

Online Supplementary Material

A Log-Gaussian Cox process framework for zone-free probabilistic seismic hazard assessment in Sumatra

Irwan Endrayanto Aluicius Nanang Susyanto Wiwit Suryanto
Dwi Ertiningsih Fajar Adi-Kusumo

Contents

S1 Thomas Process: Minimum Contrast Algorithm	2
S2 Gardner–Knopoff Declustering Algorithm	2
S3 Stochastic Declustering: ETAS Background Comparison	3
S4 Omori–Utsu MLE Algorithm	3
S5 Spatial ETAS Grid-Search MLE	3
S5.1 Anisotropic ETAS Kernel	4
S6 PSHA Vectorised Algorithm	4
S7 RELM Eight-Step Algorithm	5
S8 Fault-Distance Covariate: Megathrust Proxy and Covariate Confounding	7
S9 Engineering Hazard Products and Code Reconciliation	7

S1 Thomas Process: Minimum Contrast Algorithm

The following algorithm implements the minimum contrast estimation of the Thomas process parameters (κ, σ) described in Section 3.3 of the main text.

Input: Declustered mainshock locations $\{\mathbf{s}_i\}_{i=1}^{n_K}$ ($M \geq 5.0$, $n_K = 1,638$); window $\mathcal{W}$; integration limits $r_{\min} = 2$ km, $r_{\max} = 400$ km; r -grid of 200 equally spaced values.

Step 1. Compute the non-parametric K -function estimate $\hat{K}(r)$ using the isotropy edge-correction weights $e_{ij}(r) = |\mathcal{W}|/|\mathcal{W}_{ij}(r)|$ on the 200-point r -grid (Equation 3 of main text).

Step 2. Compute the theoretical Thomas K -function $K_{\text{Th}}(r; \kappa, \sigma)$ on the same grid for a 50×50 parameter grid: $\kappa \in [10^{-9}, 10^{-5}]$ km $^{-2}$ (log-spaced) and $\sigma \in [0.5, 200]$ km (log-spaced).

Step 3. For each (κ, σ) pair evaluate the contrast

$$Q(\kappa, \sigma) = \int_{r_{\min}}^{r_{\max}} [\hat{K}(r)^{1/2} - K_{\text{Th}}(r; \kappa, \sigma)^{1/2}]^2 dr,$$

approximated by the trapezoidal rule.

Step 4. Return $(\hat{\kappa}, \hat{\sigma}) = \arg \min Q$.

Step 5. Compute $\hat{\nu} = \hat{\lambda}/\hat{\kappa}$ where $\hat{\lambda} = n_K/|\mathcal{W}|$.

Fitted values: $\hat{\kappa} = 8.47 \times 10^{-7}$ km $^{-2}$, $\hat{\sigma} = 3.2$ km, $\hat{\nu} \approx 1,000$.

S2 Gardner–Knopoff Declustering Algorithm

Declustering follows Gardner and Knopoff (1974) with the magnitude-dependent space–time windows:

$$L(M) = 10^{0.1238M+0.983} \text{ km}, \quad T(M) = \begin{cases} 10^{0.5409M-0.547} \text{ days} & M < 6.5, \\ 10^{0.7507M-1.670} \text{ days} & M \geq 6.5. \end{cases}$$

Step 1. Sort the 5,356-event catalogue by origin time.

Step 2. For each event i in decreasing magnitude order:

- (a) If event i is already flagged as an aftershock, skip.
- (b) Flag as aftershock every event $j \neq i$ satisfying $|\mathbf{s}_j - \mathbf{s}_i| \leq L(M_i)$ and $|t_j - t_i| \leq T(M_i)$.

Step 3. Events not flagged are mainshocks.

Result: 693 mainshocks, 4,663 aftershocks removed (87% reduction).

The exclusion window for $M = 9.1$ reaches $L \approx 840$ km and $T \approx 14$ yr, so all aftershocks of the 2004 megathrust within the study window are removed.

S3 Stochastic Declustering: ETAS Background Comparison

As a sensitivity check, the stochastic declustering method of Zhuang et al. (2002) was applied using an ETAS background model fitted to the full catalogue. Each event i is assigned a continuous probability p_i of belonging to the background:

$$p_i = \frac{\mu(\mathbf{s}_i)}{\lambda(\mathbf{s}_i, t_i)},$$

where μ is the spatially uniform background rate and λ is the full ETAS conditional intensity. Events with $p_i > 0.5$ are retained as background.

