

The Geometry of Earthquake Shaking: A Probabilistic Model for Response-Spectrum Directionality

Rajesh Rupakhety

rajesh@hi.is

University of Iceland

Victor Moises Hernández-Aguirre

University of Iceland

Research Article

Keywords: response spectra, rotation-invariant measures, covariance anisotropy, ground-motion directionality, effective sample size

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10988530/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

The Geometry of Earthquake Shaking: A Probabilistic Model for Response-Spectrum Directionality

Rajesh Rupakhety^{1*}, Víctor M. Hernández-Aguirre¹

Abstract:

Directional variability of horizontal response spectra is commonly characterized through rotation-invariant measures such as RotD50 and RotD100, yet the statistical behaviour of the underlying directional field remains poorly understood. This study shows that a substantial fraction of the observed directional variability can be explained by a finite-sample effect: when a record contains few statistically independent response cycles, covariance estimates become inherently variable, producing apparent anisotropy even under isotropic excitation. Using 4690 strong-motion records from the Engineering Strong Motion database, we develop a probabilistic model for the RMS anisotropy parameter κ_{rms} . The model expresses the effective sample size, N_{eff} , as a function of the duration-to-period ratio, ψ , through three physically motivated regimes: a lower floor associated with the few-cycle regime of covariance estimation, an intermediate narrowband-scaling regime, and a short-period saturation imposed by the finite bandwidth of the excitation. Despite containing only a single calibrated parameter, the model reproduces the observed κ_{rms} distributions across the full period range considered. Residual analyses indicate that much of the apparent dependence of anisotropy on earthquake magnitude and source-to-site distance is removed after conditioning on this stochastic baseline. Pulse-like motions and site response contribute genuine physical anisotropy but remain secondary to the duration-controlled stochastic baseline. When combined with an existing stochastic model for directional peak-factor variability, the framework enables simulation of directional response spectra and reproduces the principal features of observed directional response. The results support reinterpreting directional variability as a bounded stochastic geometric phenomenon governed primarily by the duration-to-period ratio. This perspective provides a physically grounded basis for directional ground-motion modelling, record selection, and probabilistic seismic hazard analysis, and suggests that compact, low-dimensional stochastic representations of horizontal ground-motion fields are an achievable goal.

Keywords: response spectra, rotation-invariant measures, covariance anisotropy, ground-motion directionality, effective sample size

1. Introduction

Horizontal ground motion is a two-component vector process. Consequently, the intensity experienced along a structural axis depends on the orientation of that axis with respect to the horizontal ground-motion field, and this dependence can be substantial (Hong and Goda, 2007; Poulos and Miranda, 2022). Such behaviour arises from a combination of source, path, and site effects, including rupture directivity (Huang et al., 2008; Somerville et al., 1997), radiation pattern

¹ Earthquake Engineering Research Centre, Faculty of Civil and Environmental Engineering, University of Iceland, Austurvegur 2a, 800 Selfoss, Iceland

^{1*} Corresponding author, rajesh@hi.is

(Kotha et al., 2019; Poulos and Miranda, 2023), wave-propagation complexity, directional topographic amplification (Spudich et al., 1996), and heterogeneities in the near surface (Baena-Rivera et al., 2024; Pischiutta et al., 2018). As a result, horizontal response spectra are orientation dependent rather than intrinsically scalar. These effects are increasingly considered in record selection (Nievas and Sullivan, 2018), seismic assessment (Skoulidou et al., 2019), and loss estimation (Feng et al., 2021; Lagaros, 2010; Poulos and Miranda, 2025). Yet probabilistic seismic hazard analysis (PSHA), ground-motion modelling, and structural design are formulated almost entirely in terms of scalar intensity measures. The directional response field is therefore typically condensed into one or a few scalar summaries that retain only selected features of the underlying directional structure.

Rotation-invariant measures provide a useful and reproducible basis for this reduction. RotD50, defined as the median pseudo-spectral acceleration (PSA) over all orientations (Boore, 2010), and RotD100, the corresponding maximum, have become the standard horizontal-component definitions in modern ground-motion modelling and in the design procedures built upon them (Boore et al., 2014; Bozorgnia et al., 2014). Their ratio, RotD100/RotD50, is bounded between unity for an isotropic response and $\sqrt{2}$ for a perfectly polarized response, and is widely interpreted as a measure of ground-motion directionality.

However, these quantities are summaries of the directional response field (Rupakhety, 2026). Geometrically, replacing an anisotropic response field by a rotation-invariant measure is equivalent to replacing it with a circle of prescribed radius. RotD50 represents a typical directional level, whereas RotD100 corresponds to a maximum-direction or resultant-envelope quantity. Neither preserves the shape or orientation of the original response field. Preserving orientation dependence is also relevant in urban risk analysis, where common building orientations can increase the correlation and variance of aggregate losses (Poulos and Miranda, 2025).

More fundamentally, RotD100/RotD50 does not isolate a unique form of anisotropy. Directional PSA can be separated into a root mean square (RMS) response field and an orientation-dependent peak factor (Rupakhety and Hernández-Aguirre, 2026a), analogous to the classical random-vibration-theory separation between RMS response and finite-duration peak formation (Vanmarcke and Lai, 1980). The RMS field is determined by the covariance structure of the two horizontal oscillator-response histories and therefore possesses an exact elliptical second-order geometry. The peak factor introduces an additional angular modulation associated with the stochastic formation of temporal extremes (Rupakhety and Hernández-Aguirre, 2026b). Consequently, RotD100/RotD50 reflects the combined influence of RMS anisotropy and stochastic variability of directional peak formation rather than a single underlying state variable.

Existing empirical treatments of ground-motion directionality can be organised into three broad groups. The first relies on conversion models between scalar horizontal-component definitions. Period-dependent relationships for RotD100/RotD50 and related quantities allow intensity measures to be translated from one rotational percentile to another (Boore and Kishida, 2017; Shahi and Baker, 2014a). These models are practical and compatible with established ground-motion modelling workflows, but their targets are selected scalar functionals of the directional field rather than the underlying geometric field. Because RotD100/RotD50 is bounded, additional

assumptions are required when modelling its residual variability probabilistically, particularly in the tails of the distribution where physical bounds become relevant.

The second group characterises spectral ordinates at arbitrary orientations relative to the realised direction of maximum response. Poulos and Miranda (2022) developed probabilistic models for spectral ordinates normalised by RotD50 or RotD100 as functions of angular distance from the RotD100 axis. Their framework models the means, dispersions, correlations, and bounded distributions of these ratios and therefore retains substantially more angular information than a conversion between two scalar percentiles. However, the formulation remains descriptive: directional distributions are modelled directly rather than through an underlying state variable governing their evolution. As a result, directional variability is characterised as a function of oscillator period, but the physical quantities controlling that variability are not explicitly identified.

The third group incorporates directional variability into the aleatory variability of ground-motion models (GMMs) through component-to-component variability terms (Baker and Cornell, 2006; Beyer and Bommer, 2006; Girmay and Miranda, 2026). This approach is convenient for prediction and hazard calculations because directional effects can be represented through additional variability within existing GMM frameworks. However, directional variability remains represented as an additional source of variability rather than as the consequence of an underlying state variable governing the directional response field.

Taken together, these approaches provide increasingly detailed descriptions of directional peak response, but none supplies a predictive model for geometric anisotropy that exists before peak extraction. A natural variable for this purpose is the normalised contrast between the principal eigenvalues of the two-component response covariance matrix, denoted here by κ_{rms} (Rupakhety and Hernández-Aguirre, 2026). It is bounded and scale invariant, is zero for a circular RMS field, and increases as the response covariance becomes concentrated along one direction. It therefore isolates geometric elongation independently of the overall response amplitude and of the subsequent angular modulation introduced by peak formation.

Importantly, non-zero κ_{rms} does not necessarily imply a persistent physical preferred direction. Even when the population excitation is isotropic, a covariance matrix obtained from a finite-duration response history generally has unequal eigenvalues. Finite covariance sampling therefore produces realisation-dependent anisotropy. Under the isotropic Gaussian reference model, the distribution of κ_{rms} is governed by the amount of effectively independent covariance information in the oscillator response, represented by N_{eff} (Rupakhety and Hernández-Aguirre, 2026).

Previous work established the distributional basis of κ_{rms} and derived a first-order narrowband approximation relating N_{eff} to oscillator period, damping, and energetic response duration (Rupakhety and Hernández-Aguirre, 2026). That approximation captures much of the observed period and duration dependence over an intermediate period range, but it is not a population-level predictive model for recorded ground motions. At short periods, N_{eff} cannot increase indefinitely because the oscillator inherits the finite bandwidth of the excitation. At long periods, it cannot decrease indefinitely because a two-dimensional covariance estimate possesses a minimum admissible number of degrees of freedom. Recorded motions also differ in duration,

bandwidth, temporal correlation, and possible persistent physical anisotropy, so their variability cannot generally be represented by one common effective sample size.

The remaining problem is therefore to model how N_{eff} , and hence the conditional distribution of κ_{rms} , varies across real strong-motion records. Such a model should preserve the intermediate-period scaling implied by the covariance framework while accommodating the finite-bandwidth behaviour at short periods and the few-cycle behaviour at long periods.

This study develops a probabilistic model for the geometric anisotropy of directional response spectra. It derives a composite model for the effective sample size governing covariance fluctuations, calibrated using a large strong-motion database. The resulting formulation yields period- and duration-dependent distributions of κ_{rms} , establishes a stochastic baseline for directional anisotropy, and enables a generative representation of the full directional-spectrum field. The model is validated through distributional diagnostics and used to examine departures associated with persistent physical directionality.

2. Directional field and anisotropy

At a fixed oscillator period, PSA can be viewed as a stochastic field defined on the circle. The purpose of this section is to identify the geometric quantity governing that field and to distinguish it from anisotropy measures based directly on spectral peaks. Directional anisotropy enters first through the response covariance, and once the covariance structure is specified, the dominant geometry of the directional field follows immediately.

