SCEC Project 25341 — Final Report

Combining hierarchical MCMC and dynamic rupture simulations in a supercomputing framework to infer off-fault damage parameters

Zihua Niu (Scripps Institution of Oceanography, UC San Diego), Alice-Agnes Gabriel (Scripps Institution of Oceanography, UC San Diego), Yehuda Ben-Zion (University of Southern California)

Project Period: February 1, 2025 – January 31, 2026

Funding	Organization Budget	Investigators
Scripps Institution of Oceanography, UC San Diego	$25,572	Zihua Niu, Alice-Agnes Gabriel
University of Southern California	$0	Yehuda Ben-Zion

Abstract

Theoretical earthquake source models predict large stress concentrations at the rupture front that must be accommodated by anelastic, off-fault deformation. While decades of laboratory experiments have constrained on-fault frictional parameters, parameters governing off-fault brittle damage remain poorly known because they are difficult to measure directly in the field at seismogenic depths. This project performed, for the first time, a Bayesian inversion for both on-fault and off-fault rheology parameters in fully 3D dynamic rupture simulations of the 2019 M_w 7.1 Ridgecrest mainshock. We coupled SeisSol dynamic rupture simulations with a prefetching multilevel delayed acceptance (MLDA) MCMC sampler through the UM-BRIDGE software framework, deployed across a hierarchy of forward models with a Gaussian-process surrogate at the cheapest level and high-fidelity 3D simulations at the finest level. Using ten near-fault GNSS time series and along-strike fault-parallel offsets from satellite imagery, we constrained four nonlinear parameters: linearly along-strike variable scaling factors of off-fault Drucker–Prager plastic cohesion (γ_0, γ_1) and of the on-fault direct-effect parameter in strong-velocity-weakening rate-and-state friction (α_0, α_1). Eight Markov chains, run in parallel after pre-training a Gaussian-process surrogate, converged to posteriors with rank-normalized $\hat{R} < 1.05$ after 80 samples per chain. The preferred posterior models reproduce the observed seismic moment magnitude (M_w = 6.98 ± 0.01 vs. 7.01 in the reference model), match independently determined kinematic moment release rates, and provide quantified posterior uncertainties on the spatially variable off-fault damage and on-fault frictional weakening across the conjugate Ridgecrest fault network. The project delivered an open, supercomputer-ready Bayesian inversion workflow for 3D dynamic rupture, the first application of MLDA to a real complex earthquake, and a refereed publication in Earth and Planetary Science Letters.

Intellectual Merit

This award produced the first Bayesian inversion for off-fault damage rheology parameters in 3D dynamic rupture models of a real earthquake. Three advances stand out. First, methodologically, we demonstrated that prefetching MLDA, in which most likelihood evaluations are spent on a Gaussian-process surrogate and a polynomial-order-2 SeisSol forward model and only the most promising proposals are passed to the polynomial-order-3 fine model, makes Bayesian inference tractable for forward problems whose single-evaluation cost would otherwise preclude MCMC. The level-3 model alone is roughly 48 times cheaper than the reference high-resolution model used by Taufiqurrahman et al. (2023), yet retains accuracy sufficient for the static and low-frequency observables that drive the inversion. Second, the inversion simultaneously constrained on-fault rate-and-state friction and off-fault Drucker–Prager cohesion, exposing parameter trade-offs that single-parameter sensitivity studies cannot resolve and that would be invisible to deterministic best-fit searches. Third, by parameterizing both on- and off-fault rheology with along-strike linear variations, we recovered different effective strengths on the northwestern and southeastern segments of the Ridgecrest mainshock fault, consistent with independently observed surface-rupture complexity. The inversion workflow, built on the open-source SeisSol and UM-BRIDGE codes, is portable across CPU (Frontera, TACC) and GPU (LUMI) supercomputers.

Broader Impacts

This project supported PhD student Zihua Niu, visiting at SIO/UCSD, who led the inversion design, forward modeling, and analysis and presented the work at the 2025 SCEC Annual Meeting. The award also enabled close collaboration between SIO/UCSD, USC, the Karlsruhe Institute of Technology, and Ludwig-Maximilians-Universität München, fostering exchange between the dynamic-rupture and Bayesian-inference communities. The HPC-optimized workflow that resulted from the project lowers the barrier for the broader earthquake science community to perform uncertainty-quantified dynamic rupture inversions: SeisSol, UM-BRIDGE, and the prefetching MLDA implementation of Kruse et al. are open source, documented, and reusable. Quantified posterior uncertainties on off-fault damage and on-fault friction directly inform physics-based seismic hazard assessment and ground-motion prediction by exposing how confidently each rheology parameter can be inferred from the available seismic and geodetic data.