Result: 631 background events (versus 693 by G–K). The 9% difference arises because stochastic declustering assigns partial membership probabilities near large mainshocks, retaining more events in those zones. Both counts are carried through the LGCP fitting; the hazard results change by less than 5% between the two declusterings, well within the GMPE epistemic band.

S4 Omori–Utsu MLE Algorithm

Omori–Utsu parameters (p, c) are estimated by conditional log-likelihood maximisation on the full 5,356-event catalogue, applied separately to each identified mainshock sequence.

For mainshock m with N_m aftershocks at times $t_1 < \dots < t_{N_m}$ (measured from the mainshock origin), the conditional log-likelihood over window $[0, T_m]$ is:

$$\ell_m(p, c) = \sum_{k=1}^{N_m} \log \frac{K_0}{(t_k + c)^p} - K_0 \int_0^{T_m} \frac{dt}{(t + c)^p} = N_m \log K_0 - p \sum_{k=1}^{N_m} \log(t_k + c) - \frac{K_0}{1-p} [(T_m + c)^{1-p} - c^{1-p}].$$

The joint log-likelihood $\ell(p, c) = \sum_m \ell_m(p, c)$ is maximised over the grid $p \in [0.50, 1.50]$ (50 points) and $c \in [10^{-2.5}, 1.0]$ day (28 log-spaced points).

Fitted values: $\hat{p} = 0.643$, $\hat{c} = 0.0032$ day.

Physical interpretation. The sub-unity exponent ($\hat{p} = 0.643 < 1$) implies slower aftershock decay than the $p \approx 0.9$ – 1.2 typical of crustal sequences. Low p -values are well documented for large subduction-zone megathrust sequences, where stress relaxation is slowed by the low geothermal gradient of the subducting slab; reported values reach $p \approx 0.6$ – 0.8 for the 2004 Sumatra and 2011 M_w 9.0 Tōhoku sequences (Utsu et al. 1995). The fitted value is thus consistent with a catalogue dominated by the prolonged aftershock train of the 2004 M_w 9.1 Aceh–Andaman rupture, whose triggered seismicity remained elevated above background for roughly a decade. Over the 475-year return period used for engineering design these post-seismic contributions average out, so the persistent low- p decay does not materially bias the long-term seismicity rate drawn from the declustered background catalogue.

S5 Spatial ETAS Grid-Search MLE

The four spatial ETAS parameters $(\mu_0, A, \alpha, \sigma_{\text{sp}})$ are estimated by Poisson MLE with Omori parameters fixed at $\hat{p} = 0.643$, $\hat{c} = 0.0032$ day.

The spatial conditional intensity at event i is

$$\lambda(\mathbf{s}_i) = \frac{\mu_0}{|\mathcal{W}|} + \sum_{j < i} \frac{A e^{\alpha(M_j - M_c)}}{2\pi\sigma_{\text{sp}}^2} e^{-|\mathbf{s}_i - \mathbf{s}_j|^2 / (2\sigma_{\text{sp}}^2)}.$$

The Poisson log-likelihood $\ell = \sum_i \log \lambda(\mathbf{s}_i) - \int \lambda d\mathbf{s}$ is maximised over a $30 \times 30 \times 30 \times 20$ parameter grid: $\mu_0 \in [10^{-5}, 10^{-1}]$, $A \in [10^{-4}, 10^1]$, $\alpha \in [0.5, 3.0]$, $\sigma_{\text{sp}} \in [2, 80]$ km (all log-spaced).

The spatial integral $\int \lambda d\mathbf{s}$ is evaluated on the 60×60 quadrature grid used for the LGCP.

S5.1 Anisotropic ETAS Kernel

Model (4) replaces the isotropic Gaussian triggering kernel with an elliptical kernel oriented along the local fault strike θ_j for event j :

$$g(\mathbf{s}; \mathbf{s}_j, \sigma_{\text{sp}}, \rho, \theta_j) = \frac{1}{2\pi\sigma_{\text{sp}}^2\rho} \exp\left(-\frac{d_{\parallel}^2}{2\sigma_{\text{sp}}^2\rho^2} - \frac{d_{\perp}^2}{2\sigma_{\text{sp}}^2}\right),$$

where $d_{\parallel}$ and $d_{\perp}$ are the components of $\mathbf{s} - \mathbf{s}_j$ parallel and perpendicular to the fault strike, and $\rho > 1$ is the aspect ratio (major/minor axis).