2.1 Geometric anisotropy from the response covariance

Let $\mathbf{a}(t)$ be two orthogonal horizontal components of an oscillator's pseudo-acceleration response at period T . Along orientation θ , measured from the direction of the first component, the response is the projection

$$a(t, \theta) = \mathbf{u}(\theta)^T \mathbf{a}(t), \quad \mathbf{u}(\theta) = (\cos \theta, \sin \theta)^T. \quad (1)$$

The variance of the projected response depends on orientation because the two horizontal components are generally correlated and need not have equal variance. The variance over the record is $\sigma_a^2(\theta; T) = \mathbf{u}(\theta)^T \mathbf{C}(T) \mathbf{u}(\theta)$, where $\mathbf{C}(T)$ is the response covariance matrix. Writing its eigenvalues as λ_1 and λ_2 , with $\lambda_1 \geq \lambda_2$, and choosing θ_0 along the principal axis associated with λ_1 , the variance assumes the familiar elliptical form

$$\sigma_a^2(\theta, T) = \overline{\sigma_a^2}(T) [1 + \kappa_{\text{rms}}(T) \cos 2(\theta - \theta_0)] \quad (2)$$

with

$$\overline{\sigma_a^2} = \frac{\lambda_1 + \lambda_2}{2}, \quad \kappa_{\text{rms}} = \frac{\lambda_1 - \lambda_2}{\lambda_1 + \lambda_2}, \quad 0 \leq \kappa_{\text{rms}} < 1 \quad (3)$$

The parameter κ_{rms} measures the geometric anisotropy of the covariance ellipse. Because it depends only on the covariance structure, it is unaffected by how response peaks are extracted. This makes it a natural candidate for a state variable describing directional anisotropy.

2.2 The directional spectrum as a modulated ellipse

Directional PSA, defined as $PSA(\theta, T) = \max_t |a(t, \theta)|$, can be decomposed into a geometric component associated with the covariance ellipse, and a stochastic component associated with the directional variation of response peaks: $PSA(\theta, T) = \sigma_a(\theta, T)p(\theta, T)$.

The logarithm of the directional peak factor, $p(\theta, T)$, can be in turn decomposed into a mean level and a directional fluctuation, $\ln p(\theta) = \mu_p + \varepsilon(\theta)$, where $\varepsilon(\theta)$ is π -periodic, has zero median, and is dominated by its first angular harmonic (Rupakhety and Hernández-Aguirre, 2026b). Substituting this representation into the directional spectrum and expressing it in squared form gives

$$S(\theta, T) = S_0(T) [1 + \kappa_{\text{rms}}(T) \cos(2(\theta - \theta_0))] \exp(2\varepsilon(\theta, T)) \quad (4)$$

Equation (4) is the fundamental representation used throughout this paper. The bracketed term is the RMS ellipse, governed entirely by the covariance anisotropy κ_{rms} . The exponential term is the peak-factor field, which introduces directional fluctuations around that ellipse. The scale factor $S_0(T)$ controls the overall amplitude of the field and is distinct from its directional structure. The present work is concerned with the latter. In practical applications, the scale may be supplied externally, for example by a GMM for RotD50 or another rotation-invariant intensity measure. The remaining quantities, κ_{rms} , θ_0 , and $\varepsilon(\theta)$, then determine how that amplitude is distributed over orientation. The directional PSA field may therefore be viewed as the product of a scalar intensity model and a stochastic directional-shape model.

2.3 Two measures of anisotropy

The covariance anisotropy κ_{rms} is not the only way to quantify directional variability. An anisotropy measure can also be defined directly from the directional spectrum by comparing amplitudes along its principal axes,

$$\kappa_{\text{PSA}} = \frac{S(\theta_0) - S(\theta_0 + \pi/2)}{S(\theta_0) + S(\theta_0 + \pi/2)} \quad (5)$$

This quantity measures the anisotropy of the directional spectrum itself. However, as shown by Rupakhety and Hernández-Aguirre (2026), κ_{PSA} reflects both the covariance geometry and the directional fluctuations of the peak-factor field; these

two sources of variability cannot be separated from spectral amplitudes alone. On the other hand, κ_{rms} isolates the geometric component associated with the covariance structure. Since the directional spectrum can be reconstructed from κ_{rms} and the peak-factor field through Eq. (4), κ_{rms} provides the more fundamental quantity to model. The remainder of this paper is therefore concerned with its stochastic behaviour.

3. Effective degrees of freedom

The covariance model of Rupakhety and Hernández-Aguirre (2026) reduces the problem of geometric anisotropy to a single quantity: the effective number of independent samples contributing to the covariance estimate, N_{eff} . For an isotropic process of finite duration, the Wishart distribution (Muirhead, 2009) leads to

$$\kappa_{\text{rms}}^2 \sim \text{Beta}\left(1, \frac{N_{\text{eff}} - 1}{2}\right). \quad (6)$$

N_{eff} measures how many statistically independent fluctuations of the oscillator response contribute to the covariance estimate. Long response durations provide more information and therefore larger values of N_{eff} . Strong temporal correlation reduces the amount of independent information and therefore decreases it. Directional anisotropy is ultimately governed by the balance between these two effects. For a lightly damped oscillator with period T and damping ratio ξ observed over duration D , Rupakhety and Hernández-Aguirre (2026) show that

$$N_{\text{eff}} \approx 4\pi\xi \frac{D}{T}$$

This result explains the familiar increase of directionality with period reported by various empirical studies (Bradley and Baker, 2015; Shahi and Baker, 2014a). Within the covariance framework, the period dependence arises because longer-period oscillators provide fewer effectively independent response fluctuations from which the covariance structure can be estimated. This expression, however, applies only when the oscillator undergoes many effectively independent cycles. At short oscillation periods, the finite bandwidth of the excitation limits the growth of N_{eff} . At long periods, the effective sample size cannot decrease indefinitely, and the asymptotic narrowband approximation eventually breaks down. The derivation below provides a general expression for N_{eff} , from which the narrowband regime, the short-period ceiling, and the long-period floor emerge naturally.

3.1 The covariance window

Anisotropy is inferred from the response covariance, so the first question is over what part of the record that covariance should be estimated. The relevant duration is not the duration of the ground motion itself, but the duration over which the oscillator accumulates most of its response energy. Let $X_T(t)$ and $Y_T(t)$ denote the two horizontal response components at period T , with resultant $a_{R,T}(t)$, its energetic duration (Rupakhety and Hernández-Aguirre, 2026) is defined as

$$D_R^*(T) = \frac{\left(\int a_{R,T}^2(t) dt\right)^2}{\int a_{R,T}^4(t) dt} \quad (7)$$

This quantity is a participation ratio that characterizes how concentrated the response energy is in time. It is large when response energy is distributed across many cycles and small when that energy is concentrated into a few dominant excursions. The covariance is evaluated over the window of length $D_R^*(T)$ containing the largest response energy. Because the resultant response is rotation invariant, this window is independent of the choice of horizontal axes. Within this window beginning at time t_0 , the cross-covariance estimate is

$$\widehat{C}_{XY}(T) = \frac{1}{D_R^*(T)} \int_{t_0}^{t_0 + D_R^*(T)} X_T(t) Y_T(t) dt. \quad (8)$$

3.2 Effective sample size from covariance uncertainty

The effective sample size answers how many independent samples would produce the same uncertainty in the covariance estimate as the correlated oscillator response. The variance of the

covariance estimator under the isotropic Gaussian process (null) is used as a measure of such uncertainty. Within the energetic window, let the response components satisfy $\mathbb{E}[X_T(t)] = \mathbb{E}[Y_T(t)] = 0$, $\text{Var}[X_T(t)] = \text{Var}[Y_T(t)] = \sigma_T^2$, and $\text{Corr}[X_T(t), X_T(s)] = \text{Corr}[Y_T(t), Y_T(s)] = \rho_T(t - s)$. Under the isotropic Gaussian reference model, the two horizontal response components are taken as independent realisations of the same process, giving

$$\text{Cov}[X_T(t)Y_T(t), X_T(s)Y_T(s)] = \sigma_T^4 \rho_T^2(t - s).$$

The square of the correlation function appears because covariance estimation depends on products of the response process rather than on the response itself. Correlation therefore enters through pairwise products, causing the memory of the covariance estimator to decay as ρ_T^2 . Using Eq. (8), and changing variables to $\tau = t - s$, the variance of the covariance estimator becomes

$$\text{Var}[\widehat{C}_{XY}(T)] = \frac{\sigma_T^4}{(D_R^*(T))^2} \int_{-D_R^*(T)}^{D_R^*(T)} (D_R^*(T) - |\tau|) \rho_T^2(\tau) d\tau \quad (9)$$

The weighing factor $D_R^*(T) - |\tau|$ counts the number of observation pairs separated by lag τ . For $\tau = 0$, every sample contributes a pair. As $|\tau|$ increases, fewer pairs remain within the window. The weighting therefore decreases linearly to zero at the window boundaries. For N independent samples, the corresponding covariance variance is simply σ_T^4/N . Equating this expression with Eq. (9) therefore defines the effective number of independent samples,

$$\frac{D_R^*(T)}{N_{\text{eff}}(T)} = \int_{-D_R^*(T)}^{D_R^*(T)} \left(1 - \frac{|\tau|}{D_R^*(T)}\right) \rho_T^2(\tau) d\tau \quad (10)$$

Equation (10) shows that the effective sample size is determined not by the number of recorded points, but by the amount of effectively independent covariance information contained within the energetic response window.

