g, because of their location close to the fault, whereas for the latter a threshold value of 0.25g was used.
This choice is made for illustrative purposes—the results and conclusions do not change for different threshold values.

In subregions with differing epicentral azimuth values and heterogeneous soil conditions (top row), the isotropic model E
predicts stronger correlations, and thus heavier tailed distributions (i.e., higher probabilities of jointly exceeding the threshold
340 value at a high proportion of sites). This is especially apparent if subregions one and two are combined (Fig. 11c), where the
probability of jointly exceeding the threshold value at least at 40% of all sites is 12% for model E and only 8% for model EAS.
If subregions three and four are combined (Fig. 11f) the exceedance probability curves from both models are similar, because

these subregions are located at similar azimuths from the epicenter and the assigned correlation coefficients are similar for both models (see Fig. 4a). For subregion four (Fig. 11e), the curve obtained via model EAS has slightly heavier tails compared to

345 the one of model E, which may be explained by the homogeneous soil conditions in this subregion (see Fig. 9).
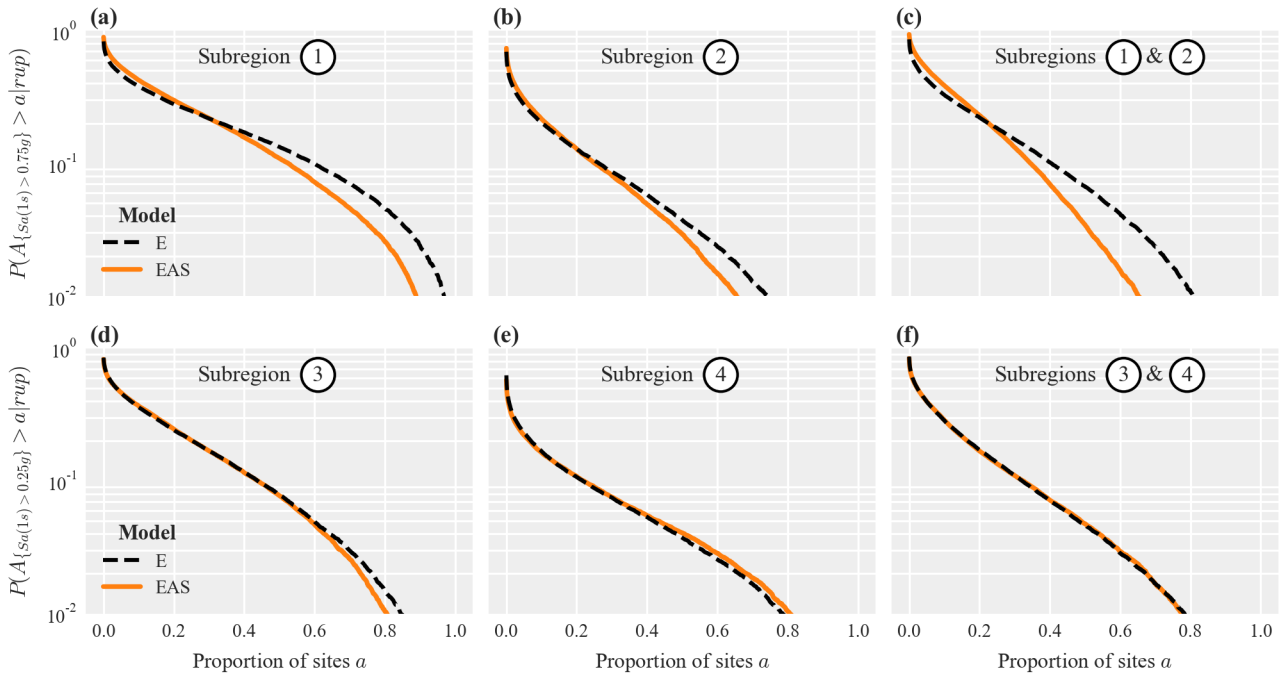
Figure 11. Exceedance probability curves for the proportion of sites a in the considered subregions where $Sa(1s)$ jointly exceeds a threshold value sa: For subregions one and two separately, (a) and (b), and for one and two combined (c) using a threshold value of 0.75g, and in the bottom row for subregions three and four using a threshold value of 0.25g.

7 Conclusions

This study explored the role of spatial proximity, local site effects and path effects on spatial correlations of recorded ground-motion intensity measures. The motivation for this work came from the substantial event-to-event variability found in the correlation model parameters estimated in previous studies, and questions as to whether such variability was due to event-

350 specific characteristics or due to model and estimation uncertainty. Site and path effects are qualitative contributors to spatial correlations but were not captured by the isotropic correlation models used in previous studies: thus, our focus on the path and site effects to explain the observed model parameter variability.

