Reduced-order modelling of Cascadia's slow slip cycles

Yohai Magen, Dave A. May, Alice-Agnes Gabriel

This manuscript has been submitted for publication in *Geophysical Journal International*. This version is a non-peer-reviewed preprint submitted to *EarthArXiv*.

Please note that the manuscript has yet to be formally accepted for publication. Subsequent versions of this manuscript may have slightly different content.

November 20, 2025

This is a non-peer reviewed preprint submitted to EarthArXiv.

Reduced-order modelling of Cascadia's slow slip cycles

Yohai Magen [1], Dave A. May [1], and Alice-Agnes Gabriel [1,2]

[1] *Institute of Geophysics and Planetary Physics, Scripps Institution of Oceanography, University of California, San Diego, CA, USA*

[2] *Department of Earth and Environmental Sciences, Ludwig-Maximilians-Universität München, Munich, Germany*

November 20, 2025

Abstract

Slow-slip events (SSEs) modulate the earthquake cycle in subduction zones, yet understanding their physics remains challenging due to sparse observations and high computational cost of physics-based simulations. We present a scientific machine-learning approach using a data-driven reduced-order modelling (ROM) framework to efficiently simulate the SSE cycle governed by rate-and-state friction in a Cascadia-like 2D subduction setting. Our approach projects fault slip, slip-rate, and state-variable trajectories onto a spline-based latent space, which is subsequently emulated using proper-orthogonal decomposition and radial-basis-function interpolation. Achieving a speedup of $\sim 360,000\times$ compared to volumetric simulations, the ROMs enable comprehensive parameter exploration and Bayesian Markov chain Monte Carlo (MCMC) inversion. By smoothly interpolating between the physics-based simulations, the ROMs reveal complex dependencies that might be overlooked with coarser parameter space sampling. Our analysis reveals complex, non-linear dependencies of SSE characteristics on the width and magnitude of the deep, low-effective-normal-stress region while holding frictional parameters constant. We show that, to first order, the recurrence time of SSEs is linearly dependent on the normalized fault width, defined as the SSE zone width divided by the characteristic nucleation length, in agreement with previous work. However, at second order, the recurrence interval can change more rapidly with small variations in the SSE zone width. We identify a region of steep, non-linear dependence of the recurrence interval on the normalized fault width, which we attribute to the extent of the velocity-weakening portion of the subduction interface that produces SSEs. Our MCMC inversion constrained by Northern Cascadia SSEs observations indicates near-lithostatic pore fluid pressure ($99.6 \pm 0.17\%$ lithostatic) and positions the upper frictional transition zone at 30.4 ± 2.8 km depth, consistent with geophysical observations. The inversion resolves the deep SSE-portion of the slab spanning 45 ± 16 km with low effective normal stress of 3.8 ± 1.4 MPa. We discuss how varying the fault frictional parameters affects the MCMC-inverted parameter values. This framework provides a new tool for advancing the physics-based understanding of SSEs and subduction zone faulting mechanics. By systematically linking megathrust properties such as fluid pressure and fault strength to rate-and-state friction governed slow slip cycle characteristics, such as recurrence interval, our approach helps to constrain the first- and second-order physics-based controls and the uncertainties of how subduction zones slip.

1 Introduction

The conventional view of the earthquake cycle as a simple stick-slip phenomenon, where accumulated tectonic strain is released only through seismogenic earthquakes (Ried, 1911), has evolved dramatically over the past two decades. Advanced seismological and geodetic observation networks have revealed a rich spectrum of fault slip behavior (Gomberg et al., 2016; Ide and Beroza, 2023). Among these, slow slip events

(SSEs) have emerged as an important component of the seismic cycle, accommodating a large portion of tectonic plate motion and thus contributing to the moment budget at large continental faults (Linde et al., 1996; Rousset et al., 2019) and in many subduction zones worldwide (Behr and Bürgmann, 2021; Schwartz and Rokosky, 2007). For example, SSEs have been documented at the Cascadia Subduction Zone (CSZ) (Schmidt and Gao, 2010), offshore from Boso Peninsula (Ozawa et al., 2007), in the Nankai Trough, Japan (Araki et al., 2017; Obara et al., 2004), and along the Hikurangi margin in New Zealand (Wallace et al., 2012). Slow slip events typically manifest as transient reversals in the direction of surface deformation, detected by high-precision Global Navigation Satellite System (GNSS) networks. SSE displacement amplitudes reach several millimeters to centimeters which is 10–100 times faster than the tectonic loading and durations range from days to years (Bartlow et al., 2011; Dragert et al., 2001; Peng and Gomberg, 2010). The CSZ has been particularly instrumental for studying SSEs (Bartlow et al., 2011; Schmidt and Gao, 2010). There, transients recur quasi-periodically every $\approx$14 months and have been monitored for the past $\sim$25 years (Gomberg et al., 2016; Schmidt and Gao, 2010).

The physical mechanisms controlling SSEs remain enigmatic, as their potential role as precursors of imminent megathrust earthquakes (Li and Gabriel, 2024; Obara and Kato, 2016; Ruiz et al., 2014), despite extensive research (Bürgmann, 2018). Their occurrence at depths ranging from $\sim$25-40 km introduces uncertainties about the structure, material properties, and pressure and temperature conditions at these transitional depths (Behr and Bürgmann, 2021). This, combined with the small surface deformation signals produced by SSEs and the rapid loss of resolution with depth, limits the constraints that observational data alone can place on their governing physics (Liu and Rice, 2007). Consequently, numerical simulations incorporating laboratory-derived friction laws have become essential for investigating the mechanics of these phenomena (Liu and Rice, 2007; Segall et al., 2010). Rate-and-state friction has emerged as the predominant framework for simulating sequences of earthquakes and aseismic slip (SEAS, e.g., Barbot et al. (2012); Erickson et al. (2023); Jiang et al. (2022); Kato (2002); Rice and Tse (1986)). These constitutive laws describe fault strength as dependent on both slip rate and a state-variable, representing contact properties at the fault interface (Dieterich, 1979; Ruina, 1983).

Traditionally, two classes of rate-and-state friction models have been used to reproduce the key characteristics of SSEs. In both frameworks, SSEs arise spontaneously from specific combinations of frictional stability regimes along the fault, particularly, including a conditionally stable region near the transition between fully locked and continuously creeping sections. In the first class of models, SSEs emerge from a narrow velocity-weakening (VW) region embedded within a velocity-strengthening (VS) zone, where the ratio between VW patch width and nucleation size governs SSEs recurrence and propagation (Rubin, 2008). The second class, introduced by Liu and Rice (2007), focuses on transitional frictional stability near the down-dip limit of the seismogenic zone. In this approach SSEs nucleate within conditional stable fault portions characterized by VW behavior that is stabilized by low effective normal stress and proximity to the VW-VS transition. To model spontaneous SSEs, these models incorporate a fault portion with reduced effective normal stress (e.g. Li and Liu, 2016b; Liu and Rice, 2005, 2009; Luo and Ampuero, 2018; Matsuzawa et al., 2010; Rubin, 2008). In subduction zones, this may result from elevated pore fluid pressure at the megathrust interface (Audet et al., 2009; Suppe, 2014). The Liu and Rice (2007) framework demonstrates that the ratio between the width of this low effective stress zone and the nucleation length scale (h^*) critically controls both the occurrence and characteristics of SSEs, with higher ratios producing more frequent events of smaller magnitude (Cattania, 2019; Liu and Rice, 2009).

More recent efforts have focused on incorporating rate-and-state friction fault, along with fluid flow, permeability evolution, and changes in pore fluid pressure to produce SSEs (Ozawa et al., 2024; Perez-Silva et al., 2023). This aims to provide a more realistic physical explanation that aligns with observations. Despite their success in reproducing observed characteristics of SSEs, physics-based forward models are often simplified, e.g. by reducing dimensionality or idealizing fault geometry, and typically focus on limited aspects of the parameter space due to computational challenges. Simulating SSEs requires capturing processes

across vastly different timescales, from the evolution of individual slow slip events to the periods where fault slip rate and slip rate gradients are low. This multi-scale character necessitates repeated fine temporal resolution to capture SSE nucleation dynamics, while simultaneously requiring long simulation times to capture multiple SSE cycles. The computational costs become problematic when attempting to conduct sensitivity analyses or quantify uncertainties in model predictions, motivating the need for more efficient simulation approaches.

Scientific machine learning (ML) methods have recently advanced the analysis of SSE observations. Deep learning models, in particular, have proven effective in detecting low signal-to-noise ratio signatures of SSEs from geodetic data and their seismic counterparts, low-frequency earthquakes (LFEs) and non-volcanic tremor (Hulbert et al., 2022). For instance, Lin et al. (2024) applied Convolutional Neural Networks (CNNs) to continuous seismic records to build comprehensive LFE catalogs, revealing thousands of previously uncataloged events, suggesting hidden SSEs which were unraveled previously. Münchmeyer et al. (2024) demonstrated that such models can generalize across different subduction zones, suggesting that LFEs may share universal waveforms characteristics that are learnable by deep networks. From a geodetic perspective, Costantino et al. (2023) developed a multi-station deep learning detector that operates on raw GNSS time series to identify subtle deformation transients associated with SSEs. To overcome the scarcity of labeled training data, they used a synthetic training simulation data set. Early efforts to apply ML directly to the governing physics of fault slip include Physics-Informed Neural Networks (PINNs), embedding the governing equations into the loss function of neural networks to estimate fault properties (e.g. Fukushima et al., 2023, 2025; Okazaki et al., 2022; Rucker and Erickson, 2024). Recently, Fourier Neural Operators (FNOs) have been used to emulate the nonlinear equations governing dynamic rupture propagation (Tainpakdipat et al., 2025).

Reduced-order models (ROMs) have proven to be an effective method to accelerate geoscience simulations, including modelling slow slip events (Kaveh et al., 2024), seismic shake maps (Rekoske et al., 2023) and wave fields (Rekoske et al., 2025), thermal structures in subduction zones (Hobson and May, 2025a), geothermal geodynamic processes (Degen et al., 2023) and magnetotellurics (Quiaro et al., 2025). By projecting high-dimensional systems onto lower-dimensional subspaces that retain the primary characteristics, ROMs make it possible to capture the overall evolution of complex physical systems at a fraction of the cost of full-order models (FOMs). Reported speedups range from several hundred to over 100,000 times faster, depending on the application. Such reductions in computational cost enable global exploration of the parameter space and facilitate uncertainty quantification that are impractical with physics-based models.

In this study, we introduce and evaluate a new reduced-order modelling framework designed to efficiently simulate SSE cycles governed by rate-and-state friction. We utilize the volumetric, scalable discontinuous Galerkin library (Uphoff et al., 2022) on unstructured meshes to simulate sequences of aseismic slip in a Cascadia-like model setup following Liu and Rice (2009), and expand their exploration of the model parameter space. Our ROM methodology combines an efficient spline-based latent representation of rate-and-state friction SSE cycle models outputs, leveraging their phase-space characteristics, with Proper Orthogonal Decomposition (POD, Bui-Thanh et al. (2003)) and coefficient-interpolation using Radial Basis Functions (RBFs, Audouze et al. (2009)). This approach reduces the computational cost by $\sim 3.6 \times 10^5$ times compared to performing a physics-based tandem simulation, enabling comprehensive parameter space exploration and uncertainty quantification of SSE characteristics. We perform a full Bayesian Markov chain Monte Carlo (MCMC) inversion constrained by Northern Cascadia SSEs characteristics, constraining the width (44.7 ± 16.2 km) and magnitude (3.8 ± 1.4 MPa) of a deep low effective normal stress region. We demonstrate the accuracy and efficiency of our ROM framework and highlight its potential for advancing the physics-based understanding and simulation of complex fault slip behavior.

Table 1: Summary of tandem SSE cycle forward simulation parameters.

Symbol	Definition	Value	Reference
W_l	Up-dip extent of low effective normal stress zone	varied $[152, 189.5]$ km	
W_r	Down-dip extent of low effective normal stress zone	252 km along dip	
σ_0	Fault normal stress outside of $[W_l, W_r]$	50 MPa	Liu and Rice (2009)
σ_W	Fault normal stress within $[W_l, W_r]$	varied [1, 6] MPa	See Eq. (3)
L_0	Characteristic slip distance outside of $[W_l, W_r]$	13 mm	Liu and Rice (2009)
L_W	Characteristic slip distance within $[W_l, W_r]$	0.16 mm	Liu and Rice (2009)
a	Direct effect parameter	varies along-fault, see Figure 1	He et al. (2006)
b	Evolution effect parameter	0.0045	He et al. (2006)
μ	Shear modulus	33.91 GPa	
ν	Poisson's ratio	0.25	
f_0	Reference coefficient of friction	0.6	
$\dot{s}_0$	Reference slip rate	1×10^{-6} m/s	

2 Full-order models of slow slip cycles in Cascadia

To investigate the time-dependence of repeating SSEs in a Cascadia-like subduction zone, we first describe a full-order model (FOM) simulation framework. The fault stress evolution in this model follows the regularized form of the rate-and-state friction law (Dieterich, 1979; Lapusta et al., 2000; Rice and Ben-Zion, 1996; Ruina, 1983), given by

$$\tau(\dot{s}, \psi) = a \operatorname{arcsinh}\left(\frac{\dot{s}}{2\dot{s}_0} \exp\left(\frac{\psi}{a}\right)\right) \tag{1}$$

where the fault shear stress τ is a function of the slip rate $\dot{s}$ (with $\dot{s} = \|\dot{s}\|$) and a state-variable ψ, with a representing an empirical friction parameter to describe the "direct effect" and $\dot{s}_0$ a reference slip rate. The state-variable evolves temporally according to the Dietrich-Ruina aging law

$$\frac{d\psi}{dt}(\dot{s}, \psi) = \frac{b\dot{s}_0}{L}\left(\exp\left(\frac{f_0 - \psi}{b}\right) - \frac{\dot{s}}{\dot{s}_0}\right) \tag{2}$$

where b is an empirical frictional parameter describing the time-dependent "evolution" effect, L denotes the characteristic slip distance, and f_0 is a reference friction coefficient.