3.3 Asymptotic regimes and composite model

Evaluation of Eq. (10) requires knowledge of the response autocorrelation function. Additional insight can be obtained by examining its limiting behaviour. When the energetic window contains many correlation lengths, the factor $D_R^*(T) - |\tau|$ varies slowly over the range where $\rho_T^2(\tau)$ is significant. The function $D_R^*(T) - |\tau|$ may then be approximated by the constant value $D_R^*(T)$, reducing Eq. (10) to

$$N_{\text{eff}}(T) \approx \frac{D_R^*(T)}{I_T}, \quad \text{with } I_T = \int_{-\infty}^{\infty} \rho_T^2(\tau) d\tau \quad (11)$$

The quantity I_T is the correlation timescale of the covariance estimator. It measures how long covariance information remains redundant. The effective sample size is therefore the ratio of available duration to covariance memory. A lightly damped oscillator retains memory over many oscillation cycles. Its autocorrelation is approximately

$$\rho_T(\tau) \approx e^{(-\alpha|\tau|)} \cos(\omega_d \tau) \quad (12)$$

249 with $\alpha = 2\pi\xi/T$, $\omega_d = \omega\sqrt{(1-\xi^2)}$, and $\omega = 2\pi/T$. Substituting into Eq. (11) gives $I_T \approx T/(4\pi\xi)$,
 250 and therefore

$$N_{\text{eff}}(T) \approx 4\pi\xi \frac{D_R^*(T)}{T} = 4\pi\xi e^\psi, \quad \psi = \ln\left(\frac{D_R^*(T)}{T}\right) \quad (13)$$

251 This recovers the narrowband result of Rupakhety and Hernández-Aguirre (2026). In this regime,
 252 the effective sample size is proportional to the number of effectively independent oscillator
 253 cycles contained within the energetic window. This asymptotic form is denoted by N_{osc} . The
 254 opposite limit occurs at short periods. As T decreases, the oscillator becomes increasingly
 255 transparent to the excitation, and the response approaches the ground acceleration itself. The
 256 response autocorrelation therefore approaches that of the ground motion, $\rho_T(\tau) \rightarrow \rho_g(\tau)$, and
 257 the covariance timescale converges to

$$I_g = \int_{-\infty}^{\infty} \rho_g^2(\tau) d\tau = \frac{1}{B_{\text{eff},g}} \quad (14)$$

258 Here $B_{\text{eff},g}$ is the effective bandwidth of the ground motion,

$$B_{\text{eff},g} = \frac{2\sigma_g^4}{\int_0^\infty P_g^2(f) df} \quad (15)$$

259 computed from the combined horizontal power spectral density $P_g(f) = P_U(f) + P_V(f)$ within
 260 the energetic window of the record. The equivalence of the time-domain and frequency-domain
 261 expressions follows from Parseval's theorem. The effective sample size therefore approaches the
 262 finite limit

$$N_\infty = D_R^* B_{\text{eff},g}. \quad (16)$$

263 This ceiling reflects a simple physical constraint. A ground motion record contains only a finite
 264 number of statistically independent fluctuations. Once the oscillator becomes faster than the
 265 correlation timescale of the excitation, further reductions in period cannot create additional
 266 independent samples.

267 The effective sample size is therefore constrained by two mechanisms. Oscillator memory limits
 268 the number of statistically independent response cycles at intermediate and long periods, while
 269 ground-motion bandwidth imposes a finite ceiling at short periods. The effective sample size
 270 must satisfy both constraints simultaneously. Since $N_{\text{eff}} \approx D_R^*/I_T$, reciprocal effective sample
 271 size is proportional to covariance memory per unit duration. If oscillator memory and ground-
 272 motion memory contribute independently to the total covariance memory, their contributions
 273 add. This leads naturally to,

$$\frac{1}{N_{\text{eff}}} = \frac{1}{N_{\text{osc}}} + \frac{1}{N_\infty} \quad (17)$$

274 Using Eq. (16) and the definition of N_{osc} gives

$$N_{\text{eff}}(\psi) = \frac{A e^{\psi}}{1 + B e^{\psi}}, \quad A = 4\pi\xi, \quad B = \frac{4\pi\xi}{N_{\infty}} \quad (18)$$

This expression is the simplest form consistent with both asymptotic limits. When oscillator memory dominates, it reduces to the narrowband scaling. When ground-motion bandwidth dominates, it approaches the ceiling N_{∞} .

The large-window approximation underlying Eq. (18) eventually breaks down when $D_R^*(T)/T$ becomes small. In this regime, the energetic window contains only a fraction of an effective response cycle, and the large-window asymptotic scaling is no longer valid. The effective sample size instead approaches a finite lower bound. The lower limit follows directly from the anisotropy distribution.

The asymptotic interpretation of N_{eff} as an effective number of independent covariance observations loses meaning once the available information becomes comparable to the three independent parameters of the two-dimensional covariance matrix. This transition naturally occurs at $N_{\text{eff}} = 3$. Consistently, the isotropic-null anisotropy distribution becomes uniform in κ_{rms}^2 at $N_{\text{eff}} = 3$, while smaller values lead to a singular concentration towards perfectly polarised covariance ellipses. We therefore adopt $N_{\text{min}} = 3$ as a natural lower limit of the effective-sample-size closure. Incorporating this constraint gives the final model,

$$N_{\text{eff}}(\psi) = \max\left(N_{\text{min}}, \frac{A e^{\psi}}{1 + B e^{\psi}}\right), \quad A = 4\pi\xi, \quad B = \frac{4\pi\xi}{N_{\infty}}, \quad N_{\text{min}} = 3 \quad (19)$$

3.4 From the effective sample size to a distribution

Once the effective sample size has been specified, the distribution of anisotropy follows directly. From Eq. (6), the squared anisotropy is Beta distributed, with shape parameter determined by N_{eff} . Transforming from κ_{rms}^2 to κ_{rms} gives a Kumaraswamy (1980) distribution,

$$\kappa_{\text{rms}} \sim \text{Kumaraswamy}(2, b), \quad b = \frac{N_{\text{eff}}(\psi) - 1}{2}. \quad (20)$$

Its cumulative distribution is

$$F(\kappa) = 1 - (1 - \kappa^2)^b. \quad (21)$$

The parameter b controls the concentration of the distribution near isotropy. As N_{eff} increases, the distribution contracts toward $\kappa_{\text{rms}} = 0$, reflecting the increasing stability of the covariance estimate. Records containing many effectively independent response fluctuations therefore tend to exhibit weaker apparent anisotropy than records containing only a few.

Equations (20)–(21) define the distribution for an individual record once its energetic duration, and hence its value of ψ , is known. The population distribution is obtained by recognising that different records have different durations and therefore different effective sample sizes. At a fixed oscillator period, the observed anisotropy is therefore a mixture of Kumaraswamy distributions

$$F(\kappa | T) = \frac{1}{n} \sum_{i=1}^n [1 - (1 - \kappa^2)^{b(\psi_i)}], \quad b(\psi_i) = \frac{N_{\text{eff}}(\psi_i) - 1}{2}. \quad (22)$$

The central problem therefore shifts from modelling anisotropy directly to modelling the distribution of effective sample sizes across records. This mixture distribution represents the stochastic baseline expected under the isotropic null. The extent to which observations depart from it provides a direct measure of anisotropy beyond that expected from finite-sample covariance fluctuations alone.

4 Data and processing

This study is based on 4690 horizontal-component strong-motion recordings obtained from the European Strong Motion (ESM) database (Lanzano et al., 2021). These recordings correspond to 417 earthquake events and 1304 recording stations, spanning a broad range of magnitudes, source-to-site distances, faulting styles, and site conditions. Figure 1 shows the distribution of the dataset in terms of moment magnitude M_w and Joyner–Boore distance R_{JB} .

Records are classified according to the ESM fault-mechanism code. Normal-faulting events constitute the largest subset (2564 records, 55%), followed by strike-slip events (1200 records, 26%) and thrust events (926 records, 20%). The broad coverage of magnitudes and distances allows anisotropy to be examined across a wide range of source and propagation conditions.

A subset of 121 recordings is identified as pulse-like using the procedure of (Shahi and Baker, 2014b). These records occur predominantly at moderate-to-large magnitudes and short-to-intermediate distances (Figure 1). Because pulse-like motions represent a physically distinct class of ground motion, they are examined separately in later sections.

Microtremor HVSr (horizontal to vertical spectral ratio) measurements are available for 421 stations, corresponding to approximately 32% of the recording sites. These data are used to investigate possible associations between site resonance, frequency-dependent amplification, and spectral anisotropy.

For each record, directional PSA fields are computed using linear elastic single-degree-of-freedom response analysis. One hundred logarithmically spaced oscillator periods between 0.01 s and 10 s are considered, with a viscous damping ratio of 5% of critical. For each record and period, the squared directional spectrum $S(\theta; T)$ is evaluated, the principal axes of the response covariance are determined, and the covariance anisotropy κ_{rms} is calculated according to Eq. (3).

To avoid artefacts associated with long-period filtering, periods exceeding $T_{\text{max}} = 0.8/f_{\text{hp}}$ are excluded from the analysis, where f_{hp} is the corner frequency of the high-pass filter applied during record processing. This criterion ensures that anisotropy estimates at long periods are based only on spectral ordinates that remain reliably constrained by the recorded signal.




Figure 1. Distribution of the 4690 ESM records analysed in this study in terms of moment magnitude M_w and Joyner-Boore distance R_{JB} . Records are colored by faulting mechanism (normal, strike-slip, thrust). Pulse-like near-fault records ($n = 121$) are indicated by open circles. For visualization on a logarithmic scale, records with $R_{JB} = 0$ km are plotted at 0.05 km.

5 Calibration and validation of the stochastic anisotropy model

Section 3 derived a stochastic model for covariance anisotropy in terms of a dimensionless variable $\psi = \log(D_R^*(T)/T)$, which controls the effective sample size and, through it, the distribution of κ_{rms} . The purpose of this section is to examine whether the behaviour predicted by the theory is present in the data, to calibrate the parameters of the model, and to assess the extent to which the resulting distribution reproduces the statistics of the ESM dataset.