We proposed a novel correlation model, EAS, that accounts for path and site effects in addition to spatial proximity. The EAS model assigns decreasing correlation coefficients for sites with increasing Euclidean distance, increasing angular distance,

355 and increasing soil dissimilarity. These three model components reflect the role of spatial proximity, path effects, and site

effects, respectively, on spatial ground-motion correlations. Compared to an isotropic model, the proposed model has increased complexity and more parameters (five instead of one or two for the isotropic model). To account for this increase in model complexity, we employ Bayesian inference to estimate the parameters, and we assume that the same parameters describe the correlation for all events in the considered ground-motion database (i.e., a pooled model).

360   For each event in the NGA-West2 training data set, we then computed the predictive accuracy of the proposed EAS model, as well as of two isotropic models E, where the parameters of one model were estimated from the pooled data set, and the others exclusively consider data from that specific event. For most events, we found that the event-specific models E have higher predictive accuracy then the pooled model E, thus confirming the presence of some event-to-event variability in correlation model parameters. However, the pooled model EAS outperforms the event-specific models E for the majority of events and, es-

365   pecially, for the well-recorded events. This indicates that the event-to-event variability in estimated isotropic model parameters found in previous studies is an apparent variability due to estimation uncertainty and the lack of accounting for site and path effects, rather than a true variability. Data from the 2019 Ridgecrest earthquake sequence was then used to compare the models in terms of their out-of-sample performance. The results showed a higher predictive accuracy for model EAS, compared to the isotropic model E, further highlighting the benefit of accounting for site and path effects in correlation models.

370   We then used a case study to explore the implications of using the different correlation models for regional seismic risk simulations. First, we found that generating correlated ground-motion samples using the mean values from the posterior distribution of each parameter instead of sampling from the posterior predictive distribution produces ground-motion fields with practically equivalent distributions. This is helpful, because it is much less computationally expensive to use mean parameter values. Second, we saw that the isotropic model E predicts substantially stronger correlations than model EAS in regions with

375   heterogeneous soil conditions and varying epicentral azimuths. This may lead to an overestimation of regional seismic risk tails (low probability, high consequence events), particularly in regions located close to the earthquake source.

The proposed model and analysis could benefit from some further study. This could include a refined model parameterization to consider an azimuth metric that accounts for finite-fault effects, or to consider other metrics of dissimilarity in site conditions. The refined EAS model could also be tested on more complex risk analysis problems, to further understand the practical impact

380   of these refinements. Despite those opportunities for further study, the proposed EAS model form and the proposed techniques for evaluating model performance should be of general use for analysts interested in studying and improving the prediction of spatial correlations in ground motions.

Code and data availability.   The code for Bayesian inference, the post-processing and the case study application is available at www.doi.org/ 10.5281/zenodo.7124213 (Bodenmann, 2022). The PEER NGA-West2 data set (Ancheta et al., 2014) is available at www.apps.peer.berkeley.

385   edu/ngawest2/databases/ (last accessed on April 20, 2022). To compute residuals we used the code from Baker and Chen (2020) available at www.github.com/bakerjw/spatialCorrelationEstimation/. The data set from the 2019 Ridgecrest earthquake sequence (Rekoske et al., 2020) is available at www.strongmotioncenter.org/specialstudies/rekoske_2019ridgecrest/ (last accessed on May 11, 2022). The ergodic GMM employed for this data set has been provided to us by the authors of Liu et al. (2022a).

Author contributions. LB: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Visualization, Writ-
390 ing – original draft preparation. JWB: Conceptualization, Supervision, Writing – review & editing. BS: Resources, Project administration,
Funding acquisition, Writing – review & editing.

Competing interests. The contact author has declared that neither they nor their co-authors have any competing interests.

Acknowledgements. We thank Chenying Liu and Jorge Macedo for their feedback and help with the 2019 Ridgecrest sequence data set. The
calculations were run on the Euler cluster of ETH Zürich. The first author gratefully acknowledges support from the ETH Risk Center (Grant
395 2018-FE-213) and the ETH Doc.Mobility fellowship.

References

Ancheta, T. D., Darragh, R. B., Stewart, J. P., Seyhan, E., Silva, W. J., Chiou, B. S.-J., Wooddell, K. E., Graves, R. W., Kottke, A. R., Boore, D. M., Kishida, T., and Donahue, J. L.: NGA-West2 Database, Earthquake Spectra, 30, 989–1005, https://doi.org/10.1193/070913EQS197M, 2014.