The evolution of and interaction between slip rate and state-variable can be analyzed in a two-dimensional phase-space defined by these variables, which offers a simplified framework for examining the system's inherent instability (Ranjith and Rice, 1999). For VW faults, trajectories in this phase-space may form either closed or open orbits (i.e., limit cycles, e.g., (Pranger et al., 2022)), reflecting stable oscillatory behavior or runaway slip, respectively. The geometry of phase-space trajectories provides a clearer, lower-dimensional view of the complex faulting behavior than the time-domain (Ciardo and Viesca, 2024; Viesca, 2016a). We will later utilize this phase-space representation of the rate-and-state friction SSE cycle simulations as a more suitable representation for reduced-order modelling.

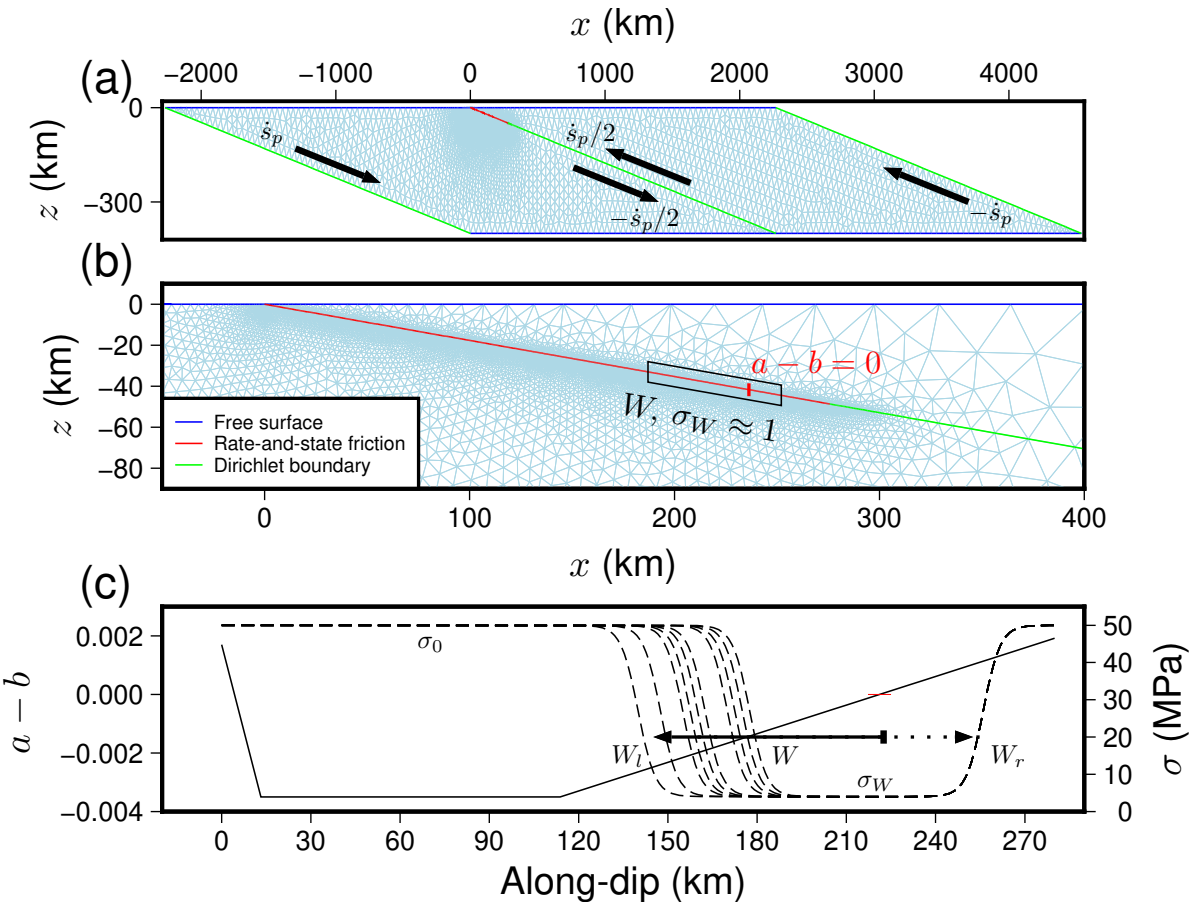

Figure 1: Model setup and unstructured triangular mesh for the full-order SSE cycle models using tandem. (a) View of the computational domain, boundary conditions and applied loading, where $\dot{s}_P$ is the applied loading rate. Not to scale, note the 15× vertical exaggeration of the y-axis. (b) Zoomed view of the rate-and-state friction fault, showing the adaptively refined mesh, with W and σ_W being the width and magnitude of a low effective normal stress zone respectively and a, b are rate-and-state friction parameters. (c) Along-fault profiles of $a-b$ derived from gabbro friction experiments (He et al., 2006) and the initial effective normal stress distribution along the fault where W_l and W_r are the up-dip and down-dip extent of the low effective normal stress zone, σ_0 is the effective normal stress outside of this zone and red horizontal bar mark the VW-VS transition point. Profiles of normal stress are shown for $\sigma_W = 4$ MPa.

All models in this study consist of two-dimensional parallelogram-shaped domains measuring 4500 km in length and 400 km in width, containing a single planar fault that bisects the domain into two geometrically identical regions (Figure 1(a), see governing equations in Appendix B). This setup follows the configuration of the BP3 SCEC community benchmark (Erickson et al., 2023). To approximate the CSZ, both the fault and lateral domain boundaries are inclined at a shallow dip angle of 10^o relative to the upper and lower domain boundaries consistent with the Slab2 model (Hayes et al., 2018). The updip edge of the fault intersects the upper free surface, while the down-dip edge terminates at the bottom boundary, which is also treated as a free surface.

Loading is prescribed as steady tectonic convergence, (Figure 1(a), Biemiller et al., 2024), imposed via Dirichlet boundary conditions, with velocities of $\dot{s}_p = (3.64, 0.64)$ cm/year in the x and z components, respectively, to the fault at depths exceeding 280 km along-dip. In addition $\pm\dot{s}_p/2$ is applied to the lateral domain boundaries. Fault slip behavior governed by rate-and-state friction can be classified into two regimes based on the difference of the frictional parameters $a-b$. When $a-b<0$, the fault exhibits VW behavior, which promotes the release of accumulated stress through rapid seismic slip. In contrast, when $a-b>0$, the fault is VS, favoring stable creep (Scholz, 1998).

SSEs are emerging in our model framework by applying a reduced effective normal stress (σ_W) to a localized subregion of the rate-and-state friction fault, which encompasses the stability transition zone ($a-b=0$) between VW (up-dip) and VS (down-dip) frictional behavior (1), following the approach established in previous studies (e.g., Liu and Rice, 2005, 2007, 2009; Rubin, 2008). The width of this region will be denoted as W.

The initial effective normal stress along the fault (Figure 1(c)) is given by

$$\sigma(d) = \left(\frac{1}{1+\exp(-\gamma(d-W_l))} - \frac{1}{1+\exp(-\gamma(d-W_r))} \right)(\sigma_W - \sigma_0) + \sigma_0, \tag{3}$$

where d is the distance along the fault in the down-dip direction with $d=0$ at the top free surface, σ_0 is the normal stress outside the low effective normal stress zone which is taken to be 50 MPa. The parameter γ is a smoothing coefficient controlling the smoothness of the transition between the two normal stress regions and taken to be 0.35, while W_l and W_r represent the up-dip and down-dip boundaries of the region of low effective normal stress. Our domain geometry and boundary conditions ensure that the effective normal stress remains approximately constant, varying by no more than a few kPa throughout our simulations.

We follow Liu and Rice (2009) and define the width of the low-effective normal stress zone (W) as the distance from the stability transition point ($a-b=0$) to its up-dip termination (W_l). In the following, we build reduced order models to represent variations in both the width (by varying W_l) and normal stress reduction (varying σ_W) of this region (Figure 1). Other model parameters remain fixed.

We use frictional parameters resembling experimentally derived values for gabbro (He et al., 2006). The parameter b is held constant throughout the domain, while a is varied to achieve the target $a-b$ values (Figure 1) consistent with the experimental data. As characteristic slip distance, we assume L_W= 0.16 mm within the low effective normal stress zone and L_0=13 mm outside of this zone. All model parameters are listed in Table 1.

We employ the open-source SEAS simulation software tandem (Uphoff et al., 2022) to conduct quasi-dynamic simulations of SSEs in this two-dimensional Cascadia-like setting (Liu and Rice, 2009). Each simulation constitutes a full-order model (FOM), where simulations are run to contain at least 20 SSE cycles, after a spin up phase. tandem utilizes a discontinuous Galerkin (DG) finite element method and supports curvilinear elements within an unstructured mesh composed of triangular elements. Preconditioned Krylov methods from the PETSc-TAO library (Balay et al., 1997, 2025a,b) are used to solve the elasticity problem arising from the DG spatial discretization. Time integration of the rate-and-state friction ODEs within tandem uses an embedded sixth-order Runge-Kutta scheme (Abhyankar et al., 2018). We implement

time step adaptivity by using an error estimator obtained from the difference between the 5th order and 6th order solution following (Dormand and Prince, 1980).

For all simulations, the on-fault resolution is chosen to resolve the smaller of the characteristic nucleation size, h^*, and the process zone size, Λ, by a factor of $20/d_{\text{poly}}$, where $d_{\text{poly}} = 6$ is the polynomial degree of DG basis functions used, thereby ensuring adequate spatial resolution (Uphoff et al., 2022). The simulations are initiated by setting the state-variable ψ to a value that satisfies the governing rate-and-state friction law for the prescribed initial stress conditions and slip rate, respectively.

In a general setting, the FOM described above is a set of coupled parametric ODEs coupled to a PDE. In any given study designed to explore model parameters, one is required to hold some parameters constant and vary others. To generalize the presentation of the reduced-order model in Section 3 we denote the parameters varied via

$$\xi = \left\{\xi_1, \ldots, \xi_{\dim(\mathscr{P})}\right\}. \tag{4}$$

In this work we consider a two-dimensional parameter space ($\dim(\mathscr{P}) = 2$) with $\xi_1 = W \in [30.5, 68]$ km and $\xi_2 = \sigma_W \in [1, 6]$ MPa. All other parameters such as a, b, L (Table 1 for a complete list) are not varied. We will denote our 2D parameter space by $\mathscr{P}$.

3 SEAS model order reduction

Our reduced-order model (ROM) methodology is designed to provide an approximate value of slip, slip rate and state-variable at locations along the fault as a function of time t and the parameters $\xi = (W, \sigma_W)$. An overview of our ROM methodology is shown in Figure 2.

The construction of our ROM starts by uniformly sampling the parameter space within prescribed minimum and maximum values for each ξ_i. We then evaluate the FOM at different parameter realizations ξ_i, $i = 1, \ldots, N$ to obtain N simulation outputs, each being denoted by $Q(\xi_i)$ (step 1 in Figure 2). The N outputs consist of time discrete values of slip, slip rate and state-variable at different locations along the fault, that is $Q(\xi_i) = (\mathbf{t}, \dot{\mathbf{S}}, \Psi, \mathbf{S})$, where $\mathbf{t}$ is a vector of simulation time stamps of length n, and $\dot{\mathbf{S}}, \Psi, \mathbf{S} \in \mathbb{R}^{n \times m}$ are matrices representing slip rate, state-variable, and cumulative slip, respectively, across m fault observation points. We will refer to the time series of length n for quantities at observation points $j = 1, \ldots, m$ via the sub-script j. That is, $\mathbf{S}_j$ denotes the time series for the slip at the j^{th} observation point and thus $\mathbf{S} = [\mathbf{S}_1, \ldots, \mathbf{S}_m]$ and $Q_j = (\mathbf{t}, \dot{\mathbf{S}}_j, \Psi_j, \mathbf{S}_j)$.

A key aspect of our ROM strategy is that, instead of reducing the dimensionality of simulations defined over fixed time intervals, we reduce the dimensionality over each SSE cycle independently. Each trajectory $Q(\xi_i)$ is segmented into p individual cycles $Q^k(\xi_i)$, $k = 1, \ldots, p$ by applying a slip rate threshold criterion at a pre-selected location on the fault (step 2 in Figure 2). This location is set to 195 km along-dip, a position chosen to be within the VW region of the low effective normal stress zone (W) for all FOM simulations. Whenever the slip rate at this location exceeds a threshold of 1×10^{-4} m/s, a new cycle boundary is defined, thereby triggering a split in the data sequence. To prevent the clustering of multiple triggers within a single event, we impose an additional temporal constraint which ensures that no new cycle boundary can be defined within a three-month period following the previous one. This period is well below the expected recurrence interval of the SSE and thus ensures that each distinct SSE is identified as a single cycle.