Detailed Project Report

1. Goals and approach

The proposed work targeted SCEC Milestones C1,2,3-1, C1-2, C2-2, and C3-1 by combining HPC-optimized dynamic rupture forward simulations with hierarchical Bayesian inference. The Ridgecrest 2019 sequence was selected because of the unusually dense near-field instrumental coverage — ten high-rate GNSS stations recording mainshock displacements, regional broadband seismometers, and fault-parallel offsets measured by sub-metric optical image correlation (Antoine et al., 2021) — and because of the existence of a published reference 3D dynamic rupture model (Taufiqurrahman et al., 2023) that already explained the first-order kinematics. Our objective was not to refit the reference model, but to relax the spatially uniform assumptions on its on- and off-fault rheology parameters and to recover, with quantified posterior uncertainty, the along-strike variations consistent with the data.

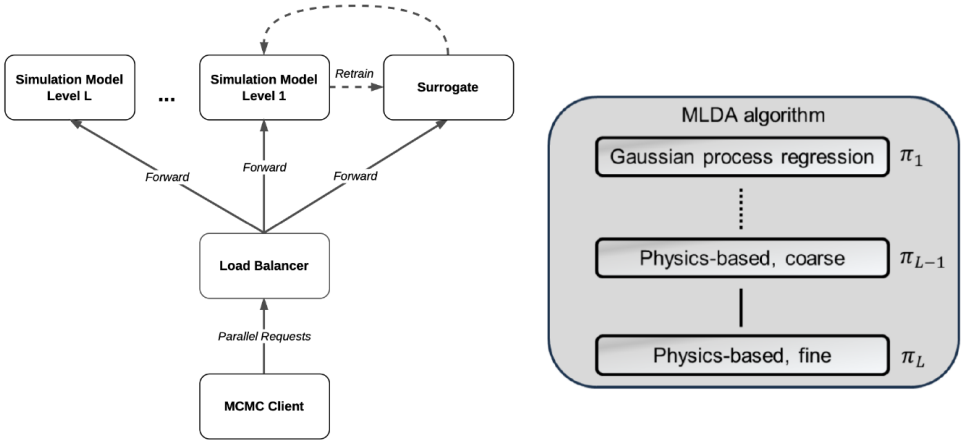

Figure 1. Computational setup linking HPC dynamic rupture simulations to UM-BRIDGE for hierarchical Bayesian inversion. Left: the MCMC client issues parallel proposal requests to a load balancer that distributes evaluations across simulation models of increasing fidelity (Level 1 to Level L) and to a Gaussian-process surrogate that is retrained adaptively. Right: in the prefetching MLDA hierarchy, level π_1 is the Gaussian-process regression surrogate, intermediate levels are coarse SeisSol simulations, and level π_L is the high-order fine SeisSol simulation. Adapted from the project proposal.

2. Forward model and inversion framework

Forward simulations were performed with SeisSol on the geometrically complex four-segment Ridgecrest fault system of Taufiqurrahman et al. (2023), with strong-velocity-weakening rate-and-state friction on the faults and non-associated Drucker–Prager elasto-viscoplasticity in the off-fault medium (Wollherr et al., 2018). The computational domain spans 200 × 200 × 100 km, discretized with approximately 4.0 million tetrahedral elements and a 200 m maximum element size near the faults. Two simulation levels were used during inversion — polynomial order 2 (level 2) and polynomial order 3 (level 3) — resolving ground motions up to 0.5–1.0 Hz, which is sufficient for the static and low-frequency-displacement observables that drive the likelihood.

The likelihood function combines (i) static three-component displacements at twenty-one GNSS stations, (ii) along-strike fault-parallel surface offsets from optical image correlation (Antoine et al., 2021) on either side of the epicenter, and (iii) high-rate displacement time series at ten near-fault GNSS stations (Melgar et al., 2020). We chose the standard deviations in the covariance matrix to reflect epistemic model uncertainty rather than measurement precision, calibrated against the misfits achievable in a Gaussian-process pre-training pass; this prevents the surrogate from being over-confident in regions of parameter space where simple linear parameterizations cannot reproduce the data. Eight independent Markov chains, each of 80 samples after a 20-sample burn-in, were run in parallel using the prefetching MLDA implementation of Kruse et al. (2025). The four model parameters are the off-fault plastic-cohesion linear-scaling factors (γ_0, γ_1) and the on-fault SVW-RS direct-effect linear-scaling factors (α_0, α_1) on the northwestern and southeastern fault segments of F1, respectively.