5.1 Empirical effective sample size

The theory developed in Section 3 predicts that the effective sample size should exhibit three distinct regimes: a lower floor associated with finite information content of covariance estimation, a narrowband scaling regime in which N_{eff} grows approximately as e^ψ , and an upper ceiling imposed by the finite bandwidth of the ground motion. Before calibrating the model, it is therefore useful to examine whether these behaviours are present in the observations.

For a sample κ_j , the maximum-likelihood estimator of the Beta shape parameter b is

$$\hat{b} = -\frac{n}{\sum_{j=1}^n \log(1 - \kappa_j^2)}, \quad \hat{N}_{\text{eff}} = 2\hat{b} + 1 \quad (23)$$

The estimator is evaluated within bins of ψ , aggregating all valid observations across records and oscillator periods that fall within each bin. Because both κ_{rms} and ψ are computed from the same energetic response window but arise from independent parts of the theory, agreement between the empirical estimates and the narrowband prediction provides a stringent test of the covariance framework.

Figure 2 compares the empirical estimates $\hat{N}_{\text{eff}}(\psi)$ with this prediction. The agreement is good over the intermediate range $2 \lesssim \psi \lesssim 3$, where the empirical estimates closely follow the predicted exponential trend. This confirms the central theoretical result that the effective sample size is governed by the ratio of energetic duration to oscillator period. The empirical trend

therefore supports $\psi = \ln(D_R^*/T)$ as the primary state variable organising the observed anisotropy.

Outside this range, however, systematic departures emerge. At large ψ , corresponding to short oscillator periods or long response durations, the growth of $\widehat{N}_{\text{eff}}$ slows and eventually approaches a finite ceiling. At small ψ , corresponding to long periods or short durations, the effective sample size no longer decreases and instead approaches a lower bound near three.

These departures are not random scatter around the narrowband prediction. Rather, they reveal two distinct saturation regimes that are absent from the asymptotic scaling. Figure 2 therefore serves two purposes. First, it provides direct empirical confirmation of the narrowband scaling over its range of validity. Second, it demonstrates the existence of upper and lower saturation regimes that require a more general description. These observations provide the empirical basis for the composite effective-sample-size model developed in Section 3.




Figure 2. Empirical effective sample size $\widehat{N}_{\text{eff}}$ estimated by maximum likelihood from binned values of κ_{rms} as a function of ψ . The dashed line shows the narrowband prediction $N_{\text{eff}} = 4\pi\xi e^\psi$, and the dotted horizontal line indicates $N_{\text{min}} = 3$. Three regimes are evident: a lower saturation regime at small ψ , an intermediate regime consistent with the narrowband scaling, and an upper saturation regime at large ψ .

5.2 Model calibration and interpretation of the saturation level

The model is calibrated from the individual κ_{rms} observations. With $A = 4\pi\xi$ and $N_{\text{min}} = 3$ fixed, the composite model contains a single free parameter, N_∞ . This parameter is estimated by maximising the likelihood of the full set of κ_{rms} observations across all records and oscillator periods as inferred by Equation (20). The resulting estimate is $N_\infty = 30.4$. Figure 3 compares the calibrated model with the empirical effective sample sizes of Figure 2. Although these binned estimates are not used in the calibration itself, they provide a convenient visual summary of the agreement between the model and the observations. The model reproduces the narrowband regime as well as the transition into both saturation limits.




Figure 3. Empirical effective sample size $\hat{N}_{eff}$ and fitted composite model as functions of ψ .

The theoretical prediction $N_\infty = D_R^* \cdot B_{eff,g}$ for isotropic null provides the expected saturation level. The population mean of this quantity across the dataset is approximately 59, with a harmonic mean of approximately 36, both substantially exceeding the fitted $\hat{N}_\infty \approx 30$. This indicates that real ground motions depart systematically from the ideal isotropic reference process. The fitted saturation parameter should therefore be interpreted as an effective population-level quantity rather than a direct estimator of the isotropic bandwidth limit of individual records.

The predictor ψ simultaneously governs the magnitude of the stochastic finite-sampling baseline and the detectability of physical directional structure. At small ψ , N_{eff} is small and stochastic variability is large, masking any physical directional structure. At large ψ , N_{eff} is large and the stochastic baseline becomes progressively weaker, increasing the visibility of directional structure that is not explained by finite-sample covariance fluctuations alone. The fitted N_∞ therefore characterises the effective information capacity of the observed ground-motion population rather than the theoretical isotropic limit of individual records. Residual departures from the calibrated model, examined in the following sections, provide a means of identifying systematic directional effects beyond those captured by the duration-controlled stochastic baseline. Accordingly, the calibrated model should be viewed as a population stochastic reference model rather than a pure isotropic null model.

5.3 Validation of the anisotropy distribution

Figure 4 compares the empirical histogram of κ_{rms} against the predicted mixture density at three representative oscillator periods. At $T = 0.05s$, where $\psi_{med} = 5.1$, the predicted mixture closely matches the empirical distribution. The distribution is concentrated at small κ_{rms} , with mode near 0.2, reflecting the large effective cycle count at short periods. At $T = 0.50s$, the distribution broadens and shifts rightward; the model captures the peak location well but slightly underestimates the right tail. At $T = 3.05s$, the distribution is broad and right-skewed; the model overestimates the concentration of records at large κ_{rms} . The progressive rightward shift of the distribution with increasing period, reflecting the decrease in ψ and the associated reduction in effective cycle count, is correctly reproduced by the model in all three panels.

Figure 5 presents a global comparison of the model and empirical 16th, 50th, and 84th percentiles of κ_{rms} as functions of oscillator period. The model tracks the empirical percentiles closely, confirming that the composite $N_{\text{eff}}(\psi)$ model reproduces the principal features of the observed anisotropy distribution. Residual discrepancies are expected because the model provides a duration-controlled stochastic population baseline and therefore cannot capture all sources of directional structure present in individual records.




Figure 4. Empirical distributions of κ_{rms} at three oscillator periods compared against the predicted mixture density (Equation 22). At each period, the mixture is computed from the per-record values of ψ_i through the fitted $N_{\text{eff}}(\psi)$ model. The progressive rightward shift of the distribution with increasing period reflects the decrease in effective cycle count as ψ decreases.




Figure 5. Period-dependent 16th, 50th, and 84th percentiles of κ_{rms} : model predictions (shaded band and solid black line) versus empirical estimates from the dataset (dashed blue band and solid blue line). Model percentiles are derived from the mixture cumulative distribution function (Equation 22) evaluated at each period using the record-specific ψ_i values.

5.4 Probability integral transform diagnostics


The preceding analyses assess the model in the original κ_{rms} space. A more stringent test is provided by the probability integral transform (PIT), which evaluates whether the fitted distribution correctly describes the observations across the full range of oscillator periods. For each valid observation $(\kappa_{\text{rms},i}, \psi_i)$, the Kumaraswamy cumulative distribution function is evaluated at the observed value using the record-specific shape parameter $b(\psi_i)$ derived from the fitted composite $N_{\text{eff}}(\psi)$ model:

$$U_i = F(\kappa_{\text{rms},i}; 2, b(\psi_i)) = 1 - (1 - \kappa_{\text{rms},i}^2)^{b(\psi_i)} \quad (24)$$

438

439 Under correct model specification, U_i should be uniformly distributed on (0,1). Mapping to
 440 normal scores via $z_i = \Phi^{-1}(U_i)$ yields a variable that should follow a standard normal
 441 distribution. This transformation removes the period-dependent marginal structure including
 442 boundedness, skewness, and heteroscedasticity and places all observations on a common
 443 unbounded scale suitable for subsequent physical analysis.

444 Figure 6 presents the resulting distributions of z at three representative oscillator periods and for
 445 the dataset across all valid periods. The transformed variables are approximately centred near
 446 zero and exhibit standard deviations close to unity. The aggregated distribution yields $\mu = -0.093$
 447 and $\sigma = 1.032$, indicating close but not exact agreement with the standard normal form. The
 448 close agreement with the standard normal form indicates that the composite model provides a
 449 satisfactory marginal description of κ_{rms} across the full oscillator period range. The transformed
 450 residual z therefore provides a convenient measure of anisotropy beyond the stochastic baseline.
 451 Unlike κ_{rms} , it is largely free of period-dependent scaling and bounded-support effects, making it
 452 suitable for investigating the influence of source, path, and site characteristics. This residual is
 453 used in the physical analyses that follow.



454

455 *Figure 6. Probability integral transform diagnostics for the fitted Kumaraswamy marginal model. Histograms show normal*
 456 *scores $z = \Phi^{-1}(F_{\kappa}(\kappa_{\text{rms}}; 2, b(\psi)))$ at three representative oscillator periods and collected across all valid periods. The black*
 457 *curve denotes the standard normal density.*

6. Physical regime diagnostics

Sections 3–5 showed that a large fraction of the variability in spectral anisotropy can be explained as a finite-sampling effect governed by the effective sample size $N_{\text{eff}}(\psi)$. The probability integral transform provides a convenient way to remove this stochastic baseline. The resulting normal score z_ψ measures anisotropy relative to that expected under the calibrated stochastic baseline model. Systematic departures from $N(0, 1)$ therefore indicate directional structure not captured by the stochastic baseline

The analyses that follow are diagnostic rather than predictive. Their purpose is not to develop new regression models, but to identify where physically meaningful directional structure remains after accounting for the dominant duration-controlled component of anisotropy.

6.1 Near-fault pulse effects

Near-fault directivity pulses arise from constructive interference when rupture propagates toward a recording site, producing coherent velocity transients that concentrate energy along a preferred orientation (Rupakhety et al., 2011; Somerville et al., 1997). Such records are expected to exhibit elevated κ_{rms} because the contrast between the principal response variances is amplified when energy is preferentially directed. In the following, the pulse period T_p^* is defined as the oscillator period at which pseudo-spectral velocity along the principal response axis attains its maximum (Rupakhety et al., 2011); it characterises the dominant timescale of the velocity pulse. The dataset contains 121 pulse-like records from 29 events, with a median $T_p^* = 0.87\text{s}$ and a range of 0.16–8.70 s.