The construction of the ROM corresponding to each individual k SSE cycle is detailed in Sections 3.1.1 and 3.1.2. Once constructed, each cycle-specific ROM defines a mapping

$$\text{ROM}^k(\xi^*) = Q^k(\xi^*), \quad \text{for } k = 1, \ldots, p, \tag{5}$$

where $\text{ROM}^k(\cdot)$ represents the reduced-order model for the k^{th} cycle, $\xi^* \in \mathscr{P}$ denotes an arbitrary parameter vector in which $\xi^* \neq \xi_i$, $i = 1, \ldots, N$ and $Q^k(\xi^*)$ is the resulting ROM approximation for the complete time

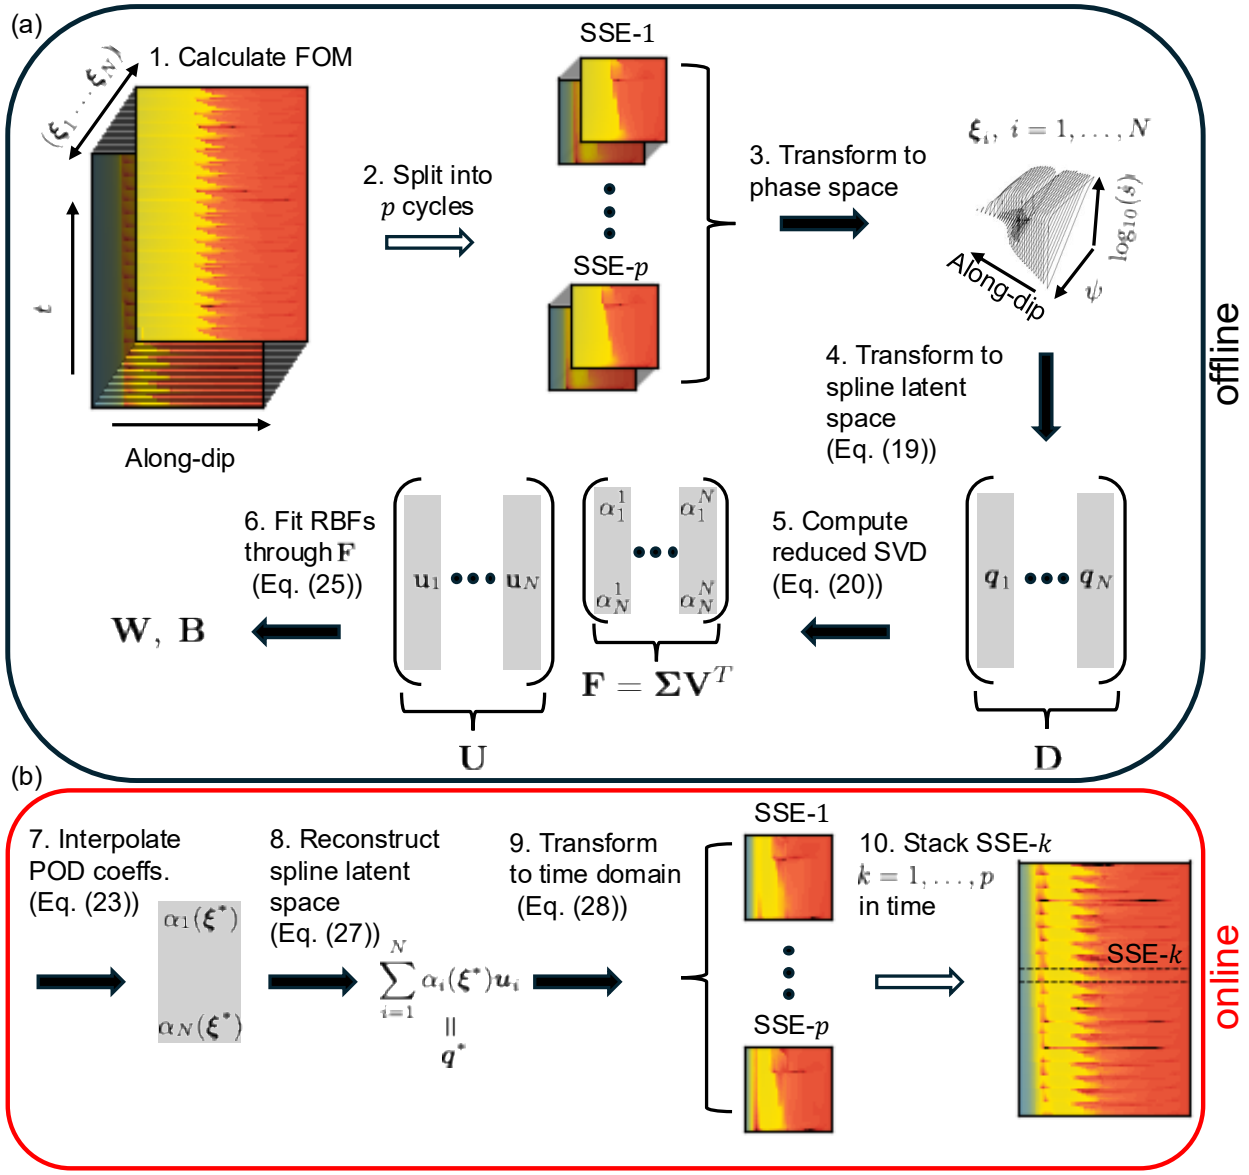

Figure 2: Overview of the SEAS ROM methodology. (a) The offline component of the ROM framework, which is performed once during initialization. This computationally intensive setup stage primarily involves data collection by performing N FOM simulations with parameters vectors ξ_i (step 1). This is followed by latent space encoding (q_i) via spline interpolation (step 4). Subsequently, dimensionality reduction using POD of $\mathbf{D}$, the stacked latent vector matrix, via SVD, where $\mathbf{D} = \mathbf{U}\Sigma\mathbf{V}^T$ (step 5). The columns of $\mathbf{U}$ ($\mathbf{u}_i$) form the POD basis. The ROM is constructed by interpolating the POD coefficients (α_r^i) along the rows of the POD coefficient matrix $\mathbf{F}$ using RBFs with polynomial terms, defined by RBF weights $\mathbf{W}$ and polynomial coefficient $\mathbf{B}$ matrices. (b) The online component of the ROM framework, which enables fast inference of ROM-based simulations for any new parameter vector ξ^*, where the * superscript distinguishes ROM inputs and outputs from their FOM counterparts, within the bounds of the training set. Steps 3 through 9 are performed for each $k = 1, \ldots, p$ cycles and marked with solid arrows.

series of the slip rate, state and slip at all fault observation points m. The superscript $*$ is used to distinguish ROM inputs from the full-order model inputs ξ_i, $i = 1,\ldots,N$ used to construct the ROM. To reconstruct a full multi-cycle simulation sequence for a given parameter ξ^*, we evaluate each of the p individual cycle ROMs to obtain $Q^1(\xi^*),\ldots,Q^p(\xi^*)$ and concatenate them to form a continuous approximation of the complete time series for slip, slip rate and state-variable (step 10 in Figure 2).

While the uniform sampling of our parameter space ensures broad coverage, it does not account for the possibility that the underlying function approximated by the ROM may locally exhibit high gradients, or non-smooth behavior. In such regions of the parameter space, the ROM may yield large approximation errors. To address this, we adopt an empirical refinement strategy wherein the uniformly sampled parameter space is iteratively refined in regions exhibiting high ROM approximation errors. Errors are estimated using leave-one-out cross-validation (LOOCV). Regions which are identified as having a high error are resampled, e.g., refined in parameter space, by performing additional FOM simulations, and rebuilding the ROM (Section 3.2). The iterative refinement continues until the ROM leave-one-out errors fall within the range of SSEs recurrence interval and potency (P_0) standard deviations of the corresponding FOM simulations. The potency is defined as

$$P_0 = \int_A \|s^+ - s^-\|_2 \, dA, \tag{6}$$

where A is the fault surface area and s^- and s^+ are the cumulative slip before and after each SSE event respectively (Ben-Zion and Zhu, 2002).

3.1 Reduced-order models of slow slip cycles in Cascadia

3.1.1 Simulation latent space representation

Directly reducing the order of the FOM is challenging for two reasons. First, the simulations are complex, spanning tens of orders of magnitude in slip rate and involving adaptive time steps ranging more than 6 order of magnitude. Second, each N simulation output ($Q(\xi_i)$) amounts to approximately 2 GB of floating-point data, making it impractical to apply a ROM scheme directly to the raw simulation output. Instead, we propose to use a B-spline interpolant (spline latent space hereafter) to represent the simulation output, which in general constitutes a lossy compression framework.

Casting spline interpolation as a compression framework aligns with a broader challenge in scientific computing, where high-performance computing (HPC) applications generate increasingly large datasets, making storage and analysis difficult. Data compression provides a viable solution by reducing the data size. While lossless compression preserves all original information, its compression is usually not efficient for scientific simulation data, which lack the necessary redundant patterns (Lindstrom, 2014). In contrast, lossy compression, such as the spline latent space approach, achieves higher compression factors by selectively discarding less critical information under user-defined error constraints (Tao et al., 2017). This enables efficient representation while preserving essential features for further analysis.

To formalize the lossy compression framework, we define a transformation of simulation data Q_j (the parameter vector is omitted for notational clarity) into a latent space representation $q_j \in \mathbb{R}^l$ at a single observation point (j) along the fault via

$$G(Q_j) = q_j, \quad G^{-1}(\mathbf{t}, q_j) \approx Q_j \quad \text{such that} \quad \mathscr{F}(Q_j, G^{-1}(\mathbf{t}, q_j)) > 1 - \varepsilon, \tag{7}$$

where $\mathscr{F}(\cdot,\cdot)$ is the fidelity criterion with tolerance $\varepsilon << 1$, $G(\cdot)$ and $G^{-1}(\cdot,\cdot)$ are the transformations defined in this subsection (Eqs. (13), (14), (15)) and G^{-1} is evaluated over $\mathbf{t} \in Q_j$. For this study, this criterion is met if the Pearson Correlation Coefficient (PCC) between its original slip rate time series $\dot{\mathbf{S}} \in Q_j$, and its corresponding reconstructed slip rate time series $\tilde{\dot{\mathbf{S}}}_j \in G^{-1}(\mathbf{t}, q_j)$ is greater than $1 - \varepsilon$, where we have used $\varepsilon = 0.001$.

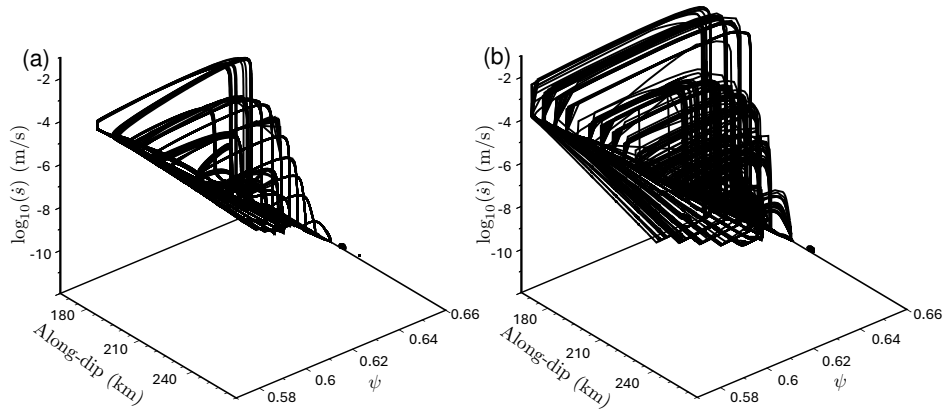

Figure 3: Phase-space H (Eq. (10)) as a function of the along-dip distance of FOM simulation results. (a) FOM results using $\xi_1 = (W = 37.375$ km, $\sigma_W = 2$ MPa). (b) FOM results using $\xi_2 = (W = 48$ km, $\sigma_W = 4$ MPa). Both panels show trajectories plotted for observation points sampled every 6 km. For each observation point, the trajectory forms a nearly closed loop. The time-domain behavior of these results is depicted in Figure 4.

Seismic cycle simulations often span a wide range of time steps and slip velocities, covering several orders of magnitude. This variability makes direct spline placement and coefficient estimation over $t \to \|\dot{s}(t)\|$ particularly challenging. However, the simulation data in phase-space, along the parametric curve

$$H(t) = (\log_{10}(\dot{s}(t)), \psi(t)) \tag{8}$$

exhibit a well-behaved, loop-like structure (Figure 3). Consequently, SEAS simulation data can be compressed more efficiently via splines when the output is represented in phase-space rather than time. We achieve this by first parametrizing time t with a variable ϕ which is defined by:

$$\phi(t) \in [0,1] \quad \text{s.t.} \quad \phi(t_1) = 0 \text{ and } \phi(t_n) = 1. \tag{9}$$

such that ϕ represents the time progression along the trajectory in the phase-space. As this phase-space trajectory $H(t)$ is defined by the physically continuous variables $\log_{10}(\dot{s}(t))$ and $\psi(t)$, $\phi(t)$ is itself a continuous and monotonic function of time. We then define

$$\begin{aligned} H(\phi) &= H(\phi(t))) \\ &= \Big(\log_{10}\Big(\dot{s}(\phi(t))\Big), \psi(\phi(t))\Big). \end{aligned} \tag{10}$$

Using a spline interpolation, we define an interpolant for ϕ in terms of t. Similarly, we also can define a spline interpolant for slip rate, state and slip in terms of ϕ. This two-step procedure results in the construction of four B-splines, constructed via least-squares fitting. These splines will be used to define the forward transformation G. The inverse transformation G^{-1} is then defined by first evaluating spline interpolant to approximate $\phi(t)$, and this result is used to evaluate the splines approximating the slip rate, state and slip (as functions of ϕ).

Definition of G and G^{-1}

Given a set of data points whose elements consist of paired values of independent (x) and dependent variables ($f(x)$) denoted by $\{\hat{x}_i, \hat{f}_i\}_{i=1}^n = (\hat{\mathbf{x}}, \hat{\mathbf{f}})$ where $\hat{f}_i = f(\hat{x}_i)$, then the construction of the B-Spline is defined by

$$\mathscr{B} : \hat{\mathbf{x}}, \hat{\mathbf{f}} \to \mathbf{k}_x, \mathbf{c}_f, \tag{11}$$

where $\mathbf{k}_x, \mathbf{c}_f$ are the spline knots and spline coefficients respectively. The B-spline approximation of the function $f(\cdot)$ is given by

$$f(x) \approx \tilde{f}(x) = \mathscr{B}^{-1}(\mathbf{k}_x, \mathbf{c}_f, x), \quad x_1 \leq x \leq x_n. \tag{12}$$

We note that subscript on the knot vector $\mathbf{k}$ indicates the independent variable, and the subscript on the coefficient vector $\mathbf{c}$ indicates the dependent variable.

In this work we use B-spline interpolants to define a latent space transformation for the SSE simulation output. Given Q_j, the transformations $G : Q_j \to q_j$ is defined as:

$$\begin{aligned} G\left(Q_j\right) &= \left(\mathscr{B}(\mathbf{t}, \Phi_j), \mathscr{B}(\Phi_j, \dot{\mathbf{S}}_j), \mathscr{B}(\Phi_j, \Psi_j), \mathscr{B}(\Phi_j, \mathbf{S}_j)\right) \\ &= \left(\mathbf{k}_t, \mathbf{c}_\phi, \mathbf{k}_\phi, \mathbf{c}_{\dot{s}}, \mathbf{c}_\psi, \mathbf{c}_s\right) \\ &= q_j, \end{aligned} \tag{13}$$

where $\Phi_j = (\phi(t_1) \dots \phi(t_n))$ for $\mathbf{t} = (t_1, \dots, t_n) \in Q_j$, $\mathbf{k}_t, \mathbf{c}_\phi \in \mathbb{R}^{K_0}$ are the spline knot and coefficient vector for $\phi(t)$ and $\mathbf{k}_\phi, \mathbf{c}_{\dot{s}}, \mathbf{c}_\psi, \mathbf{c}_s \in \mathbb{R}^{K_1}$ are the spline knot and coefficient vector for $\dot{s}(\phi), \psi(\phi), s(\phi)$ respectively. As the same independent variable ϕ is used for the slip rate, state and slip, the knot vector $\mathbf{k}_\phi$ is only stored once in Eq. (13). In Section 4.2 we discuss how K_0, K_1 were chosen. Furthermore, in Appendix D a description of how the knot placement is defined within $\mathbf{k}_t, \mathbf{k}_\phi$ (given K_0, K_1) is provided.