The fault strike θ_j at each event location is interpolated from the GEM/EOS Southeast Asia active-fault trace of the GSF. The aspect ratio ρ is added as a fifth free parameter and estimated jointly with $(\mu_0, A, \alpha, \sigma_{\text{sp}})$ by grid search over $\rho \in [1.0, 5.0]$ (10 values).

Result: $\hat{\rho} = 1.0$, confirming the declustered catalogue carries no fault-parallel directional signal ($\Delta\text{AIC} = +4.0$ relative to isotropic ETAS).

S6 PSHA Vectorised Algorithm

The Cornell (1968) hazard integral (Equation 8 of the main text) is discretised and vectorised as follows.

Grids:

- Source: 40×40 cells ($N_{\text{src}} = 1,600$), centred at $\{(x_k, y_k)\}$. This grid is independent of the 60×60 grid used to fit the LGCP trend coefficients; both sample the same continuous fitted intensity.
- Site: 28×28 nodes ($N_{\text{site}} = 784$), covering $\mathcal{W}$ at ~ 50 -km spacing.
- Magnitude: nine bins, $\Delta m = 0.5$, centred at $m \in \{4.75, 5.25, \dots, 8.75\}$.
- PGA: 60 levels log-spaced in $[0.001, 3.0]$ g.

Algorithm:

Step 1. For each source cell k , compute the annual mainshock rate: $\dot{n}_k = \hat{\lambda}_{\text{LGCP}}(\mathbf{x}_k) / T_{\text{obs}} \cdot \Delta A_k$.

Step 2. For each source cell k , site node ℓ , and magnitude bin m_b :

- (a) Compute source-to-site distance $R_{k\ell} = \|\mathbf{s}_{\ell} - \mathbf{x}_k\|$.

- (b) Evaluate median PGA and standard deviation ($\ln \overline{\text{PGA}}, \sigma_{\ln}$) from the Youngs et al. (1997) subduction-interface GMPE at $(m_b, R_{k\ell})$.
- (c) Compute $P(\text{PGA} > im \mid m_b, R_{k\ell}) = \Phi((\ln \overline{\text{PGA}} - \ln im)/\sigma_{\ln})$ for all 60 im levels.

Step 3. Sum contributions: $\lambda(\mathbf{s}_\ell, im) = \sum_k \sum_b \dot{n}_k f_M(m_b) \Delta m P(\text{PGA} > im \mid m_b, R_{k\ell})$.

Step 4. Convert to return period: $T_R = 1/\lambda$.

The truncated Gutenberg–Richter density is $f_M(m) \propto b \ln(10) 10^{-b(m-M_c)}$, truncated at $M_c = 4.5$ and $M_{\max} = 9.0$, with $\hat{b} = 1.05$.

S7 RELM Eight-Step Algorithm

Step 1. Define each of the 20 GSF segments as the straight-line trace connecting the NW and SE endpoints of Sieh and Natawidjaja (2000), with a ± 30 km perpendicular buffer.

Step 2. For each segment j , integrate the scaled LGCP intensity over the buffer area A_j : $\lambda_j^{\text{LGCP}} = T_{\text{obs}} \int_{A_j} \hat{\lambda}_{\text{LGCP}}(\mathbf{x}) d\mathbf{x}$.

Step 3. Compute the uniform-null expected count: $\lambda_j^{\text{unif}} = \bar{\lambda} A_j T_{\text{obs}}$, where $\bar{\lambda} = n_{\text{main}}/|\mathcal{W}|$.

Step 4. Count observed mainshocks in the buffer: ω_j .

Step 5. L-Test. Draw $N_{\text{sim}} = 2,000$ Monte Carlo realisations $\tilde{\omega}_j^{(s)} \sim \text{Poisson}(\lambda_j^{\text{LGCP}})$, compute the joint log-likelihood $\mathcal{L}^{(s)}$ for each, and compute the observed log-likelihood $\mathcal{L}_{\text{obs}}$. Quantile: $\gamma = N^{-1} \sum_s \mathbf{1}(\mathcal{L}^{(s)} \leq \mathcal{L}_{\text{obs}})$.

Step 6. N-Test. Simulate $N_{\text{tot}}^{(s)} \sim \text{Poisson}(\sum_j \lambda_j^{\text{LGCP}})$ and compute $\delta = N^{-1} \sum_s \mathbf{1}(N_{\text{tot}}^{(s)} \leq \omega_{\text{tot}})$.

Step 7. R-Test. Compute $\Delta\text{LL} = \mathcal{L}_{\text{LGCP}}(\boldsymbol{\omega}) - \mathcal{L}_{\text{unif}}(\boldsymbol{\omega})$ and the significance α_R by comparing to the distribution of ΔLL under the uniform model.