The left panel of Figure 7 compares pulse-like and non-pulse records through the mean anisotropy difference

$$\Delta\kappa_{\text{rms}}(T) = \mathbb{E}[\kappa_{\text{rms}}(T) \mid \text{pulse}] - \mathbb{E}[\kappa_{\text{rms}}(T) \mid \text{non-pulse}].$$

Pulse-like records are more anisotropic at all periods. The difference is approximately 0.05 at short periods, rises sharply between $T \approx 0.1\text{ s}$ and $T \approx 1\text{ s}$, reaches a maximum of about 0.18 near $T \approx 1\text{ s}$, and gradually decreases thereafter. The location of the peak is consistent with the median pulse period of the dataset ($T_p^* = 0.87\text{ s}$), suggesting that directional effects are strongest when the oscillator period matches the dominant pulse timescale.

The elevated anisotropy of pulse-like records appears to rise from two contributions. The first is genuine directional polarization arising from rupture directivity. The second is a duration effect. Pulse-like motions concentrate their energy into a shorter time window than ordinary broadband records, resulting in smaller values of $D_R^*(T)$ and therefore smaller values of ψ , implying fewer effective samples and hence larger anisotropy even under the isotropic null.

To separate these contributions, two residual measures are considered. The first, z_ψ , is the duration-conditioned normal score defined above. Because the model already accounts for the reduced effective sample size associated with short-duration motions, z_ψ isolates anisotropy that exceeds the duration-adjusted expectation. The second, denoted z_T , is obtained by transforming κ_{rms} to normal scores separately at each oscillator period using empirical period-specific distributions (see Figure 4). This removes the overall period dependence of anisotropy

but does not condition on duration. Consequently, z_T contains both duration-mediated and polarization-mediated contributions, whereas z_ψ retains only the latter.

The right panel of Figure 7 shows the mean values of z_T and z_ψ for pulse-like records as a function of the normalised period T/T_p^* . Scaling by T_p^* collapses records with different pulse periods onto a common axis, so that $T/T_p^* = 1$ corresponds to resonance with the pulse timescale. Both residuals exhibit a clear maximum near $T/T_p^* \approx 1$, confirming that pulse-induced anisotropy is strongly period matched. The peak mean value of z_T is approximately 1.3, whereas the peak mean value of z_ψ is approximately 0.5. Thus, a substantial fraction of the observed pulse anisotropy is absorbed by the duration-conditioned stochastic baseline. The remaining signal is smaller but clearly non-zero, indicating a genuine polarization effect associated with rupture directivity.

At periods shorter than the pulse period ($T/T_p^* < 0.5$), the duration-conditioned residual becomes negative. This occurs because pulse-like records have relatively small ψ , causing the stochastic model to expect elevated anisotropy at all periods. At sub-pulse periods, however, directional concentration is weak and the observed anisotropy falls below that expectation.

The decomposition shown in Figure 7 suggests that much of the anisotropy commonly attributed to near-fault directivity is mediated indirectly through energetic duration. Directivity not only concentrates energy along a preferred direction, but also concentrates it in time, reducing the effective sample size available to estimate covariance structure. The duration-mediated contribution appears to dominate the total signal, while a smaller period-matched residual remains after conditioning on ψ . The persistence of a residual peak near $T/T_p^* = 1$ after conditioning on ψ indicates that duration alone cannot explain the pulse effect. Some component of the observed anisotropy is therefore directly associated with the directional structure of the pulse itself.

This result may help explain why the principal axis of observed near-fault anisotropy is often less stable than would be expected from a purely deterministic directivity picture. When the effective response duration is short, covariance anisotropy can be amplified by finite-sampling effects, making both its magnitude and orientation more sensitive to stochastic variability. The fault-normal direction in strike slip earthquakes may therefore not always emerge as the dominant axis in a single-record observation, even when rupture directivity is present.




Figure 7. (Left) Mean difference $\Delta\kappa_{\text{rms}}(T)$ between pulse-like ($n = 121$) and non-pulse records across oscillator period. (Right) Mean normalised anisotropy for pulse records as a function of T/T_p^* , where T_p^* is the pulse period. The T -conditioned residual z_T retains all duration and polarization contributions; the ψ -conditioned residual z_ψ isolates polarization above the stochastic baseline. Both peak near $T/T_p^* \approx 1$. Error bars are ± 1 standard error.

6.2 Source and path dependence

Earthquake magnitude and source-to-site distance are among the strongest predictors of ground-motion intensity and duration. The question addressed here is whether they exert an independent influence on spectral anisotropy, or whether their apparent effects are already captured by the duration-controlled stochastic baseline developed in Section 5. Pulse-like records are excluded from this analysis. As shown in Section 6.1, their residual anisotropy contains a directivity-related polarization component that is distinct from the source and path effects considered here. The analysis is therefore restricted to the remaining 4569 non-pulse records. For each oscillator period, the linear model

$$z_\psi(T) = c(T) + \beta_R(T)\log R_{JB} + \beta_M(T)(M_w - 5.5) + \varepsilon \quad (25)$$

is fitted by ordinary least squares. The coefficient $\beta_R(T)$ measures the sensitivity of z_ψ to the logarithm of distance, and $\beta_M(T)$ measures the sensitivity to magnitude. The coefficient of determination R^2 quantifies the fraction of total variance in z_ψ explained jointly by these two predictors. The same regression is applied to the T -conditioned residual z_T from Section 6.1, which serves as the reference.

Figure 8 compares $\hat{\beta}_R(T)$, $\hat{\beta}_M(T)$, and $R^2(T)$ between the two residual definitions. In the T -conditioned case, both fields show clear and physically interpretable period dependence. The distance coefficient $\hat{\beta}_R(T)$ is negative throughout, reaching approximately -0.8 near $T = 0.1$ s: near-field records are more anisotropic and far-field records less so. The magnitude coefficient $\hat{\beta}_M(T)$ is negative at short periods and increases monotonically with period, reaching approximately $+0.3$ at long periods: larger earthquakes are associated with reduced short-period anisotropy and enhanced long-period anisotropy. These patterns are broadly consistent with the known dependence of ground-motion duration on earthquake magnitude and source-to-site distance (Afshari and Stewart, 2016). Larger earthquakes and more distant records generally produce longer-duration response, corresponding to larger $D_R^*(T)$, larger ψ , and therefore larger N_{eff} , which reduces the stochastic baseline anisotropy predicted by the model.

The median R^2 for z_T is 0.038, indicating that magnitude and distance together explain approximately 4% of the variance in the period-conditioned residual. However, after ψ -conditioning, both coefficient fields collapse substantially. The median R^2 decreases from 0.038 to only 0.0014. This near-complete collapse in explanatory power has a direct physical interpretation. The apparent dependence of κ_{rms} on source and path predictors appears to be overwhelmingly mediated through energetic duration and effective sample size. In effect, the primary pathway is $(M_w, R_{JB}) \rightarrow D_R^*(T) \rightarrow \psi \rightarrow N_{\text{eff}} \rightarrow \kappa_{\text{rms}}$. Magnitude and distance influence κ_{rms} mainly through their control on energetic duration, which determines the stochastic baseline anisotropy through $N_{\text{eff}}(\psi)$. Once this is accounted for by the ψ -conditioned model, very little direct dependence on M_w or R_{JB} remains.

A residual distance sensitivity survives ψ -conditioning across the short-period range. At the shortest periods, $\hat{\beta}_R(T)$ for both z_T and z_ψ converge to approximately -0.4 , indicating that ψ -conditioning absorbs little of the distance effect there. At periods longer than approximately $T \approx 0.1$ s, most of the distance dependence is absorbed by duration. The origin of the remaining short-period distance effect is not resolved by the present analysis. Possible contributing mechanisms include effects related to the frequency content and coherence of near-field ground motion that are not captured by the scalar measure $N_{\text{eff}}(\psi)$. Back-transforming $\hat{\beta}_R \approx -0.4$ in z_ψ space to κ_{rms} units yields a conditional mean shift of approximately 0.03 – 0.05 . Despite its systematic character, however, the effect remains weak in practical terms.




Figure 8. Period-dependent sensitivity of normalised anisotropy to source and path predictors. Shown are the distance coefficient $\hat{\beta}_R(T)$, magnitude coefficient $\hat{\beta}_M(T)$, and coefficient of determination $R^2(T)$ from regression of the T -conditioned residual z_T (grey) and the ψ -conditioned residual z_ψ (blue) on $\log R_{JB}$ and M_w . ψ -conditioning reduces the median R^2 from 0.038 to 0.001 , indicating that source and path dependence of κ_{rms} is largely mediated by ground motion duration.

6.3 Station persistence and site amplification

Substantial variability remains in z_ψ after excluding pulse records and establishing that source and path predictors explain negligible additional variance. To assess whether this remaining variability is structured by recording site or seismic event, we decompose the total variance of z_ψ at each oscillator period into three components: a between-station component capturing persistent differences among recording sites across multiple events, a between-event component capturing persistent differences among earthquakes across multiple stations, and a record-level component representing the variability attributable to neither. Formally:

$$z_\psi(T) = \mu(T) + u_s(T) + v_e(T) + \eta_r(T) \quad (26)$$

where $u_s(T)$ and $v_e(T)$ are independent station and event random intercepts and $\eta_r(T)$ is the record-level residual. Variance components are estimated at each period by restricted maximum likelihood (REML).

Figure 9(a) shows the fraction of total variance attributed to each component. Station-level variability accounts for approximately 25–30% of the total at short periods and declines steadily with increasing period, becoming small beyond $T \approx 1$ s. Event-level variability is negligible throughout the full period range. This is consistent with the weak residual dependence of anisotropy on magnitude and distance after ψ -conditioning observed in Section 6.2. Record-to-record variability dominates at all periods. Figure 9(b) shows that the total variance of $\kappa_{\text{rms}}(T)$

increases monotonically with oscillator period, consistent with the reduction in effective sample size at long periods arising from the decrease of ψ .