3. Posterior inversion results

The chains converged with rank-normalized $\hat{R}$ dropping below 1.05 after the burn-in window and autocorrelation functions falling below 0.1 within roughly ten samples, indicating well-mixed posteriors over all four parameters. The 1D marginal posteriors peak at $\gamma_0 \approx 0.6$ and $\gamma_1 \approx 0.7$ for off-fault plastic cohesion, and at $\alpha_1 \approx 1.2$ for the southeastern direct-effect scaling factor, with somewhat broader posterior support on α_0. The 2D marginals reveal a positive trade-off between γ_0 and α_0 on the northwestern segment, consistent with the expected coupling between on-fault dynamic strength drop and the strength of the off-fault medium that absorbs the radiated stress concentrations. Posterior models reproduce the observed fault-parallel surface offsets of Antoine et al. (2021) (Fig. 2k), match the kinematically inverted seismic moment release rate of Goldberg et al. (2020), and yield a moment magnitude M_w = 6.98 ± 0.01, in close agreement with the reference value of 7.01.

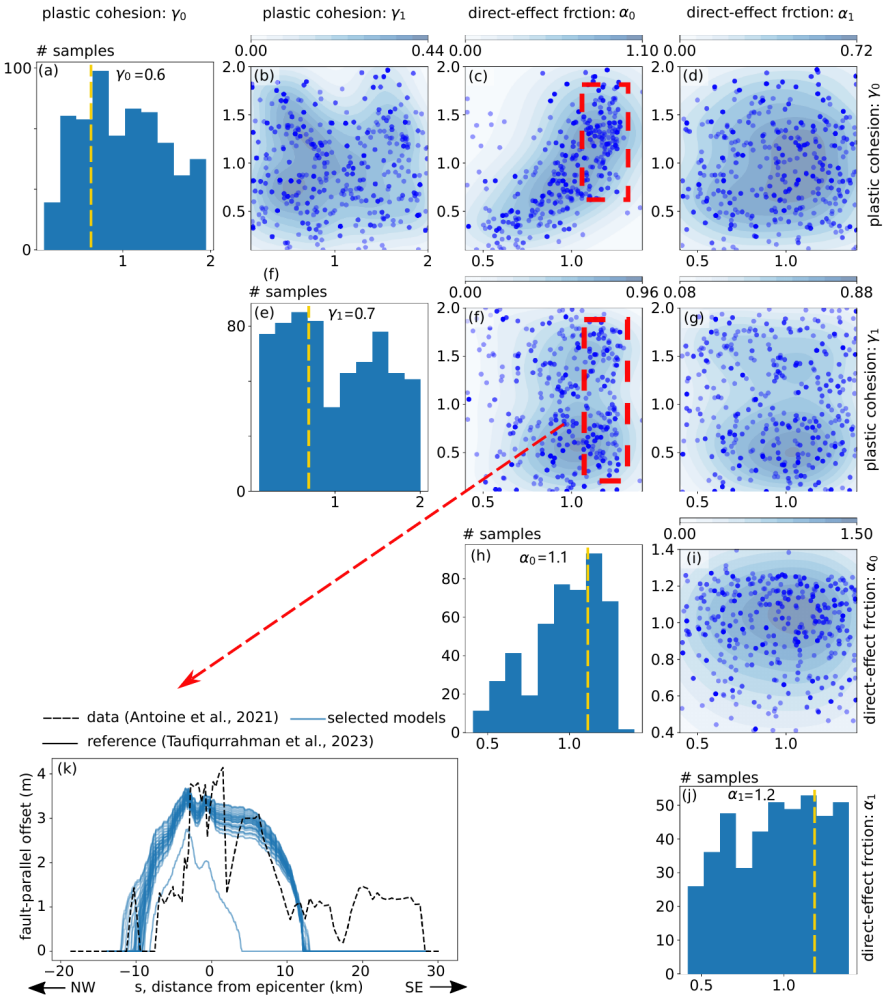

Figure 2. Bayesian posterior $\rho(m|d_{obs})$ from MLDA inversion. Diagonal panels (top to bottom): 1D marginal distributions of off-fault plastic cohesion scaling factors γ_0, γ_1 and on-fault direct-effect scaling factors α_0, α_1. Super-diagonal panels: 2D marginal posteriors for each parameter pair, with effective samples from the eight Markov chains shown as blue dots and the posterior density estimated by Gaussian kernel density. Panel (k): fault-parallel surface offsets predicted by the posterior models (blue) compared to the optical-image correlation data of Antoine et al. (2021, dashed black) and the reference model of Taufiqurrahman et al. (2023, solid black). Adapted from the EPSL paper (Niu et al., in press).