Baker, J., Bradley, B., and Stafford, P.: Seismic Hazard and Risk Analysis, Cambridge University Press, https://doi.org/10.1017/9781108425056, 2021.

Baker, J. W. and Chen, Y.: Ground motion spatial correlation fitting methods and estimation uncertainty, Earthquake Engineering & Structural Dynamics, 49, 1662–1681, https://doi.org/10.1002/eqe.3322, 2020.

Bodenmann, L.: Bayesian parameter estimation for ground-motion correlation models (v1.0.0), https://doi.org/10.5281/zenodo.7124213, zenodo [software], 2022.

Boore, D. M., Gibbs, J. F., Joyner, W. B., Tinsley, J. C., and Ponti, D. J.: Estimated Ground Motion From the 1994 Northridge, California, Earthquake at the Site of the Interstate 10 and La Cienega Boulevard Bridge Collapse, West Los Angeles, California, Bulletin of the Seismological Society of America, 93, 2737–2751, https://doi.org/10.1785/0120020197, 2003.

Chen, Y. and Baker, J. W.: Spatial Correlations in CyberShake Physics-Based Ground-Motion Simulations, Bulletin of the Seismological Society of America, 109, 2447–2458, https://doi.org/10.1785/0120190065, 2019.

Chiou, B. S.-J. and Youngs, R. R.: Update of the Chiou and Youngs NGA Model for the Average Horizontal Component of Peak Ground Motion and Response Spectra, Earthquake Spectra, 30, 1117–1153, https://doi.org/10.1193/072813EQS219M, 2014.

Esposito, S. and Iervolino, I.: PGA and PGV Spatial Correlation Models Based on European Multievent Datasets, Bulletin of the Seismological Society of America, 101, 2532–2541, https://doi.org/10.1785/0120110117, 2011.

Field, E. H., Jordan, T. H., and Cornell, C. A.: OpenSHA: A Developing Community-modeling Environment for Seismic Hazard Analysis, Seismological Research Letters, 74, 406–419, https://doi.org/10.1785/gssrl.74.4.406, 2003.

Field, E. H., Dawson, T. E., Felzer, K. R., Frankel, A. D., Gupta, V., Jordan, T. H., Parsons, T., Petersen, M. D., Stein, R. S., Weldon II, R. J., and Wills, C. J.: The Uniform California Earthquake Rupture Forecast, Version 2 (UCERF 2), Tech. rep., United States Geological Survey, Reston, Virginia, 2008.

Gelman, A., Hwang, J., and Vehtari, A.: Understanding predictive information criteria for Bayesian models, Statistics and Computing, 24, 997–1016, https://doi.org/10.1007/s11222-013-9416-2, 2014.

Goda, K.: Interevent Variability of Spatial Correlation of Peak Ground Motions and Response Spectra, Bulletin of the Seismological Society of America, 101, 2522–2531, https://doi.org/10.1785/0120110092, 2011.

Goda, K. and Atkinson, G. M.: Intraevent Spatial Correlation of Ground-Motion Parameters Using SK-net Data, Bulletin of the Seismological Society of America, 100, 3055–3067, https://doi.org/10.1785/0120100031, 2010.

Goda, K. and Hong, H. P.: Spatial Correlation of Peak Ground Motions and Response Spectra, Bulletin of the Seismological Society of America, 98, 354–365, https://doi.org/10.1785/0120070078, 2008.

Heresi, P. and Miranda, E.: Uncertainty in intraevent spatial correlation of elastic pseudo-acceleration spectral ordinates, Bulletin of Earthquake Engineering, 17, 1099–1115, https://doi.org/10.1007/s10518-018-0506-6, 2019.

Hoffman, M. D. and Gelman, A.: The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo, Journal of Machine Learning Research, 15, 1593–1623, https://doi.org/10.5555/2627435.2638586, 2014.

Jayaram, N. and Baker, J. W.: Statistical Tests of the Joint Distribution of Spectral Acceleration Values, Bulletin of the Seismological Society of America, 98, 2231–2243, https://doi.org/10.1785/0120070208, 2008.

Jayaram, N. and Baker, J. W.: Correlation model for spatially distributed ground-motion intensities, Earthquake Engineering & Structural Dynamics, 38, 1687–1708, https://doi.org/10.1002/eqe.922, 2009.

Kotha, S. R., Weatherill, G., Bindi, D., and Cotton, F.: A regionally-adaptable ground-motion model for shallow crustal earthquakes in Europe, Bulletin of Earthquake Engineering, 18, 4091–4125, https://doi.org/10.1007/s10518-020-00869-1, 2020.