Figure 9. Period-dependent variance decomposition of z_ψ for non-pulse records. (a) Fraction of total variance attributed to station, event, and record-level components estimated by crossed random-effects REML. Station-level variability is most pronounced at short periods; event-level variability is negligible throughout. (b) Total variance of $\kappa_{\text{rms}}(T)$ as a function of oscillator period.

The concentration of station variance at short periods motivates examination of site-specific predictors. We first test whether the fundamental resonance frequency f_0 derived from microtremor HVSr curves, available for 421 stations representing 1,126 records, carries predictive power. Period-wise regressions of z_ψ against $\log f_0$ show no consistent period-dependent structure, indicating that the resonance frequency is not a useful predictor of anisotropy in this dataset.

We therefore consider whether the strength of site amplification at the oscillator frequency rather than the resonance frequency itself is the relevant predictor. For each record and oscillator period T , a period-matched HVSr amplitude is defined as:

$$A_{\text{HVSr}}(T) = \frac{1}{|\mathcal{B}_T|} \int_{\mathcal{B}_T} \log \text{HVSr}(f) df \quad (27)$$

where $\mathcal{B}_T = [f_T/\sqrt{2}, f_T\sqrt{2}]$ is a one-octave band centred on $f_T = 1/T$. Stations are stratified into low, mid, and high tertiles based on their station-mean $A_{\text{HVSr}}(T_{\text{ref}})$ at the short-period reference $T_{\text{ref}} = 0.10\text{s}$.

Figure 10 shows the mean z_ψ and mean κ_{rms} for each tertile as functions of oscillator period. A clear separation is present at short periods: stations with higher short-period HVSr amplification exhibit systematically elevated z_ψ , indicating that their κ_{rms} exceeds what the isotropic null model predicts. The separation between high and low tertiles reaches approximately 0.3 z-units, corresponding to a difference of approximately 0.05–0.08 in κ_{rms} . The separation diminishes progressively with period and becomes negligible beyond $T \approx 2\text{s}$. The same ordering is confirmed directly in $\kappa_{\text{rms}}(T)$ in Figure 10(b), ruling out a transformation artifact. This behaviour is consistent with the variance decomposition of Figure 9, which identified station-to-station variability as the dominant structured component of the residual anisotropy at short periods.

The observed association is physically plausible. Directionally asymmetric site effects arising, for example, from lateral velocity contrasts, basin geometry (Baena-Rivera et al., 2024), or anisotropic near-surface structure (Pischiutta et al., 2018) may contribute to both elevated short-period amplification and increased covariance anisotropy. The confinement to short periods is consistent with the known frequency dependence of site effects: shallow site structure governs ground motion variability primarily at high frequencies, while long-period response is less sensitive to near-surface heterogeneity.

Unlike magnitude and distance, whose influence is largely absorbed by the duration-controlled stochastic baseline, site amplification leaves a measurable residual signature. However, the effect explains only part of the total variance, and HVSr measurements are available for only 421 of the 1304 stations in the dataset. The analysis should therefore be viewed as diagnostic rather than predictive. It identifies a plausible physical mechanism contributing to short-period anisotropy not explained by the duration-controlled stochastic baseline, but does not support the introduction of a site-correction model with the present data.




Figure 10. Stratification of stations by short-period HVSr amplitude. Stations are grouped into low, mid, and high tertiles based on station-mean $A_{HVSr}(T_{ref})$ at $T_{ref} = 0.10$ s. (a) Mean z_ψ for each group. (b) Corresponding mean $\kappa_{rms}(T)$. Higher short-period HVSr amplification is associated with elevated anisotropy at short periods; the effect diminishes toward longer periods.

7 Practical simulations of directional spectra

While the stochastic model for κ_{rms} presented earlier provides a probabilistic description of directional anisotropy, practical applications ultimately require a model for the directional response spectrum itself. As discussed in Section 2.2, the directional spectrum may be represented as the product of two components: an RMS ellipse governed by κ_{rms} , and a directional excess peak-factor field $\varepsilon(\theta)$. The present work provides a stochastic model for the former, while Rupakhety and Hernández-Aguirre (2026b) provide a simple stochastic representation of $\varepsilon(\theta)$. The objective of this section is to examine how far a realistic directional spectrum can be reproduced when combining both models. Any remaining discrepancies then provide insight into the relative contributions of anisotropy and peak-factor fluctuations to directional spectral variability.

7.1 Directional spectral simulations

Taking the square root of Eq. (4) and normalising by the RotD50 amplitude yields a direct expression for the directional PSA field relative to its rotation-invariant median level:

$$\frac{\text{PSA}(\theta, T)}{\text{RotD}_{50}(T)} = \sqrt{1 + \kappa_{\text{rms}}(T) \cos 2(\theta - \theta_0)} \exp(\varepsilon(\theta, T)) \quad (28)$$

The right-hand side separates into two stochastic components. In the following simulations, κ_{rms} is drawn from the Kumaraswamy $(2, b)$ distribution with $b = (N_{\text{eff}}(\psi) - 1)/2$. Following Rupakhety and Hernández-Aguirre (2026a), the peak-factor residual field is represented to first order by a single harmonic,

$$\varepsilon(\theta, T) \approx a_2(T) \cos 2\theta + b_2(T) \sin 2\theta \quad (29)$$

where the coefficients (a_2, b_2) are isotropic and normally distributed with identical variances at the population level: $(a_2, b_2) \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$. The harmonic amplitude $\rho = \sqrt{a_2^2 + b_2^2}$ is therefore Rayleigh-distributed with scale parameter $\sigma = 0.09$, reflecting the typical magnitude of azimuthal peak-factor fluctuations observed across the ESM database.

Together, these two components provide a first-order probabilistic representation of the directional PSA field anchored to RotD50. To generate a realisation at a given period, one (i) computes ψ (ii) evaluates $N_{\text{eff}}(\psi)$ and draws κ_{rms} from Kumaraswamy $(2, b)$, (iii) draws (a_2, b_2) independently from $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$, and (iv) evaluates Eq. (28) at the desired azimuths.

7.2 Illustrative examples

Three non-pulse-like records are selected from the ESM database to span the range of directional variability: one record whose κ_{rms} at $T = 1\text{s}$ falls near the 33rd percentile of the population, one near the 67th, and one near the median of the upper tertile. For each record and each of three target periods ($T = 0.1, 1, 3\text{s}$), the observed directional PSA field is aligned to the principal axis θ_0 of the response covariance matrix, normalised by RotD50, and compared against 500 Monte Carlo realisations conditioned solely on the observed ψ . The results are compared in Figure 11. The peak-factor residual field $\varepsilon(\theta; T)$ is displayed in the central inset as $1 + 2.5 \varepsilon(\theta)$ so that azimuthal fluctuations of typical amplitude $\sigma = 0.09$ are legible at the inset scale.

The most immediate observation is how strongly period governs the width of the simulation band. This behaviour emerges almost entirely from the period dependence of N_{eff} , without requiring any period-dependent tuning of the peak-factor field. At $T = 0.1\text{s}$, ψ exceeds 4.4 for all three records, placing N_{eff} near its saturation value N_{∞} . The Kumaraswamy distribution for κ_{rms} is consequently narrow and skewed toward zero, and the band is nearly circular across all three κ_{rms} tertiles. By $T = 1\text{s}$ and $T = 3\text{s}$, ψ has fallen by one to two units, N_{eff} is well below saturation, and the band has widened and elongated.

The high- κ_{rms} row reveals a progressive departure between the observed field and the model that becomes instructive on closer reading. Where ψ is large and κ_{rms} is small (the upper left panel), the model median tracks the observed field closely, the peak-factor residual inset is nearly circular, and the direction of maximum PSA coincides with the major axis of the RMS ellipse. This is the regime in which the isotropic residual assumption is well supported: the peak-factor field

adds only small, approximately symmetric fluctuations on top of the geometric anisotropy, and the RMS axis is a reliable predictor of the maximum-response direction.

As the period increases and κ_{rms} grows, this alignment deteriorates. The residual field $\varepsilon(\theta)$ develops pronounced azimuthal structure whose amplitude is concentrated away from the RMS major axis, rotating the direction of maximum PSA away from θ_0 by an amount that varies from record to record. In the most anisotropic cases, the residual field displays a multi-lobe pattern with significant energy in harmonics beyond the first, none of which are aligned with the RMS axes. The consequence is visible in the observed PSA field itself: a notched suppression perpendicular to the maximum-response direction that the single-harmonic, isotropic model cannot reproduce. The model median, which averages over an isotropic ensemble of ε realisations, remains close to the RMS ellipse and misses this structure.




Figure 11. Aligned directional PSA fields for three representative non-pulse-like records. Columns correspond to oscillator periods $T = 0.1$, 1 , and 3 s. In each panel the observed field $\text{PSA}(\theta)/\text{RotD50}$, aligned to the principal axis θ_0 , is shown in black; the RMS ellipse in red dashed; and the model median and 16th–84th percentile band from 500 Monte Carlo realisations conditioned on the observed ψ in blue. The central inset shows the peak-factor residual field displayed as $1 + 2.5 \varepsilon(\theta)$: the unit shift centres the field on the RotD50-normalised reference circle (dotted), and the scale factor of 2.5 amplifies the small azimuthal fluctuations whose typical amplitude is $\sigma = 0.09$ to make them visually distinguishable from the unit circle. Grey and blue curves show the observed and median simulated residual fields respectively. ESM event identifier and recording station code are given in the right margin; values of κ_{rms} and ψ at each period are annotated below each panel.