4. Off-fault damage and rupture-dynamics implications

Off-fault inelastic deformation in the preferred posterior models is concentrated in a flower-like damage zone whose width broadens with decreasing depth, with peak plastic strain η of order 10^{-2} within ~1–2 km of the fault and a long tail extending to a few kilometers off-fault near the surface. The map-view distribution of plastic strain in the preferred model (Fig. 3b) is similar in extent and amplitude to that of the reference model (Fig. 3a), but with reduced strain on the southeastern segment of F1, consistent with the lower posterior cohesion scaling there. This spatial pattern is consistent with observed shallow slip deficit and with the diffuse near-fault deformation reported by Antoine et al. (2021), and provides a quantitative, uncertainty-bounded baseline against which observed near-fault aftershock distributions and dense-array near-fault sensor data can now be compared. Importantly, the four-parameter linear parameterization is sufficient to fit static displacements and fault-parallel offsets, but the residuals at high-rate GNSS stations CCCC, P580, and P595 — where modeled peaks arrive earlier than observed — indicate that the rupture speed on the southeastern segment is overestimated. This points to additional structure (depth-dependent friction, depth-dependent damage, or 3D variations in stress orientation) that future inversions, with richer parameterizations enabled by the same MLDA workflow, can target.

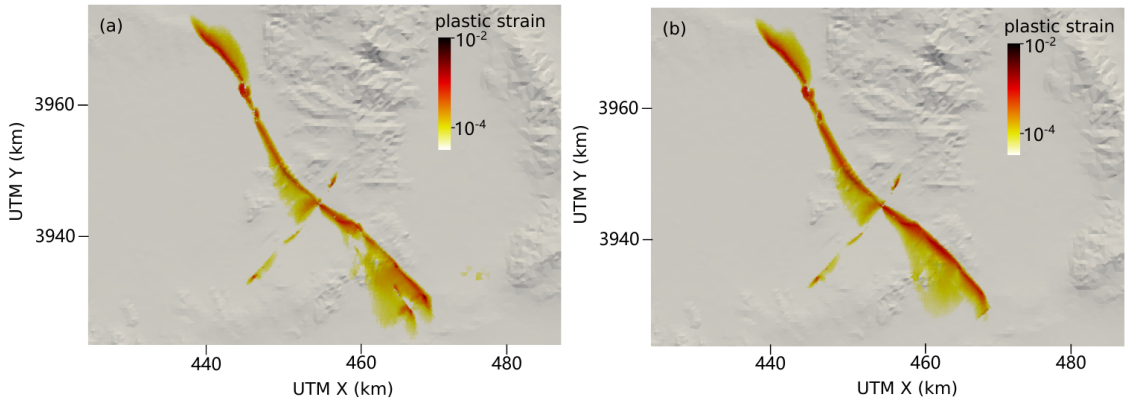

Figure 3. Map view of off-fault plastic strain η around the Ridgecrest fault system. (a) Reference model of Taufiqurrahman et al. (2023). (b) Preferred posterior model from this project. The preferred model recovers the broad spatial pattern and amplitude of off-fault damage seen in the reference model while reducing strain on the southeastern fault segment, consistent with the inferred along-strike variation of plastic cohesion and on-fault friction. Adapted from the EPSL paper (Niu et al., in press).

5. Validation, verification, and outlook

Verification of the MLDA sampler used in this study draws on the published benchmarks of Kruse et al. (2025), which document order-of-magnitude reductions in autocorrelation and corresponding gains in effective sample size on a controlled “banana” posterior, and which confirm stable two-parameter Bayesian posteriors on a simplified Ridgecrest test problem with the same SeisSol forward-model family. For the production four-parameter inversion, we relied on rank-normalized $\hat{R}$, autocorrelation functions, and posterior concentration as primary diagnostics, supplemented by independent validation against (i) the kinematically inverted seismic moment release rate of Goldberg et al. (2020), (ii) ground-velocity waveforms at five regional seismic stations bandpass filtered between