Kuehn, N. M. and Abrahamson, N. A.: Spatial correlations of ground motion for non-ergodic seismic hazard analysis, Earthquake Engineering & Structural Dynamics, 49, 4–23, https://doi.org/10.1002/eqe.3221, 2020.

Kuehn, N. M. and Stafford, P. J.: Estimating Ground-Motion Models via Bayesian Inference using Stan, https://doi.org/10.31224/osf.io/uj749, engrXiv [preprint], 2021.

Lee, R. and Kiremidjian, A. S.: Uncertainty and Correlation for Loss Assessment of Spatially Distributed Systems, Earthquake Spectra, 23, 753–770, https://doi.org/10.1193/1.2791001, 2007.

Liu, C., Macedo, J., and Kottke, A. R.: Evaluating the performance of nonergodic ground motion models in the ridgecrest area, Bulletin of Earthquake Engineering, https://doi.org/10.1007/s10518-022-01342-x, 2022a.

Liu, C., Macedo, J., and Kuehn, N.: Spatial correlation of systematic effects of non-ergodic ground motion models in the Ridgecrest area, Bulletin of Earthquake Engineering, https://doi.org/10.1007/s10518-022-01441-9, 2022b.

Moss, R. E. S. and Der Kiureghian, A.: Incorporating parameter uncertainty into attenuation relationships, in: 8th US National Conference on Earthquake Engineering, San Francisco, CA, 2006.

Neal, R. M.: Probabilistic Inference Using Markov Chain Monte Carlo Methods, Tech. rep., University of Toronto, Canada, 1993.

Padonou, E. and Roustant, O.: Polar Gaussian Processes and Experimental Designs in Circular Domains, SIAM/ASA Journal on Uncertainty Quantification, 4, 1014–1033, https://doi.org/10.1137/15M1032740, 2016.

Phan, D., Pradhan, N., and Jankowiak, M.: Composable effects for flexible and accelerated probabilistic programming in NumPyro, https://doi.org/10.48550/ARXIV.1912.11554, arXiv [preprint], 2019.

Rasmussen, C. and Williams, K.: Gaussian processes for machine learning, The MIT Press, Cambridge, MA, 2006.

Rekoske, J. M., Thompson, E. M., Moschetti, M. P., Hearne, M. G., Aagaard, B. T., and Parker, G. A.: The 2019 Ridgecrest, California, Earthquake Sequence Ground Motions: Processed Records and Derived Intensity Metrics, Seismological Research Letters, 91, 2010–2023, https://doi.org/10.1785/0220190292, 2020.

Schiappapietra, E. and Douglas, J.: Modelling the spatial correlation of earthquake ground motion: Insights from the literature, data from the 2016–2017 Central Italy earthquake sequence and ground-motion simulations, Earth-Science Reviews, 203, 103 139, https://doi.org/10.1016/j.earscirev.2020.103139, 2020.

Schiappapietra, E. and Douglas, J.: Assessment of the uncertainty in spatial-correlation models for earthquake ground motion due to station layout and derivation method, Bulletin of Earthquake Engineering, 19, 5415–5438, https://doi.org/10.1007/s10518-021-01179-w, 2021.

Sokolov, V., Wenzel, F., Wen, K.-L., and Jean, W.-Y.: On the influence of site conditions and earthquake magnitude on ground-motion within-earthquake correlation: analysis of PGA data from TSMIP (Taiwan) network, Bulletin of Earthquake Engineering, 10, 1401–1429, https://doi.org/10.1007/s10518-012-9368-5, 2012.

Stafford, P. J.: Continuous integration of data into ground-motion models using Bayesian updating, Journal of Seismology, 23, 39–57, https://doi.org/10.1007/s10950-018-9792-3, 2019.

Thompson, E. M.: An Updated Vs30 Map for California with Geologic and Topographic Constraints (ver. 2.0, July 2022), U.S. Geological
470   Survey data release, https://doi.org/10.5066/F7JQ108S, 2022.

Weatherill, G. A., Silva, V., Crowley, H., and Bazzurro, P.: Exploring the impact of spatial correlations and uncertainties for portfolio analysis
       in probabilistic seismic loss estimation, Bulletin of Earthquake Engineering, 13, 957–981, https://doi.org/10.1007/s10518-015-9730-5,
       2015.

Wesson, R. L. and Perkins, D. M.: Spatial Correlation of Probabilistic Earthquake Ground Motion and Loss, Bulletin of the Seismological
475   Society of America, 91, 1498–1515, https://doi.org/10.1785/0120000284, 2001.