For a specific time t, the spline latent vector q_j can be used to reconstruct the slip rate, state and slip. Defining this as $g^{-1} : t, q_j \to \tilde{\dot{s}}(t), \tilde{\psi}(t), \tilde{s}(t)$ we have:

$$\begin{aligned} g^{-1}(t, q_j) &= \left(\mathscr{B}^{-1}(\mathbf{k}_\phi, \mathbf{c}_{\dot{s}}, \tilde{\phi}(t)), \mathscr{B}^{-1}(\mathbf{k}_\phi, \mathbf{c}_\psi, \tilde{\phi}(t)), \mathscr{B}^{-1}(\mathbf{k}_\phi, \mathbf{c}_s, \tilde{\phi}(t))\right) \\ &= \left(\tilde{\dot{s}}(t), \tilde{\psi}(t), \tilde{s}(t)\right), \end{aligned} \tag{14}$$

where $\tilde{\phi}(t) = \mathscr{B}^{-1}(\mathbf{k}_t, \mathbf{c}_\phi, t)$ is the B-spline approximation for ϕ.

Lastly we have the inverse transformation given by

$$G^{-1} : \tilde{\mathbf{t}}, q_j \to \tilde{Q}_j = \left(\tilde{\mathbf{t}}, \tilde{\dot{\mathbf{S}}}_j, \tilde{\Psi}_j, \tilde{\mathbf{S}}_j\right), \tag{15}$$

where $\tilde{\mathbf{t}} \in \mathbb{R}^{n^*}$ is an arbitrary time vector which in general is different to $\mathbf{t}$, and $\tilde{\dot{\mathbf{S}}}_j$, $\tilde{\Psi}_j$, $\tilde{\mathbf{S}}_j$ are the length n^* vectors containing the spline approximations for slip rate, state and slip at observation point j for times $t = \tilde{t}_1, \dots, \tilde{t}_{n^*}$. The entries of $\tilde{\dot{\mathbf{S}}}_j$, $\tilde{\Psi}_j$, $\tilde{\mathbf{S}}_j$ are computed using $\tilde{\dot{s}}(t), \tilde{\psi}(t), \tilde{s}(t)$ obtained from applying g^{-1} pointwise.

The transformations from simulation output into a spline latent space at a specific observation point could easily be expanded to all simulation fault observation points by applying G on each specific observation point and letting the latent space vector q store the spline knots and coefficients of all observation points. For notation simplicity we will use from here on $q \in \mathbb{R}^l$, G, g^{-1} and G^{-1} to represent the transformation of the entire simulation to latent space with

$$G(Q) = (\mathbf{k}_t, \mathbf{c}_{\phi 1}, \mathbf{k}_{\phi 1}, \mathbf{c}_{\dot{s}1}, \mathbf{c}_{\psi 1}, \mathbf{c}_{s1}, \dots, \mathbf{c}_{\phi m}, \mathbf{k}_{\phi m}, \mathbf{c}_{\dot{s}m}, \mathbf{c}_{\psi m}, \mathbf{c}_{sm}) = q \tag{16}$$

where $Q = (\mathbf{t}, \dot{\mathbf{S}}, \Psi, \mathbf{S})$ with $\dot{\mathbf{S}}, \Psi, \mathbf{S} \in \mathbb{R}^{m \times n}$ as defined in Section 2. We note that all fields (slip rate, state and slip) and at all observation points use the same vector of time values $\mathbf{t}$, hence for efficiency the knot vector $\mathbf{k}_t$ can be reused in Eq (16) and thus only appears once. Therefore, the length l of the spline latent space representation q is given by:

$$l = (1 + m)K_0 + 4mK_1, \tag{17}$$

where K_0 is the dimension of the knot vector $\mathbf{k}_t$ and each of the m coefficient vectors $\mathbf{c}_{\phi j}$. Similarly, K_1 is the dimension of each of the m knot vectors $\mathbf{k}_{\phi j}$ and their corresponding coefficient vectors $(\mathbf{c}_{\dot{s}j}, \mathbf{c}_{\psi j}, \mathbf{c}_{sj})$. The

K_0 and K_1 values that we find to be sufficient for transforming the presented Cascadia-like SEAS simulation to the latent-space and back are presented in Section 4.2

In the same fashion we will expend g^{-1} to be

$$g^{-1}(t,q) = \left(\begin{pmatrix} \tilde{s}_1(t) \\ \vdots \\ \tilde{s}_m(t) \end{pmatrix}, \begin{pmatrix} \tilde{\psi}_1(t) \\ \vdots \\ \tilde{\psi}_m(t) \end{pmatrix}, \begin{pmatrix} \tilde{s}_1(t) \\ \vdots \\ \tilde{s}_m(t) \end{pmatrix} \right) \in \mathbb{R}^{m\times 3} \tag{18}$$

resulting in three column vectors of B-splines evaluations at time t at all m observation points.

This spline-based latent space provides a compact representation of complex seismic cycle simulations. In the following, we demonstrate how this latent encoding can be leveraged to construct a ROM that accurately emulates the full-order simulations with orders of magnitude lower computational cost.

3.1.2 One SSE cycle reduced-order model

To construct a ROM for an individual SSE cycle, we adopt a strategy similar to the approach of Rekoske et al. (2025), originally developed for seismic wave propagation. Here, we extend this methodology to one cycle of SSE simulation.

Let us consider a collection of N full-order model simulations, each corresponding to a distinct parameter realization $\xi_1, \ldots, \xi_N$. For a given SSE cycle k, these simulations yield full-order outputs $Q^k(\xi_1), \ldots, Q^k(\xi_N)$. Using Eq. (16) we map each simulation output into its spline latent representation to form a matrix $\mathbf{D} \in \mathbb{R}^{l\times N}$ (step 4 in Figure 2) given by

$$\mathbf{D} = \Big(G(Q^k(\xi_1)) \ldots G(Q^k(\xi_N)) \Big) = \Big(q^k(\xi_1) \ldots q^k(\xi_N) \Big), \tag{19}$$

where $q^k(\xi_i)$ is the spline latent column vector representation of the i^{th} simulation in the k^{th} cycle. We then compute the Singular Value Decomposition (SVD) of $\mathbf{D}$:

$$\mathbf{D} = \mathbf{U}\Sigma\mathbf{V}^T, \tag{20}$$

where $\mathbf{U}$ contains the left singular vectors (basis functions), Σ is a diagonal matrix of singular values λ_i, and $\mathbf{V}$ contains the right singular vectors. The SVD provides an optimal low-rank approximation in the least-squares sense and allows each latent vector $q^k(\xi_i)$ to be expressed as a linear combination of the orthonormal basis vectors in $\mathbf{U}$:

$$q^k(\xi_i) = \sum_{r=1}^{N} \lambda_r v_{ir} \mathbf{u}_r = \sum_{r=1}^{N} \alpha_r^i \mathbf{u}_r, \tag{21}$$

where $\alpha_r^i = \lambda_r v_{ir}$ are the Proper Orthogonal Decomposition (POD) coefficients. This decomposition is consistent with viewing the SVD as a POD of the matrix $\mathbf{D}$, capturing the most energetic modes of variability across the simulations dense representation (step 5 in Figure 2, Berkooz et al., 1993; Bui-Thanh et al., 2003; Druault et al., 2005).

We adopt a radial basis functions (RBFs) interpolation strategy to perform a mapping from an arbitrary $\xi^* \in \mathscr{P}$ to approximate POD coefficients (e.g., Audouze et al., 2009; Lazzaro and Montefusco, 2002; Rekoske et al., 2025; Xiao et al., 2015). To avoid bias due to differing parameter scales, we first normalize the parameter space:

$$\begin{pmatrix} \bar{\xi}_1 \\ \vdots \\ \bar{\xi}_N \end{pmatrix} = \left[\begin{pmatrix} \xi_1^1 & \cdots & \xi_1^b \\ \vdots & \ddots & \vdots \\ \xi_N^1 & \cdots & \xi_N^b \end{pmatrix} - \begin{pmatrix} \bar{\xi^1} & \ldots & \bar{\xi^b} \end{pmatrix} \right] \oslash \begin{pmatrix} \text{std}(\xi^1) & \ldots & \text{std}(\xi^b) \end{pmatrix}, \tag{22}$$

where $\bar{\xi}^r$ and std(ξ^r) denote the mean and standard deviation of the r^{th} parameter across the N samples, and $\oslash$ denotes element-wise division.

Given a new normalized parameter vector $\bar{\xi}^*$, the interpolated POD coefficient vector $\alpha(\bar{\xi}^*)$ (step 7 in Figure 2) is expressed as:

$$\alpha(\bar{\xi}^*) = \sum_{i=1}^{N} w_i \varphi(\|\bar{\xi}^* - \bar{\xi}_i\|_2) + \sum_{r=1}^{|V|} b_r p_r(\bar{\xi}^*), \tag{23}$$

where w_i are the RBFs weights and $\varphi(\cdot)$ is the RBF kernel defined in terms of the Euclidean distance (R) between points. In this study, we consider the linear kernel $\varphi(R) = R$, which provides a simple yet effective means of interpolating in the reduced space (see Text S2 and Figure S1). In Eq. (23) b_r are polynomial weights, and $p_r(\cdot)$ are the monomials of the input parameters up to degree d:

$$p_r(\bar{\xi}^*) = \xi_1^{a_1} \cdots \xi_{\dim(\mathscr{P})}^{a_{\dim(\mathscr{P})}}, \quad \text{for } a_1, \ldots, a_{\dim(\mathscr{P})} \geq 0 \text{ and } \sum_{l=1}^{\dim(\mathscr{P})} a_l \leq d. \tag{24}$$

The interpolation weights are determined by solving the following linear system (Step 6 in Figure 2):

$$\begin{pmatrix} \Phi & \Lambda \\ \Lambda^T & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{W} \\ \mathbf{B} \end{pmatrix} = \begin{pmatrix} \mathbf{F} \\ \mathbf{0} \end{pmatrix} \tag{25}$$

where

$$\Phi_{ij} = \varphi(\|\bar{\xi}_i - \bar{\xi}_j\|_2), \qquad \Phi \in \mathbb{R}^{N \times N} \tag{26a}$$

$$F_{ij} = \alpha_i^j = \lambda_i v_{ji}, \qquad \mathbf{F} \in \mathbb{R}^{N \times N} \tag{26b}$$

and $\Lambda \in \mathbb{R}^{N \times |V|}$ is the matrix of monomial evaluations at each $\bar{\xi}_i$.

Following the calculation of the POD coefficients and the RBF interpolant, the coefficient corresponding to a new parameter vector ξ^* can be retrieved by interpolating the POD coefficients corresponding to the training parameter realizations $\xi_1, \ldots, \xi_N$ used for the FOM simulations. To estimate a spline latent vector (step 8 in Figure 2) using the ROM scheme for a new normalized parameter vector $\bar{\xi}^*$, we perform

$$q^k(\xi^*) = \sum_{i=1}^{N} \alpha_i(\bar{\xi}^*) u_i, \tag{27}$$

where $\alpha_i(\bar{\xi}^*)$ is the i^{th} interpolated POD coefficient computed using Eq. (23).

The ROM-computed spline latent vector $q^k(\xi^*)$ can then be mapped back to the full-order output space (step 9 in Figure 2) by applying the transformation G^{-1} given by Eq. (13) to yield:

$$\begin{aligned} Q^k(\xi^*) &= G^{-1}(\mathbf{t}^*, q^k(\xi^*)) \\ &= \left(\mathbf{t}^*,\, g^{-1}(t_1^*, q^k(\xi^*)), \ldots, g^{-1}(t_k^*, q^k(\xi^*))\right) \\ &= \left(\mathbf{t}^*, \dot{\mathbf{S}}^{k*}, \boldsymbol{\Psi}^{k*}, \mathbf{S}^{k*}\right), \end{aligned} \tag{28}$$

where $\mathbf{t}^* \in \mathbb{R}^{n^*}$ yielding $m \times n^*$ matrices approximating the simulation output corresponding to the new parameter realization ξ^* over the n^* length time vector $\mathbf{t}^*$.

3.2 Training and validation: Leave-one-out cross validation

A leave-one-out cross-validation (LOOCV) approach is employed to assess the accuracy of our ROM predictions (Kohavi, 1995; Rippa, 1999). Under this framework, each parameter realization ξ_i is systematically excluded from the training set, the model is constructed using the remaining $N-1$ simulations, and prediction accuracy is evaluated against the withheld simulation. This process iterates across all N parameter realizations to obtain an unbiased estimate of generalization error.

For each ξ_i and each SSE cycle k, we consider the POD coefficient vector α_r the r^{th} row of the F matrix as our validation set and perform the procedure described in Eq. (23) to Eq. (28) with $\bar{\xi}^* = \bar{\xi}_i$. This yields the predicted coefficients $\alpha_r(\bar{\xi}_i^*)$ and the corresponding ROM output $Q^k(\bar{\xi}_i^*)$ for the k^{th} cycle.

To assess the accuracy of the ROM predictions, we:

1. Compare the recurrence intervals of the full-order and ROM simulations using the average cycle durations:
$$T_c = \frac{1}{p}\sum_{k=1}^{p}\left(\max(t_k) - \min(t_k)\right), \quad t_k \in Q^k(\xi_i), \tag{29}$$
and similarly for the predicted outputs:
$$T_c^* = \frac{1}{p}\sum_{k=1}^{p}\left(\max(t_k^*) - \min(t_k^*)\right), \quad t_k^* \in Q^k(\bar{\xi}_i^*). \tag{30}$$

2. Compare the SSEs mean potency P_0 (Eq. (6)). In a 3D context, the potency can be related to the seismic moment by $M_0 = \mu P_0$, where μ is the shear modulus. In our 2D model the potency can be estimated without further assumptions where $\int_A$ is a line integral, we can approximate the seismic moment by assuming some fault width of W_s km along strike, yielding $M_0 = \mu W_s P_0$.