0.1 and 0.3 Hz, and (iii) the moment magnitude derived from the integrated moment-rate function. The preferred posterior models satisfy all three independent checks within their assigned uncertainties. The principal output of this project is the Earth and Planetary Science Letters paper of Niu et al. (in press), which fully documents methodology, results, and validation. The MLDA–SeisSol–UM-BRIDGE workflow developed under this award has been deployed on Frontera (TACC) and LUMI (Finland), is publicly available, and is being extended to ensemble dynamic rupture inversions for the Tohoku-Oki and Cape Mendocino earthquakes (Wong, Gabriel, & Fan, 2026; Ulrich, Magen, & Gabriel, submitted) and to Bayesian uncertainty quantification of fault-geometry effects via mesh morphing (Hobson, May, & Gabriel, 2025). The four-parameter linear parameterization here is a deliberately conservative starting point: the same workflow now supports higher-dimensional inversions of depth-dependent friction, characteristic slip distance L, and bulk Drucker–Prager friction angle.

Publications and preprints acknowledging this award

1. Garagash, D. I., & Gabriel, A.-A. (2026). Localization of fast and slow slip in fault gouge and fracture energy scaling. Journal of Geophysical Research: Solid Earth, submitted; arXiv preprint.
2. Niu, Z., Kruse, M., Seelinger, L., Schliwa, N., Igel, H., & Gabriel, A.-A. (in press). Constraining on- and off-fault nonlinear dynamic rupture parameters via hierarchical Bayesian inversion for the 2019 Mw 7.1 Ridgecrest earthquake. Earth and Planetary Science Letters.
3. Wong, J. W. C., Gabriel, A.-A., & Fan, W. (2026). Dynamic restrengthening and fault heterogeneity explain megathrust earthquake complexity. Nature Communications.
4. Gabriel, A.-A., Karki, P., Magen, Y., Oryan, B., Ulrich, T., Yun, J., & May, D. A. (in press). Tandem: An open-source high-performance computing volumetric software to model sequences of earthquakes and aseismic slip across complex fault systems. Seismological Research Letters.
5. Kurapati, V., Schneller, D., Seelinger, L., Niu, Z., Gabriel, A.-A., & Bader, M. (2026). Fused ensembles of dynamic-rupture earthquake simulations to accelerate Bayesian inference. International Journal on Geomathematics, 17, 8.
6. Kurapati, V., Hillers, G., Krenz, L., Gabriel, A.-A., & Bader, M. (2026). Numerical wavefield simulations with instantaneous time mirror in a 3D elastic medium. Geophysical Journal International, 245(1), ggag031.
7. Hobson, G. M., May, D. A., & Gabriel, A.-A. (2025). Quantifying the influence of fault geometry via mesh morphing with applications to earthquake dynamic rupture and thermal models of subduction. Geochemistry, Geophysics, Geosystems, 26, e2025GC012531.
8. Biemiller, J., Gabriel, A.-A., Staisch, L. M., Ulrich, T., Dunham, A., Wirth, E. A., Watt, J. T., Lucas, M., & Ledeczi, A. (2025). Structural controls on splay fault rupture dynamics during Cascadia megathrust earthquakes. AGU Advances, 6, e2025AV001812.
9. Niu, Z., Gabriel, A.-A., & Ben-Zion, Y. (2025). Delayed dynamic triggering and enhanced high-frequency seismic radiation due to brittle rock damage in 3D multi-fault rupture simulations. Journal of Geophysical Research: Solid Earth, 130, e2025JB031632.
10. Wu, B., Oglesby, D. D., Kyriakopoulos, C., Tarnowski, J. M., & Gabriel, A.-A. (submitted). General solution scaling relations in linear elastic dynamic rupture models. Journal of Geophysical Research: Solid Earth, submitted.
11. Ulrich, T., Magen, Y., & Gabriel, A.-A. (submitted). The complex rupture dynamics of an oceanic transform fault: supershear rupture and deep slip during the 2024 Mw 7.0 Cape Mendocino earthquake. EarthArXiv preprint.

SCEC Annual Meeting presentations acknowledging this award

1. Gabriel, A., Niu, Z., Kruse, M., Seelinger, L., Schliwa, N., Igel, H., & Ben-Zion, Y. (2025). Constraining on- and off-fault nonlinear dynamic rupture parameters via hierarchical Bayesian inversion of GNSS and satellite data for the 2019 Mw 7.1 Ridgecrest earthquake. Poster #165, 2025 SCEC Annual Meeting (SCEC Contribution #14774).
2. Hobson, G., May, D. A., & Gabriel, A. (2025). Using mesh morphing and reduced-order modeling to quantify the influence of fault geometry on earthquake dynamic rupture. Poster #163, 2025 SCEC Annual Meeting (SCEC Contribution #14759).