These observations point to three related limitations of the first-order model. First, although the peak-factor residual field is approximately isotropic at the population level (Rupakhety and Hernández-Aguirre, 2026b), individual records can exhibit strong azimuthal organisation, particularly at long periods and large κ_{rms} . Second, the alignment of the residual field with the RMS axes is period- and κ_{rms} -dependent, suggesting a coupling between the two components. Third, multi-lobe patterns require higher harmonics, and the relative importance of those harmonics appears to grow with κ_{rms} .

7.3 Dataset-level distributions

Figure 12 compares the empirical and simulated distributions of two population-level quantities at three oscillator periods. The top row shows the marginal distribution of $\text{PSA}(T, \theta)/\text{RotD50}(T)$ aggregated over all valid non-pulse records and all 180 azimuths (1° increment); the bottom row shows the distribution of $\text{RotD100}(T)/\text{RotD50}(T)$. In both rows the observed distribution is shown as a filled histogram and the model as a smooth kernel density estimate conditioned on the observed ψ values, using the same Monte Carlo procedure as described in Section 7.1.

The top row shows two features that evolve with period. First, the distribution widens progressively from $T = 0.1\text{s}$ to $T = 3\text{s}$, reflecting the ψ -mediated growth of κ_{rms} as N_{eff} falls below its saturation value N_∞ . The model reproduces this widening accurately at short periods, where κ_{rms} is small and the isotropic residual assumption is well supported. Second, the observed distribution develops a pronounced right-side asymmetry at longer periods: the upper tail is increasingly truncated as it approaches the $\sqrt{2}$ limiting case of perfectly unidirectional motion, while the lower tail remains relatively extended.

The present model approaches but cannot enforce the $\sqrt{2}$ limit because the peak-factor residual $\varepsilon(\theta)$ is drawn independently of κ_{rms} . When κ_{rms} is close to unity the geometric term $[1 + \kappa_{\text{rms}} \cos 2\theta]^{1/2}$ already approaches $\sqrt{2}$ at the major axis; any positive ε at that azimuth then pushes the simulated $\text{PSA}/\text{RotD50}$ above $\sqrt{2}$. The independence assumption between κ_{rms} and ε is a likely contributor to the ceiling violation. A more sophisticated residual model that conditions ε on anisotropy may alleviate this behaviour.

The bottom row shows that the $\text{RotD100}/\text{RotD50}$ distribution broadens and flattens with period, driven by the increasing record-to-record variance in κ_{rms} as N_{eff} moves away from saturation. The model captures this progressive broadening well. The 84th percentile of the observed distribution rises from 1.31 at $T = 0.1\text{s}$ to 1.36 at $T = 3\text{s}$, remaining comfortably below the $\sqrt{2}$ ceiling throughout. The model tracks the observed 84th percentile closely at short periods (1.30 vs. 1.31 at $T = 0.1\text{s}$) but overestimates it at longer periods (1.44 vs. 1.36 at $T = 3\text{s}$), a bias of approximately 6% that grows with period as κ_{rms} increases. Both the growing asymmetry in the top row and the 84th-percentile bias in the bottom row share the same origin: the factored, independent treatment of κ_{rms} and ε breaks the joint constraint imposed by the two-dimensional nature of horizontal ground motion.




Figure 12. Distributions of $\text{PSA}(T, \theta)/\text{RotD50}(T)$ (top) and $\text{RotD100}(T)/\text{RotD50}(T)$ (bottom) at three oscillator periods for all non-pulse-like ESM records: observed (grey histogram, open-ended bins) and model (blue KDE). The dash-dot line marks the $\sqrt{2}$ physical ceiling; dashed vertical lines in the bottom row mark the 84th percentile, with values annotated alongside.

Figure 13 extends this comparison continuously across all 100 log-spaced periods from 0.01 to 10s. At short periods both panels show close agreement between observed and model bands. As period increases, the model band widens faster than the observed in both panels: in the left panel the 84th and 16th percentile curves diverge symmetrically above and below the unit median, while in the right panel the model median and 84th percentile rise progressively above their observed counterparts, with the model 84th percentile exceeding $\sqrt{2}$ at very long periods while the observed remains below it. The bias is negligible at short periods. It is below 1% at $T = 0.1\text{s}$ and grows steadily before saturating beyond $T \approx 5\text{s}$. At long periods the model overestimates the 84th percentile of $\text{PSA}(T, \theta)/\text{RotD50}(T)$ by approximately 3%, the median $\text{RotD100}/\text{RotD50}$ by approximately 3-4%, and the 84th percentile $\text{RotD100}/\text{RotD50}$ by approximately 6-7%.

The ceiling violation is informative because it identifies the principal limitation of the present factorised simulation framework rather than of the anisotropy model itself. The covariance-anisotropy model preserves the bounded geometry of the RMS response field. Exceedance occurs only after an independent peak-factor residual field is superimposed. This indicates that the remaining discrepancy is associated primarily with the representation of directional peak formation rather than with the duration-controlled anisotropy model developed in Sections 3–5. Importantly, the field-based formulation retains an explicit directional representation and therefore admits direct enforcement of geometric constraints such as the $\sqrt{2}$ ceiling. Future refinements of the peak-factor model, including higher-order harmonic representations and dependence between covariance anisotropy and peak-factor structure, provide a natural route to improving the tail behaviour while preserving the parsimonious character of the framework.

The present simulation framework is intentionally parsimonious, involving only two calibrated parameters: $N_\infty = 30.45$ and $\sigma = 0.09$. Despite its simplicity, it reproduces the principal features of directional spectral variability, including the period-dependent growth of anisotropy and the distribution of RotD amplification factors. The remaining discrepancies are confined primarily to the extreme upper tail and are attributable to the simplified treatment of the peak-factor residual field rather than to the anisotropy model itself. A practical engineering remedy is to impose the physical ceiling $\text{RotD100}/\text{RotD50} \leq \sqrt{2}$, while more refined representations of $\varepsilon(\theta)$ may further improve tail behaviour in future applications.




Figure 13. Percentile spectra of $\text{PSA}(T, \theta)/\text{RotD50}(T)$ (left) and $\text{RotD100}(T)/\text{RotD50}(T)$ (right) across all non-pulse-like records. Shaded bands span the 16th–84th percentile; solid lines show the median; dashed lines show the 16th and 84th percentile boundaries. Grey: observed; blue: model conditioned on observed ψ . The dash-dot line marks the $\sqrt{2}$ physical ceiling.

8. Conclusions

This study developed a probabilistic model for covariance anisotropy in directional response spectra based on the effective sample size N_{eff} . Analysis of 4690 strong-motion records from the ESM database demonstrated that the duration-to-period ratio $\psi = \ln(D_R^*/T)$ acts as the primary organising variable of anisotropy. Empirical estimates of N_{eff} exhibit the three regimes predicted by the theory: a long-period floor, an intermediate narrowband scaling regime, and a short-period saturation regime associated with the finite bandwidth of the excitation. A composite model containing a single calibrated saturation parameter reproduced these trends and the observed anisotropy distributions across the full oscillator-period range considered.

The calibrated model provides a duration-controlled stochastic baseline for covariance anisotropy. Despite containing only a single calibrated population parameter, it reproduces the principal features of the observed κ_{rms} distributions, including the progressive increase in median anisotropy and the associated broadening of the distribution with increasing period. Probability-integral-transform diagnostics showed close agreement with the target standard-normal distribution, with the aggregate transformed residual distribution exhibiting $\mu \approx -0.09$ and $\sigma \approx 1.03$. These results indicate that a large fraction of the observed variability in directional anisotropy can be explained as a consequence of finite covariance sampling.

Residual analyses provided insight into the physical structure remaining after removal of this stochastic baseline. The apparent dependence of anisotropy on earthquake magnitude and source-to-site distance was found to be largely duration-mediated. After conditioning on ψ , the explanatory power of magnitude and distance collapsed from a median R^2 value of approximately 0.038 to approximately 0.001. This result suggests that source and path effects influence anisotropy primarily through their control on energetic duration and therefore on the effective number of independent covariance observations. Near-fault pulse-like motions remain an important exception. Conditioning on duration removes a substantial fraction of the pulse-induced anisotropy, but a smaller period-matched residual persists, indicating a genuine directivity-related polarization component. Site-related effects also leave a detectable short-period residual signature. Station-level variability accounts for approximately 25–30% of the residual variance at short periods, and stations with stronger short-period HVSr amplification exhibit systematically elevated anisotropy relative to the stochastic baseline.

The proposed anisotropy model was combined with a previously developed stochastic representation of directional peak-factor variability to generate simulations of directional response spectra. Using only two calibrated population parameters, N_∞ and σ_p , the resulting framework reproduces the principal population-level characteristics of directional spectral variability and RotD amplification factors. Remaining discrepancies are concentrated in the upper tail at long periods, where the assumption of independence between covariance anisotropy and peak-factor variability leads to occasional violations of the physical $\sqrt{2}$ ceiling. Importantly, this limitation arises from the representation of the peak-factor field rather than from the covariance-anisotropy model itself. Because the directional field remains represented explicitly, geometric constraints such as the $\sqrt{2}$ ceiling can be imposed directly and provide a clear target for future refinements.

The principal implication of this study is a reinterpretation of directional variability. The results suggest that directional anisotropy is governed predominantly by duration-controlled covariance sampling, with physical effects such as rupture directivity, radiation-pattern asymmetry, and site amplification acting as residual modifications to that stochastic baseline. This perspective provides a physically grounded alternative to treating directionality solely through conversion factors or additional aleatory variability terms.

Two developments appear particularly important for future work. The first is the development of predictive models for the energetic response duration $D_R^*(T)$. In the present study, $D_R^*(T)$ is computed directly from recorded response histories, but practical implementation within simulation frameworks and probabilistic seismic hazard analysis will require predictive relationships linking $D_R^*(T)$ to source, path, site, and oscillator characteristics. The second is the development of joint models for covariance anisotropy and directional peak-factor structure. The simulation results suggest that dependence between these two components, together with higher-order harmonic representations of the peak-factor field, may substantially improve the representation of extreme directional response while preserving the low-dimensional character of the framework. Together, these developments point toward a compact stochastic representation of directional ground-motion fields suitable for future simulation and hazard-analysis applications.