4 Results

4.1 FOM results

We performed a total of 76 FOM evaluations using the computational resources of the Nautilus Kubernetes cluster of the National Research Platform. Our simulation suite was initiated with a relatively coarse sampling grid spanning $W \in [30.5, 68]$ km and $\sigma_W \in [1, 6]$ MPa. Subsequently, the sampling density was increased within the range $W \in [30.5, 53]$ km, and further refined within $W \in [35.5, 38]$ km. This adaptive refinement strategy aimed to improve the accuracy and robustness of the ROMs, particularly letting the leave-one-out cross-validation error estimation described in Section 3.2 stay within simulation T_c and P_0 standard deviation.

The computational expense per simulation exhibited considerable variation, depending on the specific values chosen for the W and the σ_W within the localized subregion around the stability transition zone $a - b = 0$. As is standard for seismic cycle simulations, an initial spin-up phase was removed to minimize the influence of initial conditions on the subsequent system time-dependence behavior (Rubin and Ampuero, 2005). The computation time required for this spin-up phase proved to be highly sensitive to the model parameters, particularly the normalized fault width

$$\frac{W}{h^*} = \frac{W\pi(1-\nu)(b-a)\sigma_W}{2\mu L}, \tag{31}$$

where h^* is the characteristic nucleation size, ν is the Poisson ratio and μ is the shear modulus (Table 1). This is consistent with the results of Liu and Rice (2009), who showed that lower W/h^* induces more

frequent SSE oscillations, necessitating longer spin-up times to allow the system to reach a quasi-steady state before the analysis period.

We establish a uniform simulation duration cutoff of $t = 75$ years for all model runs. The majority of the computational effort during each simulation was concentrated on resolving the fine-scale time-dependence associated with the SSEs themselves. Individual simulation run times ranged from 19.2 to 302.4 hours on 30 MPI ranks, corresponding to computational costs between 628 and 9125 CPU hours per simulation. The cumulative computational cost for executing the entire suite of 76 FOM simulations amounted to approximately 233,565 CPU hours.

To illustrate the typical time-dependence observed in our FOM simulations, we first examine two representative cases: simulation ξ_1 with parameters ($W = 37.375$ km, $\sigma_W = 2$ MPa) and simulation ξ_2 with parameters ($W = 48$ km, $\sigma_W = 4$ MPa), yielding W/h^* values of 5.44 and 16.73, respectively (Figure 4). Both simulations exhibit quasi-periodic SSEs: ξ_1 recurs approximately every $T_c = 0.73 \pm 0.004$ years with peak slip velocities $V_{max} \approx 10^{-2.7}$ m/s and accumulating a mean total slip of $\delta = 0.98 \pm 3.4 \times 10^{-5}$ cm per event, while ξ_2 shows $T_c = 1.36 \pm 0.07$ years, $V_{max} \approx 10^{-1.6}$ m/s, and $\delta = 2.7 \pm 8.8 \times 10^{-3}$ cm.

Analyzing these simulations in their phase-space (Figure 3) shows a simplified view of the underlying rate-and-state friction characteristics. For both cases, as an SSE nucleates, the representative phase-space trajectory rapidly transitions towards higher slip velocities. It then evolves through a characteristic loop, involving stages of rapid state evolution and slip rate changes that reflect the weakening and subsequent healing phases dictated by Eq. (2). The cycle is completed by the slow evolution during the long interseismic period, which, although dominant in the time-domain, represents a relatively small portion of the path length traced in the phase-space during one full cycle. This confirms the utility of the phase-space perspective for capturing SEAS cycles and highlighting the periods of rapid dynamic change during the events themselves, across different parameter regimes.

Figure 5 presents the resulting recurrence time (T_c) and average SSE slip (δ) as a function of the normalized fault width (W/h^*) for all 76 FOM simulations, juxtaposed with the results reported by Liu and Rice (2009) for gabbro rheology. Notably, our results, obtained using the volumetric discontinuous Galerkin finite element code tandem (Uphoff et al., 2022), show excellent agreement with those of Liu and Rice (2009), which were generated using a boundary element method.

The simulations reveal that spontaneous, quasi-periodic SSEs are generated over a wide range of the explored parameter space, corresponding to W/h^* values ranging from 3.75 to 32.54. For W/h^* values smaller than 3.75, these simulation settings do not produce distinct SSEs; the system either evolves towards steady creep or exhibits very fast recurrence intervals ($T_c \ll 1$ year) with the slip rate oscillating around the tectonic loading rate. Within the approximate range $3.75 < W/h^* < 18$, our findings are consistent with Liu and Rice (2009), we observe clear trends where both the recurrence time (T_c) and the mean slip (δ) appear to increase approximately linearly with W/h^*. However, for $W/h^* \gtrsim 18$, we observe a deviation from this linear trend, particularly for T_c, which tends to slightly decrease and then plateau. The specific W/h^* value where this transition begins appears influenced by σ_W, for instance, the deviation becomes noticeable around $W/h^* \approx 11$ for $\sigma_W = 3$ MPa, but closer to $W/h^* \approx 15$ for $\sigma_W = 6$ MPa. In contrast, the mean accumulated slip in those events (δ) seems to maintain an approximately linear dependence on W/h^* across the entire range where SSEs are observed.

These results indicate a complex dependency between the SSE time-dependence behavior and the model parameters W, σ_W. Although the normalized fault width (W/h^*) controls the time-dependence behavior at first order, our results suggest second-order independent dependencies on both W and σ_W, which we will explore further using our ROM approach.

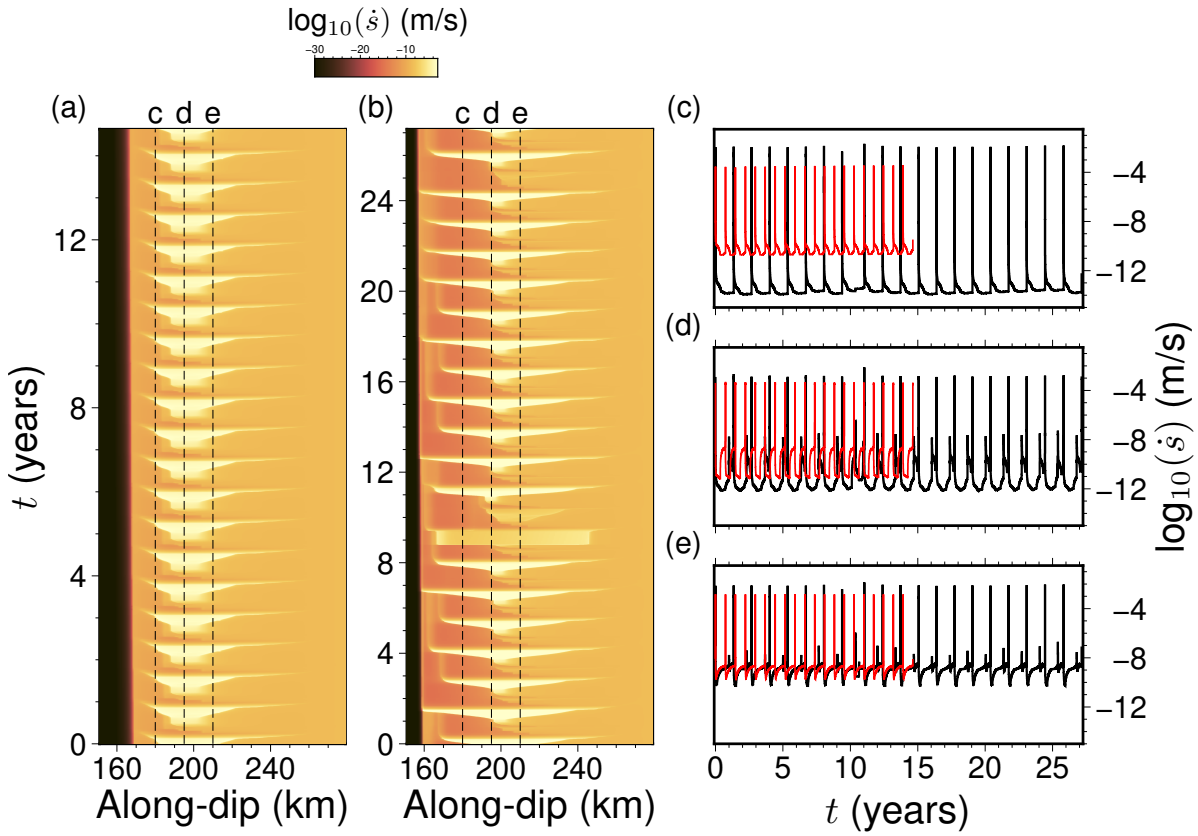

Figure 4: FOM results for two parameter vectors: $\xi_1 = (W = 37.375 \text{ km}, \sigma_W = 2 \text{ MPa})$ and $\xi_2 = (W = 48 \text{ km}, \sigma_W = 4 \text{ MPa})$. Each simulation presents 20 SSE cycles, with ξ_1 spanning a shorter time period due to its reduced recurrence interval. (a) slip rate as a function of time and along-dip distance for ξ_1. The locked portion of the subduction interface is not shown. (b) Equivalent slip rate representation for ξ_2. (c)-(e) Temporal profiles of slip rate at specific observation points along the fault (indicated by dashed lines in panels (a) and (b)), with ξ_1 shown by the red curve and ξ_2 by the black curve.

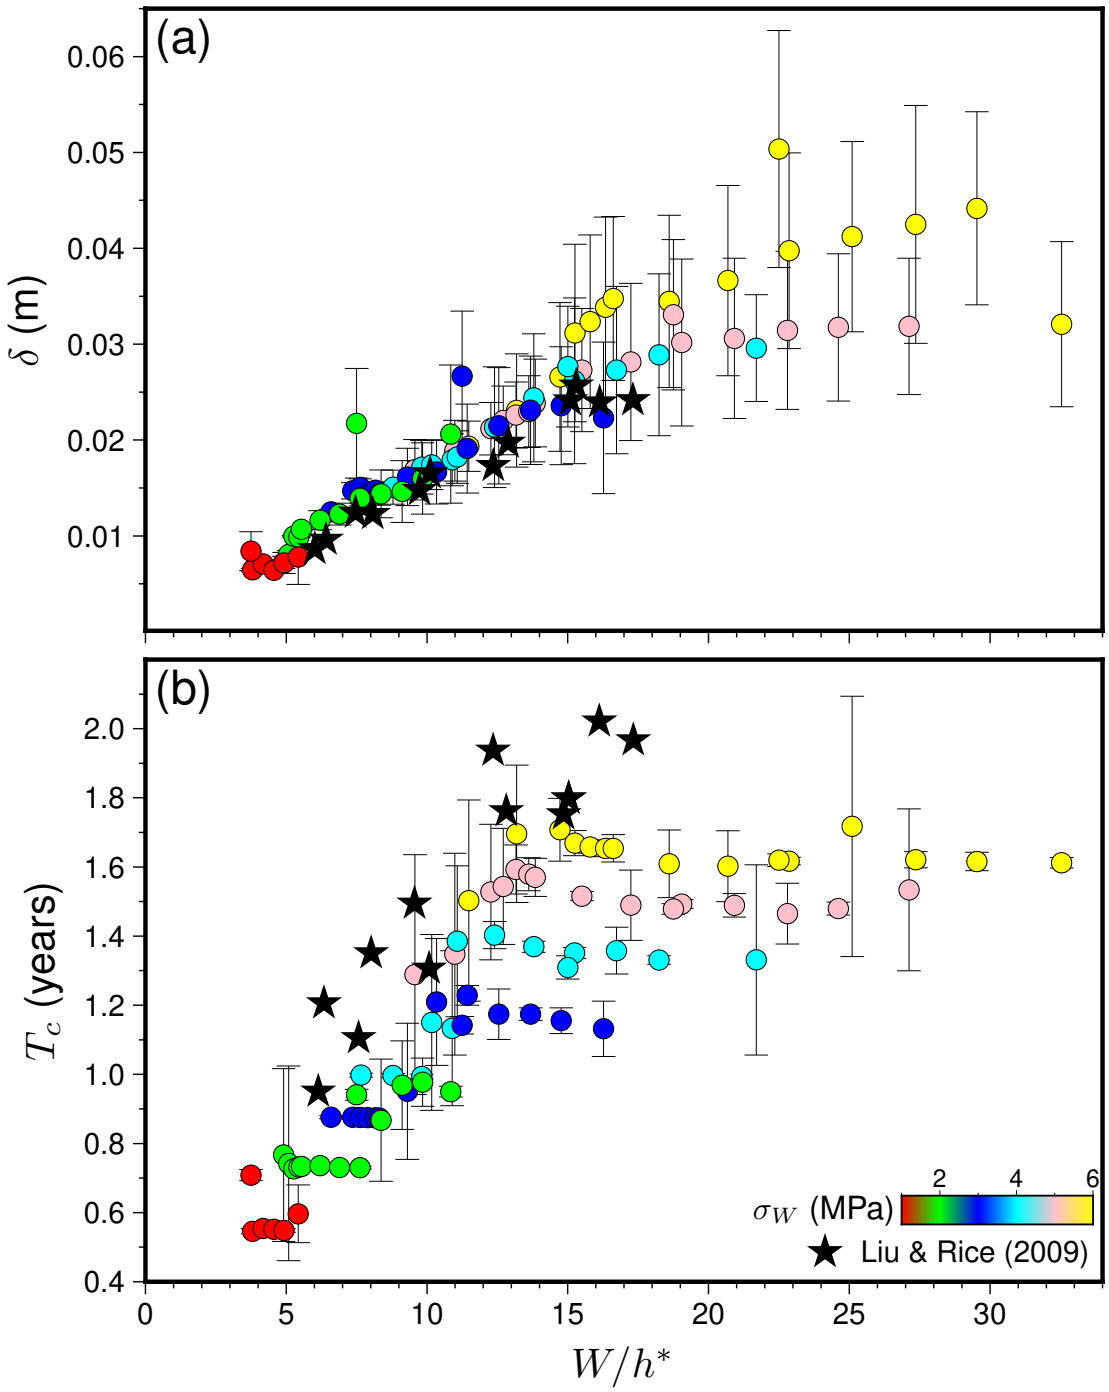

Figure 5: FOM results for key SSE characteristics as a function of normalized fault width (W/h^*). Panel (a) shows the mean SSE slip (δ) occurring within the low effective normal stress zone versus W/h^*. Panel (b) depicts the SSE recurrence time (T_c) versus W/h^*. Our FOM results using tandem (Uphoff et al., 2022) (coloured circles) are compared with the BEM results reported in Liu and Rice (2009) (black stars). The volumetric based FOM results obtained from tandem exhibit good agreement with those of Liu and Rice (2009) in their common range of W/h^*.