Step 8. Activity classification. Classify each segment by $O/E_u = \omega_j/\lambda_j^{\text{unif}}$: Highly Active (≥ 3.0), Active (1.0–3.0), Background (0.5–1.0), Quiescent (< 0.5).

S7.1 RELM Results

All values in this subsection are computed by `verify_holdout_relm.py` (Online Resource 2), which implements Steps 1–8 above; segment endpoints use the latitude bounds of Sieh and Natawidjaja (2000) with longitudes taken from the fault-trace polyline distributed with the code, so that the geometry is fully reproducible from the supplied files. Table S7.1 reports the per-segment expected counts, observed counts, and activity classification.

The unnormalised R-Test of Step 7 gives $\Delta\text{LL} = -3.3$ with $\alpha_R = 1.000$: observed counts of this concentration along the fault would essentially never arise under the uniform data-generating model, so the uniform null is rejected. The negative ΔLL itself,

however, shows that the absolute LGCP forecast is not a better description of the segment counts than the uniform one, for two reasons that the companion tests make explicit. First, the N-Test gives $\delta = 0.00$ (forecast 258 mainshocks within the buffers against 141 observed): the fitted surface, which integrates to the 1998–2024 catalogue total over the whole domain, concentrates rate near the fault and therefore over-predicts within the ± 30 -km corridor. Second, a count-normalised (purely spatial) comparison, which removes the rate component by conditioning both forecasts on the observed buffer total, gives $\Delta LL_{\text{spatial}} = -6.6$: the along-strike allocation of the smooth covariate-driven surface is no better than uniform across the 20 segments. This is the segment-scale counterpart of the main-text conclusion (Sect. 6) that the LGCP carries genuine spatial information as a continuous two-dimensional field but that this structure, smooth at the 50–200 km correlation scale, is not resolved by a coarse 20-segment partition once triggered clustering has been removed.

The activity classification of Step 8 is end-loaded along strike. Five segments are Highly Active ($O/E_u \geq 3.0$): Kumering (6.9), Manna (5.2), and Semangko (3.8) in the south-east, and Aceh (4.9) and Seulimeum (3.3) in the north-west, the latter pair consistent with Coulomb stress transfer following the 2004–2005 megathrust ruptures. At the other extreme, Toru (0.31) is Quiescent and Sumpur (0.65), Angkola (0.77), and Dikit (0.80) record counts below the uniform-null expectation; whether this reflects aseismic creep, deep locking, or the brevity of the 26-year observation window cannot be resolved from seismicity data alone.

Table 1: *

Table S7.1 Segment-level RELM summary for the 20 Great Sumatran Fault segments (± 30 -km buffers; G–K mainshocks $M \geq 4.5$, 1998–2024). λ_j^{LGCP} , λ_j^{unif} : expected counts under the LGCP forecast and the uniform null; ω_j : observed count; $O/E_u = \omega_j/\lambda_j^{\text{unif}}$.

HA = Highly Active, A = Active, B = Background, Q = Quiescent.

Segment	λ_j^{LGCP}	λ_j^{unif}	ω_j	O/E_u	Class
Seulimeum	8.0	3.7	16	4.37	HA
Aceh	11.0	4.8	20	4.16	HA
Tripa	12.9	4.8	6	1.25	A
Renun	16.6	5.4	12	2.23	A
Toru	10.6	3.1	9	2.92	A
Angkola	18.5	5.2	13	2.50	A
Sumpur	6.0	1.5	3	1.95	A
Sianok	12.8	3.3	7	2.14	A
Sumani	8.4	2.1	6	2.83	A
Suliti	10.7	2.7	4	1.48	A
Siulak	10.5	2.7	2	0.74	B
Dikit	9.6	2.5	3	1.20	A
Ketaun	10.7	2.9	2	0.69	B
Musi	7.0	1.9	2	1.04	A
Manna	10.7	3.1	5	1.62	A
Kumering	13.1	3.9	5	1.30	A
Semangko	8.4	2.5	6	2.40	A
Sunda	2.6	0.8	1	1.30	A
Barumun	11.3	3.3	0	0.00	Q
Batee	10.7	4.6	6	1.30	A

S8 Fault-Distance Covariate: Megathrust Proxy and Covariate Confounding

The large negative fault-distance coefficient $\hat{\beta}_{\text{fault}} = -1.396$ (Table 3 of the main text) acts partly as a proxy for distance to the *megathrust interface*, not only for distance to the Great Sumatran Fault (GSF). Across the Sumatra window the GSF runs approximately parallel to, and ~ 150 km landward of, the trench, so the perpendicular fault-distance field D_{fault} is empirically anti-correlated with the offshore rupture-band geometry of the great subduction events. A single fault-distance covariate therefore captures a large share of the spatial variance in log-rate without an explicit distance-to-trench predictor: the two distances are confounded by Sumatran geography.