Acknowledgements: RR acknowledges support from the University of Iceland Research Fund

Author contributions: RR: mathematical framework, modelling, derivation, coding, analysis, drafting and reviewing, VHMA: data collection and processing, drafting, editing, and reviewing

REFERENCES

- Afshari, K., Stewart, J.P., 2016. Physically Parameterized Prediction Equations for Significant Duration in Active Crustal Regions. *Earthq. Spectra* 32, 2057–2081. <https://doi.org/10.1193/063015EQS106M>
- Baena-Rivera, M., Arciniega-Ceballos, A., Sánchez-Sesma, F.J., Rosado-Fuentes, A., Pardo-Dañino, J.C., 2024. Directional HVSr at the Chalco lakebed zone of the Valley of Mexico: Analysis and interpretation. *J. Appl. Geophys.* 228, 105452. <https://doi.org/10.1016/j.jappgeo.2024.105452>
- Baker, J.W., Cornell, C.A., 2006. Which Spectral Acceleration are you Using? *Earthq. Spectra* 22, 293–312. <https://doi.org/10.1193/1.2191540>
- Beyer, K., Bommer, J.J., 2006. Relationships between median values and between aleatory variabilities for different definitions of the horizontal component of motion. *Bull. Seismol. Soc. Am.* 96, 1512–1522.
- Boore, D.M., 2010. Orientation-Independent, Nongeometric-Mean Measures of Seismic Intensity from Two Horizontal Components of Motion. *Bull. Seismol. Soc. Am.* 100, 1830–1835. <https://doi.org/10.1785/0120090400>
- Boore, D.M., Kishida, T., 2017. Relations between some horizontal-component ground-motion intensity measures used in practice. *Bull. Seismol. Soc. Am.* 107, 334–343.
- Boore, D.M., Stewart, J.P., Seyhan, E., Atkinson, G.M., 2014. NGA-West2 Equations for Predicting PGA, PGV, and 5% Damped PSA for Shallow Crustal Earthquakes. *Earthq. Spectra* 30, 1057–1085. <https://doi.org/10.1193/070113EQS184M>
- Bozorgnia, Y., Abrahamson, N.A., Atik, L.A., Ancheta, T.D., Atkinson, G.M., Baker, J.W., Baltay, A., Boore, D.M., Campbell, K.W., Chiou, B.S. -J., Darragh, R., Day, S., Donahue, J., Graves, R.W., Gregor, N., Hanks, T., Idriss, I.M., Kamai, R., Kishida, T., Kottke, A., Mahin, S.A., Rezaeian, S., Rowshandel, B., Seyhan, E., Shahi, S., Shantz, T., Silva, W., Spudich, P., Stewart, J.P., Watson-Lamprey, J., Wooddell, K., Youngs, R., 2014. NGA-West2 Research Project. *Earthq. Spectra* 30, 973–987. <https://doi.org/10.1193/072113EQS209M>
- Bradley, B.A., Baker, J.W., 2015. Ground motion directionality in the 2010–2011 Canterbury earthquakes. *Earthq. Eng. Struct. Dyn.* 44, 371–384. <https://doi.org/10.1002/eqe.2474>
- Feng, R., Yuan, W., Sextos, A., 2021. Probabilistic loss assessment of curved bridges considering the effect of ground motion directionality. *Earthq. Eng. Struct. Dyn.* 50, 3623–3645. <https://doi.org/10.1002/eqe.3525>
- Girmay, N., Miranda, E., 2026. Component-To-Component Variability of Response Spectral Ordinates from Strike-Slip Earthquakes. *Earthq. Spectra* 42, e70045. <https://doi.org/10.1002/esp4.70045>
- Hong, H.P., Goda, K., 2007. Orientation-Dependent Ground-Motion Measure for Seismic-Hazard Assessment. *Bull. Seismol. Soc. Am.* 97, 1525–1538. <https://doi.org/10.1785/0120060194>
- Huang, Y.-N., Whittaker, A.S., Luco, N., 2008. Maximum Spectral Demands in the Near-Fault Region. *Earthq. Spectra* 24, 319–341. <https://doi.org/10.1193/1.2830435>
- Kotha, S.R., Cotton, F., Bindi, D., 2019. Empirical Models of Shear-Wave Radiation Pattern Derived from Large Datasets of Ground-Shaking Observations. *Sci. Rep.* 9, 981. <https://doi.org/10.1038/s41598-018-37524-4>
- Kumaraswamy, P., 1980. A generalized probability density function for double-bounded random processes. *J. Hydrol.* 46, 79–88. [https://doi.org/https://doi.org/10.1016/0022-1694\(80\)90036-0](https://doi.org/https://doi.org/10.1016/0022-1694(80)90036-0)

- Lagaros, N.D., 2010. The impact of the earthquake incident angle on the seismic loss estimation. *Eng. Struct.* 32, 1577–1589. <https://doi.org/10.1016/j.engstruct.2010.02.006>
- Lanzano, G., Luzi, L., Cauzzi, C., Bienkowski, J., Bindi, D., Clinton, J., Cocco, M., D’Amico, M., Douglas, J., Faenza, L., Felicetta, C., Gallovic, F., Giardini, D., Ktenidou, O., Lauciani, V., Manakou, M., Marmureanu, A., Maufroy, E., Michelini, A., Özener, H., Puglia, R., Rupakhety, R., Russo, E., Shahvar, M., Sleeman, R., Theodoulidis, N., 2021. Accessing European Strong-Motion Data: An Update on ORFEUS Coordinated Services. *Seismol. Res. Lett.* 92, 1642–1658. <https://doi.org/10.1785/0220200398>
- Muirhead, R.J., 2009. Aspects of multivariate statistical theory. John Wiley & Sons.
- Nievas, C., Sullivan, T., 2018. A multidirectional conditional spectrum. *Earthq. Eng. Struct. Dyn.* 47, 945–965. <https://doi.org/10.1002/eqe.3000>
- Pischiutta, M., Cianfarra, P., Salvini, F., Cara, F., Vannoli, P., 2018. A systematic analysis of directional site effects at stations of the Italian seismic network to test the role of local topography. *Geophys. J. Int.* 214, 635–650. <https://doi.org/10.1093/gji/ggy133>
- Poulos, A., Miranda, E., 2025. Accounting for ground motion directionality and building orientations in urban seismic risk analysis. *Earthq. Spectra* 41, 1780–1800. <https://doi.org/10.1177/87552930251315751>
- Poulos, A., Miranda, E., 2023. Effect of Style of Faulting on the Orientation of Maximum Horizontal Earthquake Response Spectra. *Bull. Seismol. Soc. Am.* 113, 2092–2105. <https://doi.org/10.1785/0120230001>
- Poulos, A., Miranda, E., 2022. Probabilistic Characterization of the Directionality of Horizontal Earthquake Response Spectra. *Earthq. Eng. Struct. Dyn.* 51, 2077–2090. <https://doi.org/10.1002/eqe.3654>
- Rupakhety, R., 2026. When Rotation-Invariant Spectra Enter Structural Design: Interpreting RotD100 and Related Measures. *Earthq. Spectra* 42, e70081. <https://doi.org/10.1002/esp4.70081>
- Rupakhety, R., Hernández-Aguirre, V.M., 2026a. An Anisotropy-Aware Framework for Directional Response Spectra and Rotation-Invariant Ground-Motion Measures. *Earthq. Eng. Struct. Dyn.* <https://doi.org/10.1002/eqe.70293>
- Rupakhety, R., Hernández-Aguirre, V.M., 2026b. Directional peak factors of strong-motion response spectra: a stochastic field representation on the circle. *Bull. Earthq. Eng.* <https://doi.org/10.1007/s10518-026-02613-7>
- Rupakhety, R., Hernández-Aguirre, V.M., 2026. On the Origin of Directional Variability in Earthquake Response Spectra: A Stochastic Covariance Framework. *Earthq. Eng. Struct. Dyn.* eqe.70287. <https://doi.org/10.1002/eqe.70287>
- Rupakhety, R., Sigurdsson, S.U., Papageorgiou, A.S., Sigbjörnsson, R., 2011. Quantification of ground-motion parameters and response spectra in the near-fault region. *Bull. Earthq. Eng.* 9, 893–930. <https://doi.org/10.1007/s10518-011-9255-5>
- Shahi, S.K., Baker, J.W., 2014a. NGA-West2 Models for Ground Motion Directionality. *Earthq. Spectra* 30, 1285–1300. <https://doi.org/10.1193/040913EQS097M>
- Shahi, S.K., Baker, J.W., 2014b. An efficient algorithm to identify strong-velocity pulses in multicomponent ground motions. *Bull. Seismol. Soc. Am.* 104, 2456–2466.
- Skoulidou, D., Romão, X., Franchin, P., 2019. How is collapse risk of RC buildings affected by the angle of seismic incidence? *Earthq. Eng. Struct. Dyn.* 48, 1575–1594. <https://doi.org/https://doi.org/10.1002/eqe.3214>
- Somerville, P.G., Smith, N.F., Graves, R.W., Services, W.-C.F., Abrahamson, N.A., 1997. Modification of Empirical Strong Ground Motion Attenuation Relations to Include the Amplitude and Duration Effects of Rupture Directivity. *Seismol. Res. Lett.* 68, 199–222.
- Spudich, P., Hellweg, M., Lee, W.H.K., 1996. Directional topographic site response at Tarzana observed in aftershocks of the 1994 Northridge, California, earthquake: Implications for

947 mainshock motions. Bull. Seismol. Soc. Am. 86, S193–S208.
948 <https://doi.org/10.1785/BSSA08601BS193>
949 Vanmarcke, E.H., Lai, S.-S.P., 1980. Strong-motion duration and RMS amplitude of earthquake
950 records. Bull. Seismol. Soc. Am. 70, 1293–1307.
951