4.2 Latent space representation accuracy analysis

The primary goal of employing the spline latent space representation, as detailed in Section 3.1.1, is to achieve a substantial compression of the complex FOM simulation data while preserving the essential characteristics, a necessary requirement to construct an accurate reduced-order model. As our chosen representation employs lossy compression, it is important to quantify the quality of this representation by evaluating how accurately the original simulation variables can be reconstructed from the latent vector q. Following Eq. (7), we assess this accuracy primarily using the Pearson Correlation Coefficient (PCC) between the original FOM slip rate time series $\hat{S}$ within the low effective normal stress zone W and the corresponding reconstruction $\hat{\hat{S}}^*$ obtained via the inverse transformation $G^{-1}(\mathbf{t}, q)$. We aim to ensure the PCC remains above a threshold, $1-\varepsilon$ where $\varepsilon = 0.001$. The trade-off lies in minimizing the dimensionality l of the latent vector q (i.e., minimizing the number of spline knots and coefficients) while maintaining this high level of reconstruction accuracy.

Figure 6 provides a quantitative comparison validating the efficiency of representing the slip rate ($\dot{s}$) using our proposed phase domain spline approach versus a conventional time-domain spline, illustrated using four different FOM parameter pairs. The plot shows the number of spline knots required for the B-spline representation to achieve a PCC exceeding 0.999 when reconstructing the slip rate time series. This comparison considers the full simulation duration after the initial spin-up period, without constraining the number of SSE cycles included. As clearly demonstrated, splines parameterized by the phase progression variable ϕ consistently require considerably fewer spline knots to reach this high accuracy threshold compared to splines based directly on the time-domain. The four example FOM simulations contain approximately 2×10^5 to 4×10^5 time steps, requiring only 400 - 900 spline knots for satisfactory representation using the phase trajectory approach, compared to 7000 - 15000 spline knots needed for the time-domain approach. This empirically confirms that the simulation characteristics can be represented much more efficiently in the phase-space.

Building on the now established efficiency of phase-space parameterization, Figure 7 further assesses the performance of our spline latent space representation in reconstructing detailed simulation outputs. Although the figure presents the reconstruction of slip rate time series along one observation point (195 km along fault), the reconstruction PCC score is calculated over all observation points within W. For a lengthy 33-cycle SSE simulation, whose original outputs $\left(\hat{\mathbf{t}}, \hat{\Phi}, \hat{\mathbf{S}}, \hat{\Psi}, \hat{\mathbf{S}}\right)$ comprise 1.8×10^8 floating-point numbers, a clear relationship is observed between reconstruction fidelity and the number of spline knots utilized. With the ratio of knots fixed at $K_1 = 0.1K_0$, which was found to be adequate based on trial and error, a representation with $K_0 = 1000$ spline knots per observation point (illustrated by the blue curve in Figure 7(a)) reproduces the general characteristics of the slip rate waveform at a specific observation point, yet it fails to capture crucial details such as slip rate peaks and performs poorly during the interseismic period.

The target reconstruction accuracy, defined by a PCC score exceeding $1-\varepsilon$ (where $\varepsilon = 0.001$), is achieved with $K_0 = 12000$ spline knots (red curve, Figure 7(a)). This results in a latent representation size of $l \approx 4 \times 10^6$ floating-point numbers, corresponding to a 97.5% data compression. The performance of this latent space representation was also evaluated for individual SSE cycles (Figure 7(c)-(j)), which are fundamental to our ROM scheme (Section 3.1.2), across a range of model parameters (ξ_1 through ξ_4). The original output data for these single-cycle simulations range from 4×10^6 to 7×10^6 floating-point numbers. For these single-cycle cases, the target PCC threshold is met when the number of spline knots K_0 equals 250, 300, 450, and 350 per observation point for simulations using $\xi_1 = (30.5\text{ km}, 4\text{ MPa})$, $\xi_2 = (38\text{ km}, 6\text{ MPa})$, $\xi_3 = (48\text{ km}, 2\text{ MPa})$ and $\xi_4 = (68\text{ km}, 3\text{ MPa})$, respectively. This level of reconstruction yields data compression ratios between 97.3% and 97.8%, comparable to the multi-cycle simulation reconstruction.

Throughout the remainder of this work, we construct ROMs using $K_0 = 800$ and $K_1 = 80$ spline knots for the time to phase-progression ($\mathbf{k}_t$) and the phase to physical variable ($\mathbf{k}_\phi$) mapping respectively. This

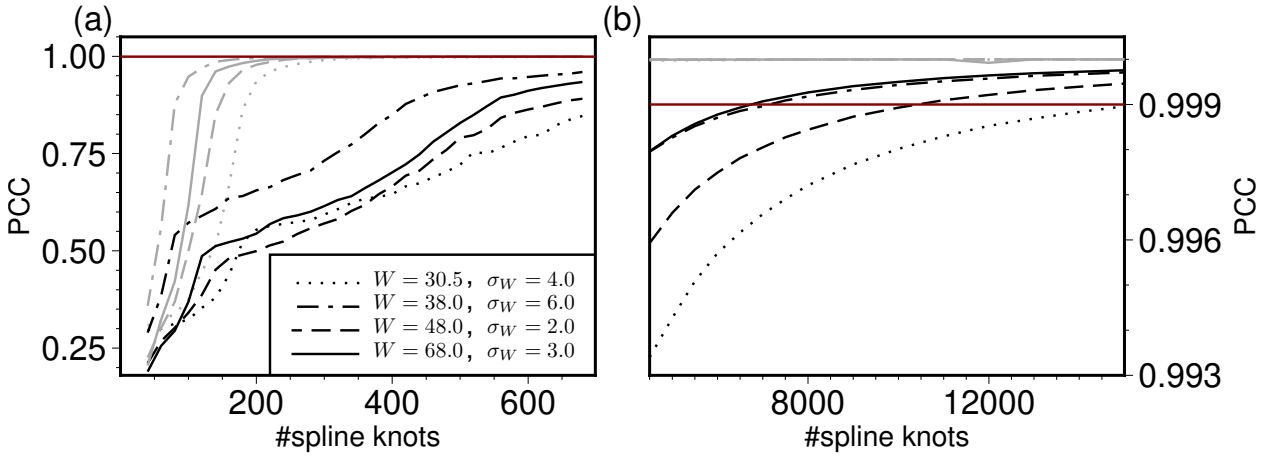

Figure 6: Comparison of spline interpolation accuracy when the slip rate ($\dot{s}$) is represented in time-domain and phase-space. The plot demonstrates the number of spline knots required to achieve a PCC exceeding 0.999 (red horizontal line) when reconstructing the slip rate. The phase-space based spline (gray lines) computed via $\mathscr{B}^{-1}(\mathbf{k}_\phi, \mathbf{c}_{\dot{s}}, \tilde{\phi}(t))$ (see Section 3.1.1, Eq. (14)) consistently requires considerably fewer spline knots than time-domain representation (black lines, computed via $\mathscr{B}^{-1}(\mathbf{k}_t, \mathbf{c}_{\dot{s}}, t)$), demonstrating the efficiency of phase-space representation of the SSE FOM simulation output. Panel (a) shows the PCC for spline calculated with number of spline knots < 700, while panel (b) illustrates the performance difference for spline calculated with number of spline knots > 5000. In the case of the gray lines, the number of knots corresponds to K_1.

choice ensures that reconstruction fidelity comfortably exceeds the desired threshold while still achieving a 88%-96% reduction in data size compared to the original FOM output ranging from 3×10^6 to 10^7 floating point numbers per cycle. With $m = 346$ fault observation points, this results in a total latent vector length of $l = 388,320$ for each SSE cycle simulation (Eq. (17)).

4.3 ROM results

The objective of developing the ROM is to accelerate the inference of SSE cycles over the parameter space defined by the width of the low effective normal stress zone (W) and its magnitude (σ_W). This section details the performance and accuracy of the constructed ROM.

The development of the ROM involves an initial offline phase, which includes running the 76 FOM simulations to generate the training dataset, transforming these simulations into their spline latent space representations, performing POD on the matrix of latent vectors $\mathbf{D}$ for each SSE cycle, and constructing the RBF interpolators for the POD coefficients (Figure 2 - offline computations). The generation of the 76 FOM simulations incurred a cumulative cost of approximately 233,565 CPU hours (Section 4.1. The subsequent offline computations for the ROM construction, of calculating spline coefficients in order to transform the FOM to their latent space require an additional $\sim$360 CPU hours all other offline steps have negligible run time compare to the first two.

Once all the offline tasks are complete, evaluating the ROM for a new parameter vector ξ^* (online task) is very fast. A single ROM evaluation, which involves interpolating the POD coefficients using the RBFs (Eq. (23)), reconstructing the latent vector q^* (Eq. (27)), and then transforming it back to the physical space using the inverse spline transformation G^{-1} (Eq. (28)) to obtain the time series for slip rate, state-variable, and cumulative slip takes approximately 30 seconds with one CPU. This represents a speedup of 3.68×10^5 compared to the average FOM runtime of $\sim$ 3000 CPUh.

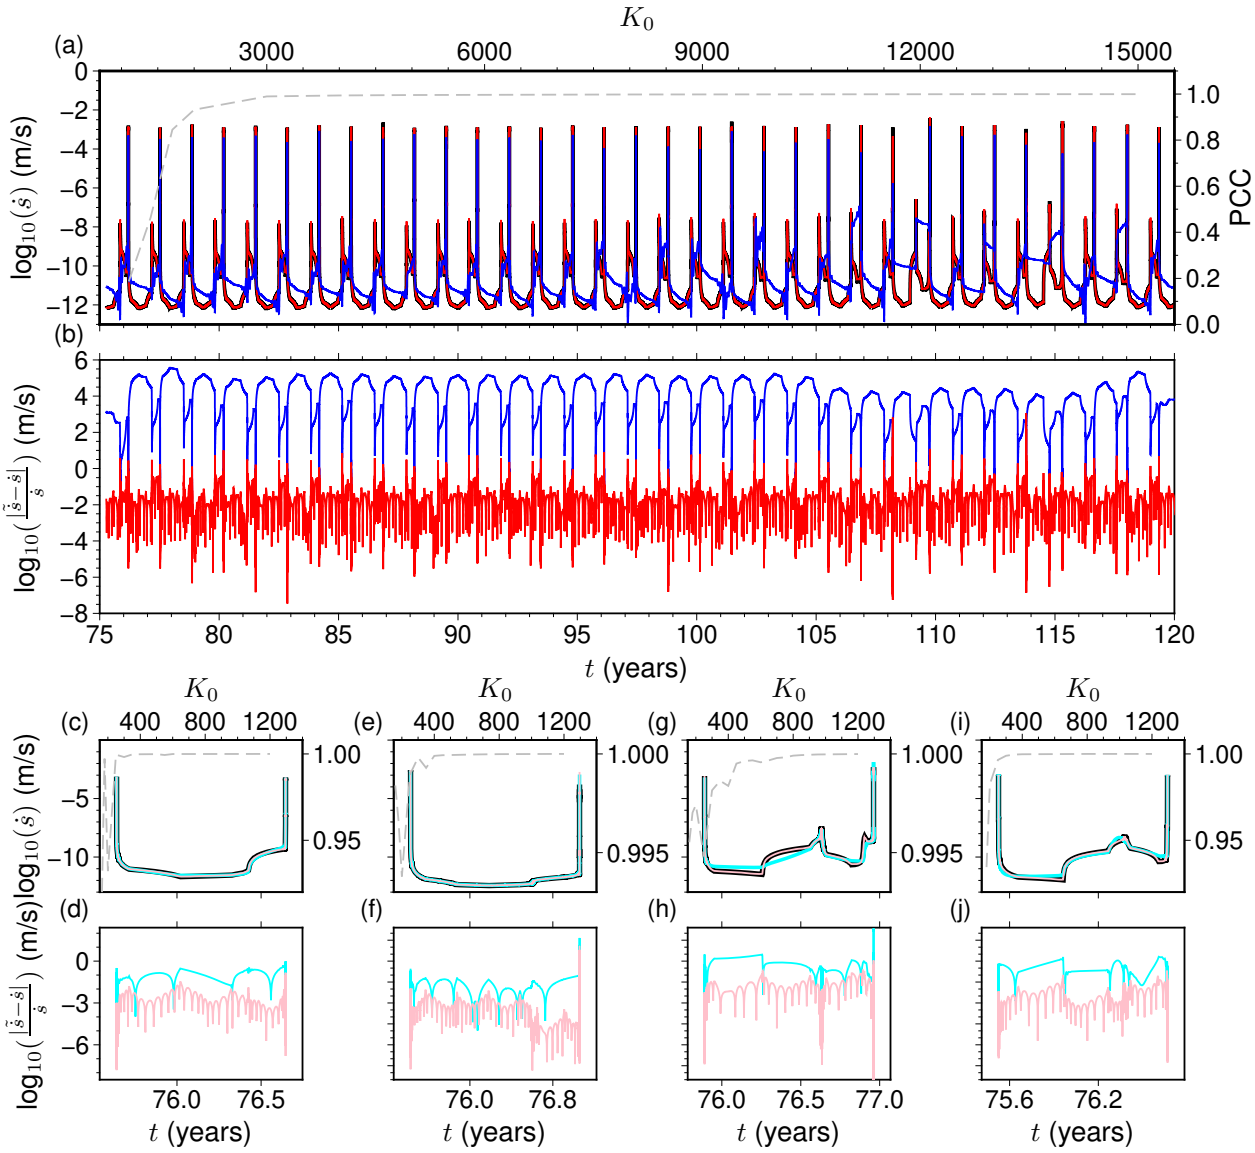

Figure 7: Accuracy of the spline latent space reconstruction of the FOM slip rate $\dot{s}(t)$. Panels (a) and (b) analyze the spline accuracy when applied to an entire time series consisting of a 33-cycle SSEs simulation with $\xi = (50.5\text{ km}, 4\text{ MPa})$ and its first 20 cycles are part of the ROM training set. Panel (a): Slip rate at an along-dip of 195 km for the FOM (black line); spline reconstruction using $K_0 = 1000$ spline knots (blue line) and $K_0 = 12000$ spline knots (red line). The gray dashed line shows the PCC (right y-axis) as a function of the number of spline knots (top x-axis). Note that the PCC is calculated for all observation points within W. Panel (b): Reconstruction errors, $|\tilde{\dot{s}}(t) - \dot{s}(t)|/\dot{s}(t)$ when using $K_0 = 1000$ (blue) and $K_0 = 12000$ (red) splines knots. Panels (c), (e), (g), (i): Spline reconstruction for slip rate when applied to a single SSE cycle using $K_0 = 200$ spline knots (cyan) and $K_0 = 800$ spline knots (pink) with ξ given by $(30.5\text{ km}, 4\text{ MPa})$, $(38\text{ km}, 6\text{ MPa})$, $(48\text{ km}, 2\text{ MPa})$ and $(68\text{ km}, 3\text{ MPa})$ respectively, all used in the ROM training data set. Panels (d), (f), (h), (j): Spline reconstruction errors associated with (c), (e), (g), (i) respectively.