Including both D_{fault} and a separate distance-to-trench covariate is not pursued, because the two predictors are strongly collinear over the study domain: adding an explicit trench-distance term would inflate the variance inflation factor without contributing independent spatial information, while leaving the fitted intensity surface essentially unchanged. The single fault-distance covariate is therefore retained as the more parsimonious and better-identified specification, with the understanding that it encodes proximity to both the GSF and the sub-parallel megathrust interface.

S9 Engineering Hazard Products and Code Reconciliation

This section collects the engineering products of the fitted hazard surface referenced in Section 6 of the main text: the magnitude disaggregation, the cross-city design-level comparison, the rare-event amplification ratio, the full eight-return-period table for Padang, and the domain-level comparison against the national code surface. All quantities are computed with the scripts `generate_fig4_hazard.py` and `generate_fig5_sni.py` provided with this supplement, using the fitted LGCP coefficients in `lgcp_parameters.csv`.

S9.1 Disaggregation, design levels, and amplification

Figure S9.1 shows three panels. Panel (A) gives the magnitude disaggregation of the Padang hazard at the 0.1 g threshold: the contribution decreases monotonically with magnitude, with the M_w 4.75 bin dominating, a direct consequence of the completeness-matched magnitude floor $M_c = 4.5$ discussed in Section 6.3 of the main text. Panel (B) compares the 475-yr and 2475-yr design accelerations across the four reference cities. Panel (C) maps the rare-event amplification ratio $\text{PGA}_{2475}/\text{PGA}_{475}$, which exceeds 2 along the high-hazard corridor because the steeper magnitude–distance tail at short source-to-site distances amplifies rare-event ground motion.

S9.2 Eight-return-period hazard table for Padang

Table S9.1 extends Table 5 of the main text to eight return periods with engineering design-level annotations. Values are computed at the same site-grid node as Table 5 (rock site, $V_{S30} = 760$ m/s, $H = 25$ km, $\hat{b} = 1.05$); epistemic bounds are $\pm 30\%$ as in the main text.




Figure 1: *

Fig. S9.1 (A) Magnitude disaggregation of the Padang hazard ($PGA > 0.1$ g). (B) 475-yr and 2475-yr design accelerations at the four reference cities. (C) Rare-event amplification ratio PGA_{2475}/PGA_{475} across the domain.

Table 2: *

Table S9.1 Site-specific hazard for Padang at eight return periods.

T_R (yr)	λ (yr^{-1})	PoE ₅₀ (%)	PGA (g)	Low _{ep}	High _{ep}	Design level
43	0.02326	68.7	0.530	0.371	0.690	Frequent (serviceability)
72	0.01389	50.1	0.705	0.494	0.917	Occasional
224	0.00446	20.0	1.256	0.879	1.633	20% in 50 yr
475	0.00211	10.0	1.791	1.254	2.328	Design basis (10% in 50 yr)
975	0.00103	5.0	2.474	1.732	3.216	5% in 50 yr
1,500	0.00067	3.3	2.983	2.088	3.877	Long return period
2,475	0.00040	2.0	3.703	2.592	4.813	MCE (2% in 50 yr)
4,975	0.00020	1.0	5.005	3.504	6.507	Critical facilities

S9.3 Domain-level comparison with SNI 1726:2019

Figure S9.2 compares the LGCP 475-yr surface (Panel A) with a surface reconstructed from published SNI 1726:2019 city anchor values by radial basis function interpolation (Panel B); Panel (C) maps their ratio. The city-level ratios computed from the site hazard curves are 2.7 at Padang, 2.5 at Pekanbaru, 2.1 at Medan, and 1.5 at Banda Aceh, as reported in Section 6.3 of the main text. On the coarser 60×60 ratio map the Banda Aceh cell reads lower (≈ 1.7) because the city sits near the domain edge where the map grid undersamples the hazard gradient; the hazard-curve values are the reference.




Figure 2: *

Fig. S9.2 (A) LGCP 475-yr PGA. (B) SNI 1726:2019 475-yr surface reconstructed from published city anchors. (C) Ratio LGCP/SNI; the dashed contour marks parity (ratio = 1).