The accuracy of the ROM was systematically assessed using the LOOCV procedure detailed in Section 3.2. In this process, for each of the 76 FOM simulations, the ROM was trained on the remaining 75 simulations, and its predictions for the held-out parameter set were compared against the actual FOM results. Figure 8 illustrates the ROM's capability in reproducing key physical characteristics of SSEs, namely the recurrence time T_c and potency P_0. For recurrence time, the ROM's predictions largely fall within the FOM's cycle-to-cycle standard deviation, with only one exception among the 76 cases. The mean absolute difference between FOM and ROM T_c was 11.56 days, relative to an average SSE recurrence time of 440 days for the dataset. Consequently, the R^2 score comparing ROM-predicted T_c to FOM T_c is 0.98, indicating a strong linear relationship and a near one-to-one correspondence.

For the potency, the majority of ROM predictions also align with the FOM results within the FOM's own standard deviation, which, as shown in Figure 8(d), can exhibit considerable cycle-to-cycle variability. However, notable discrepancies between FOM and ROM potency were observed for a few parameter vectors, specifically $\xi = (68\text{ km}, 6\text{ MPa})$, $\xi = (60.5\text{ km}, 6\text{ MPa})$, $\xi = (30.5\text{ km}, 4\text{ MPa})$, and $\xi = (30.5\text{ km}, 6\text{ MPa})$. These cases generally lie at the periphery of our sampled parameter space, often near corners where the interpolation scheme has limited surrounding data. The ROM appears to exhibit a systematic overestimation of potency. While the R^2 score for ROM predicted P_0 versus FOM P_0 is 0.59, on the other hand the PCC between the two datasets is substantially higher at 0.946, suggesting a strong underlying correlation despite the offset. This systematic overestimation, further discussed in Section 5.5, though not ideal, might be addressable through post-processing corrections.

To further illustrate the ROM's predictive capabilities, Figures 9, S2, S3 show direct comparisons of the slip rate, cumulative slip, and state-variable evolution at two specific fault observation points: 195 km and 220 km along dip. These locations correspond to a VW region within W and the stability transition zone ($a-b=0$), respectively. These predictive capabilities are presented for three distinct parameter sets: $\xi_1 = (W = 37.375\text{ km}, \sigma_W = 5\text{ MPa})$, which yielded among the best prediction results in the LOOCV; $\xi_2 = (W = 60.5\text{ km}, \sigma_W = 4\text{ MPa})$, which showed median prediction accuracy; and $\xi_3 = (W = 43.0\text{ km}, \sigma_W = 3\text{ MPa})$, which represents the cases with the largest LOOCV discrepancies.

To explore the performance of the ROM for ξ not in the LOOCV training set, Figure 10 shows a comparison of the slip rate over time as a function of along-dip distance (panels (a), (b)) and at three different positions along-dip (panels (c)-(e)) for the parameter $\xi_1 = (W = 55.5\text{ km},\ \sigma_W = 5.5\text{ MPa})$ which is part of the test set. The ROM predictions of the characteristic recurrence interval and potency are 1.56 ± 0.05 years and 5696 ± 798 m^2, respectively, both falling within one standard deviation of those of the FOM simulation, which has $T_c = 1.55 \pm 0.02$ years and $P_0 = 5207 \pm 1080$ m^2. The maximum slip rate is correctly reconstructed for an along-dip distance of 195 km, while $\log_{10}(\dot{s})$ of the FOM is 0.6 and 1.9 higher than that of the ROM for along-dip distances of 210 and 180 km, respectively. The ROM slip-onset time at 180 km along-dip closely approximates the FOM slip-onset time, although at 210 km along-dip the ROM slip-onset time is $\sim$20 seconds earlier than its FOM counterpart, both referenced to the rupture onset at 195 km along-dip. We note that we measured training accuracy based solely on recurrence time and potency, so local slip-onset times are only indirectly constrained by the ROM. This could be reduced by refining the sampling in parameter space or including additional metrics (e.g., onset time or peak slip rate) in the ROM training.

For both ξ_1 and ξ_2, the ROM mostly captures the overall SSEs timeseries characteristics successfully, including peak slip rates and recurrence times. For ξ_3, the recurrence time prediction is less accurate; although the mean recurrence time predicted by the ROM remains within one standard deviation of the FOM's mean recurrence time, the difference between the means is approximately two months, which is comparable to the cycle-to-cycle standard deviation observed in the FOM for this parameter set. Nevertheless, the prediction of event potency for ξ_3 remains reasonable, with the FOM yielding an average potency of 3153 m^2 and the ROM predicting 3097 m^2.

In machine learning applications, it is common practice to employ a validation set for hyperparameter

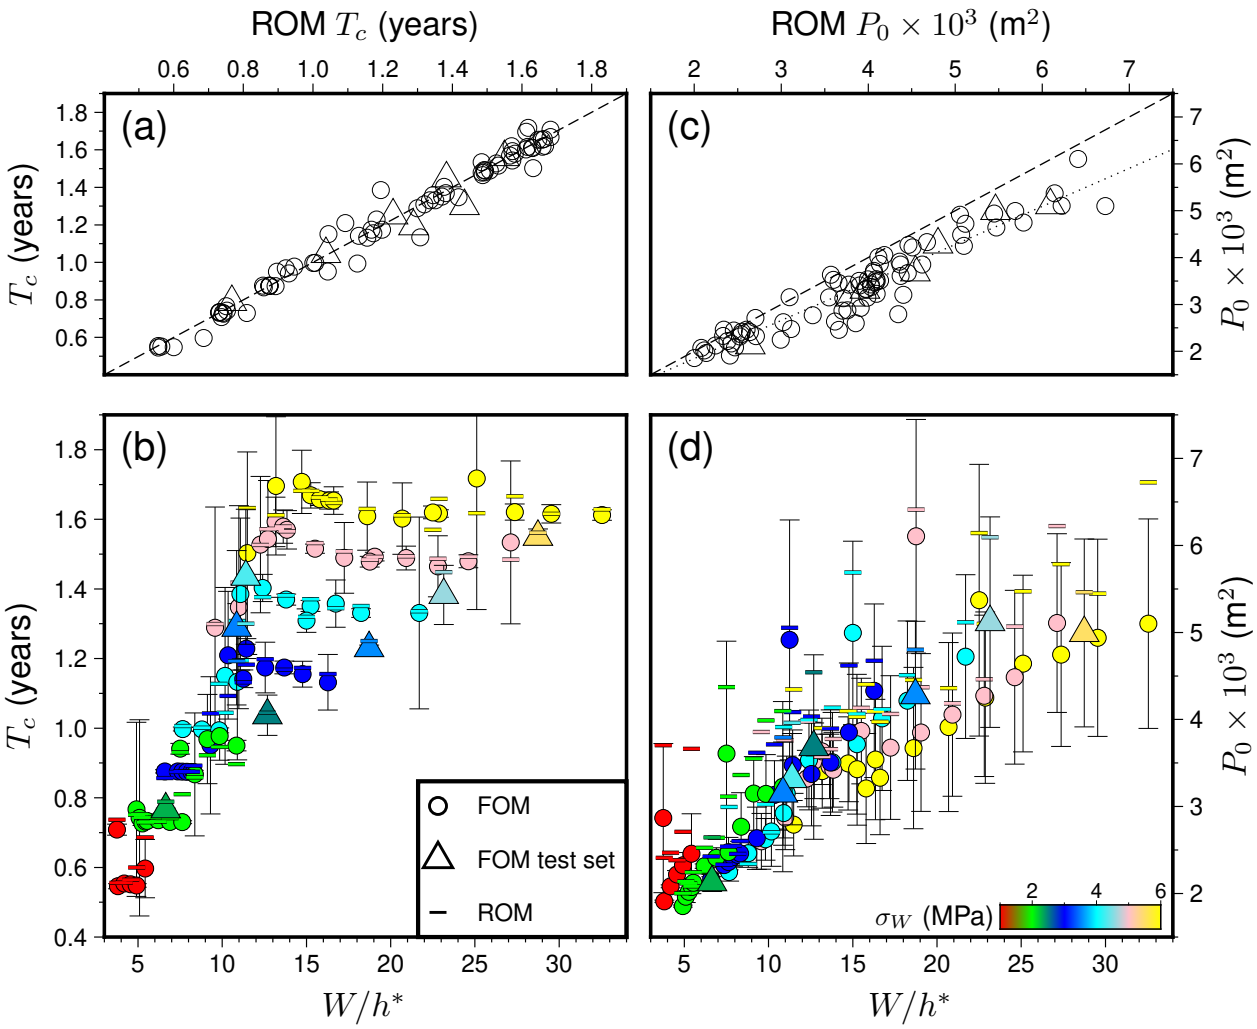

Figure 8: Results of the LOOCV test assessing ROM performance. (a) Comparison of FOM results versus ROM-predicted recurrence time (T_c). The dashed line indicates a 1:1 ratio. (b) Characteristic recurrence time (T_c) as a function of normalized fault width (W/h^*). Vertical bars represent the standard deviation derived of the FOM results, and colored horizontal bars show the corresponding ROM-calculated values. (c) Comparison of FOM results versus ROM-predicted potency (P_0). The dashed line indicates a 1:1 ratio and the dotted line shows the best linear fit between the ROM and FOM results. (d) Potency (P_0) as a function of normalized fault width (W/h^*). Vertical bars represent the standard deviation of the FOM results, and colored horizontal bars indicate the ROM-calculated values.

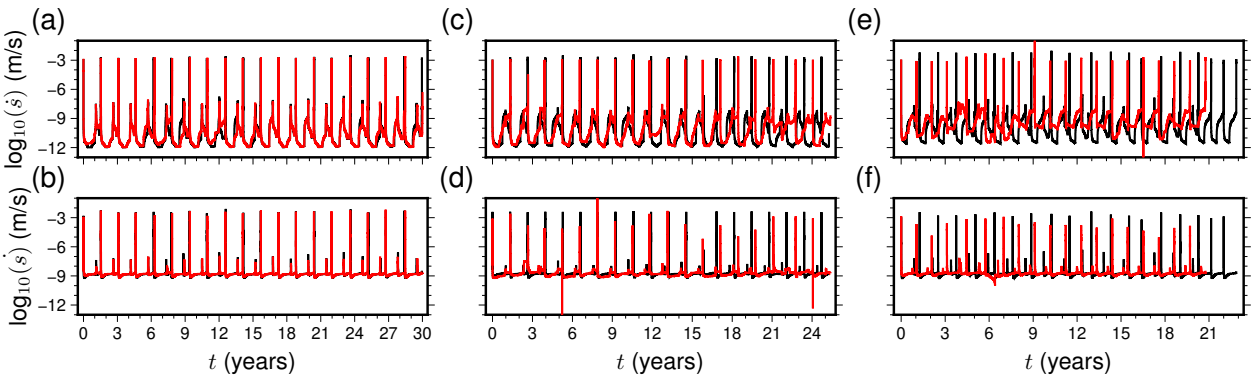

Figure 9: LOOCV comparison of slip rate predictions from the ROM (red curves) versus those from the FOM (black curves). Results are displayed for three different parameter sets at two along-dip observation points: 195 km (panels (a), (c), (e)) and 220 km (panels (b), (d), (f)). The parameters (W, σ_W) used were: $\xi_1 = (37.375 \text{ km}, 5 \text{ MPa})$ ((a)-(b)); $\xi_2 = (60.5 \text{ km}, 4 \text{ MPa})$ ((c)-(d)); $\xi_3 = (43.0 \text{ km}, 3 \text{ MPa})$ ((e)-(f)).

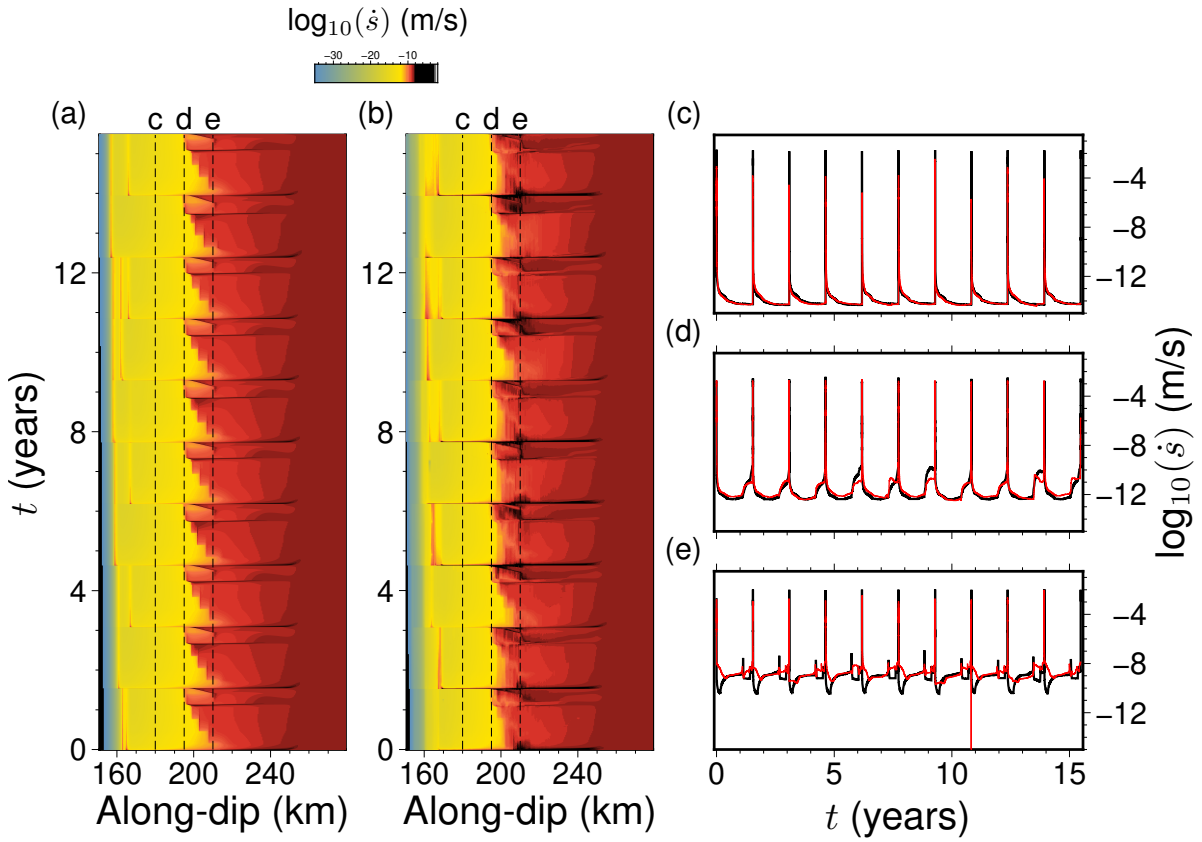

Figure 10: Comparison of FOM and ROM simulations for $\xi = (W = 55.5 \text{ km}, \sigma_W = 5.5 \text{ MPa})$ from the test set, which was not used to construct the ROM. Each simulation presents 10 SSE cycles. (a) FOM slip rate as a function of time and along-dip distance. The locked portion of the subduction interface is not shown. (b) ROM slip rate in the same representation. (c)-(e) Time histories of slip rate at specific observation points along the fault, marked by dashed lines in panels (a) and (b), with FOM results shown in black and ROM results in red.

tuning and model construction, and a separate test set for an unbiased evaluation of the finalized model (Kohavi, 1995). Although we utilized an LOOCV scheme for validation (Section 3.2), we also established an independent test dataset to further assess our ROM. For this purpose, an additional seven FOM simulations were run, distributed across our parametric domain. These simulations were not used during the ROM construction process.

The test dataset yielded prediction results for T_c and P_0 that were similar to those obtained from the LOOCV procedure (Figure 8), achieving R^2 scores of 0.92 and 0.58, respectively. The test dataset results for P_0 also showed the systematic overestimation previously observed in the LOOCV analysis, although the predicted potency maintained a strong linear correlation with the simulated potency, evidenced by a PCC of 0.98. The availability of this test dataset also facilitated the evaluation of a post-processing correction for the potency overestimation. The corrected potency, P_0^c, is calculated as:

$$P_0^c = (m-1)P_0 - n, \tag{32}$$

where m and n are the coefficients of the first-degree polynomial that best fits the potency predictions from the LOOCV analysis. Applying this correction to the potency values in the test set improved the R^2 score to 0.95 (Figure S6).

The computational efficiency and demonstrated accuracy of the ROM facilitate a detailed exploration of the parameter space. Figure 11 presents the ROM predictions for T_c and the corrected P_0 on a finely sampled grid of W and σ_W. We initially sampled the parametric space uniformly with 10000 P^* pairs, subsequently removing pairs that fell outside the convex hull of the FOM training parameters (illustrated as white space in the bottom left corner of Figure 11), which resulted in a total of 9116 ROM evaluations. The ROM smoothly interpolates between the FOM training points, revealing complex dependencies that might be overlooked with coarser parameter space sampling. For instance, analyzing the FOM results solely as a function of W/h^* might suggest that for normalized fault lengths in the range of $4-15$, SSE recurrence times increase linearly, and then gradually approach a constant for $W/h^* > 15$. In contrast, the ROM's dense parametric view of T_c as a function of both W and σ_W (Figure 11(b)) shows a more complex picture. A diagonal band, roughly delineated by the points $(W = 53\text{ km}, \sigma_W = 1\text{ MPa})$ and $(W = 30.5\text{ km}, \sigma_W = 6\text{ MPa})$, exhibits a high gradient, indicating rapid changes in recurrence time over short parametric distances. To the upper-right and lower-left of this band, smaller gradients are observed, with T_c variations appearing to be predominantly influenced by changes in effective normal stress. This detailed mapping of SSE characteristics as a function of fault properties is critical for understanding the underlying physics and for constraining these parameters against geodetic observations. The construction of these high-resolution parameter maps using the ROM required only 75 CPU hours, in stark contrast to the estimated 3×10^7 CPU hours that would have been necessary if one exclusively used a FOM.

4.4 Uncertainty quantification of width and amplitude of low effective normal stress regions governing slow slip events

The extensive parameter space exploration detailed in Section 4.3, although showcasing the ROM's efficiency, could theoretically be achieved with FOMs, if sufficient parallel computing resources were available, although at a vastly greater cost. However, certain tasks in model-based inference, such as global optimization or Bayesian parameter estimation via Markov Chain Monte Carlo (MCMC) methods, are inherently sequential or have limited parallelizability. For such methodologies, the computational cost of FOMs renders them practically infeasible. The rapid evaluation capabilities of our ROM framework, on the other hand, opens the door to utilize these techniques, allowing for robust uncertainty quantification of model parameters based on observations.

Here, we describe how the ROM can be employed within an MCMC framework to invert for the uncertainties in the fault parameters W (width of the low effective normal stress zone) and σ_W (magnitude of low

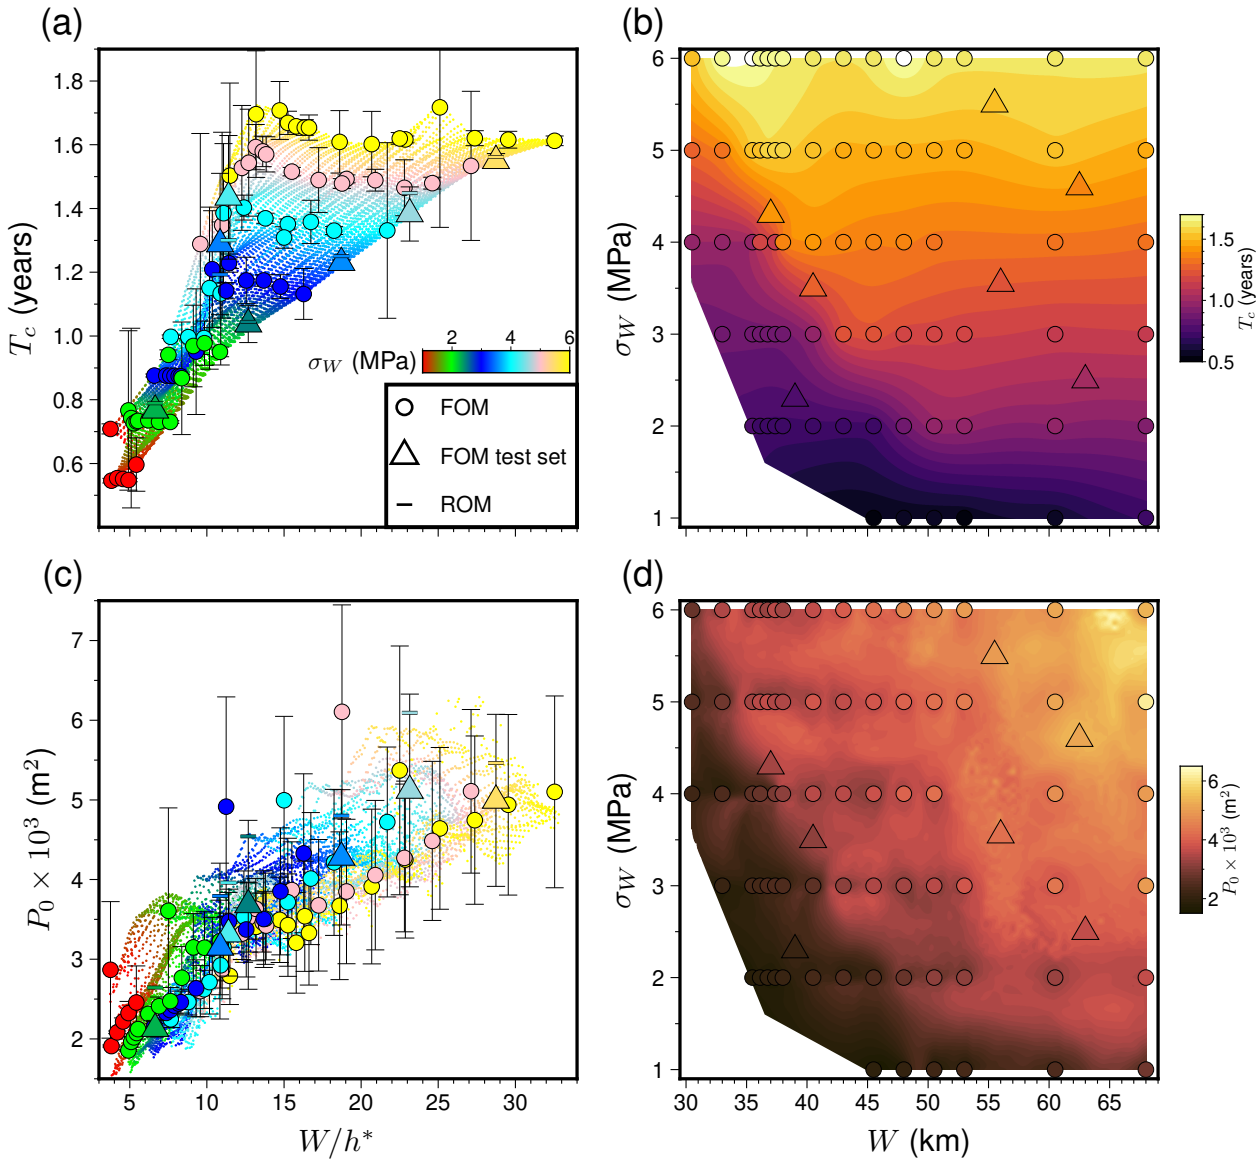

Figure 11: ROM predictions for characteristic properties of SSEs. (a) Recurrence time (T_c) as a function of the normalized fault width (W/h^*). (b) Recurrence time (T_c) as a function of the width (W, x-axis) and magnitude (σ_W, y-axis) of the low effective normal stress zone. This panel highlights complex dependencies, such as a diagonal band of high T_c gradient (c) Corrected potency (P_0) as a function of the normalized fault width (W/h^*). (d) Corrected potency (P_0) as a function of W (x-axis) and σ_W (y-axis). The uncorrected potency results are presented in Figure S7.

effective normal stress), constrained by observed characteristics of Cascadia SSEs. We will show that in our models the recurrence interval of SSEs is mainly controlled by the magnitude of effective normal stress, whereas their magnitudes primarily depend on the width of the low-stress region.

We employ the Metropolis-Hastings MCMC algorithm to sample the posterior probability distributions for W and σ_W. The observational constraints derived from Cascadia SSE studies are:

1. The mean recurrence interval $\mu_{T_c}^{obs} = 1.17$ years (14 months) with a standard deviation $\sigma_{T_c}^{obs} = 0.17$ years (2 months), assuming a normal distribution (Gomberg et al., 2016; Schmidt and Gao, 2010).

2. The seismic moment M_0 is assumed to follow a normal distribution with mean $\mu_{M_0}^{obs} = 7.28 \times 10^{18}$ Nm (corresponding to $M_w \approx 6.45$) and standard deviation $\sigma_{M_0}^{obs} = 4.22 \times 10^{18}$ Nm, encompassing the typical range of $M_w \sim 6.2 - 6.7$ for Cascadia SSEs (Behr and Bürgmann, 2021; Schmidt and Gao, 2010).

For given parameter vector $\xi = (W,\ \sigma_W)$, the ROM can be used to compute T_c^{ROM} and P_0^{ROM} (in units of m^2). The ROM potency is corrected by the best fitted linear relation to 1 to 1 relation correction (Figure 8) and then converted to seismic moment using $M_0^{ROM} = \mu W_s P_0^{ROM}$, where μ is the shear modulus (Table 1) and W_s is the assumed along-strike width of 60 km.

The likelihood function $\mathscr{L}(D^{obs} \,|\, (W, \sigma_W))$ for a proposed parameter vector, given the observed data $D^{obs} = (T_c^{obs},\ \log_{10}(M_0^{obs}))$, is:

$$\mathscr{L} \propto \exp\left(-\frac{1}{2}\left[\left(\frac{T_c^{ROM} - \mu_{T_c}^{obs}}{\sigma_{T_c}^{obs}}\right)^2 + \left(\frac{\log_{10}(M_0^{ROM}) - \mu_{\log_{10}(M_0)}^{obs}}{\sigma_{\log_{10}(M_0)}^{obs}}\right)^2\right]\right). \tag{33}$$

The formulation of the likelihood in Eq. (33) assumes that the observational uncertainties in the recurrence interval and the seismic moment are independent, as represented by a diagonal covariance matrix with zero off-diagonal terms. This is a simplifying assumption, since SSE characteristics can be correlated, as suggested by both observations Michel et al. (2019) and our model results (e.g., Figure 12). Neglecting this covariance may introduce a modest bias in the posterior distributions, potentially resulting in slightly broader or shifted parameter estimates. However, we consider this approach a reasonable and computationally tractable approximation for constraining the primary model parameters.

We initiate the chains with a uniform prior distributions to W and σ_W over the ranges $W \in [30.5, 68]$ km and $\sigma_W \in [1, 6]$ MPa. Then the chains are propagated with the Metropolis-Hastings algorithm (Hastings, 1970) which iteratively proposes new parameter vector $(W_{i+1}, \sigma_{W,i+1})$ from the current state $(W_i, \sigma_{W,i})$. The ROM is evaluated at the proposed state, the likelihood $\mathscr{L}_{i+1}$ is computed, and the proposal is accepted with probability $\alpha = \min\left(1, \frac{\mathscr{L}_{i+1} \cdot \text{Prior}_{i+1}}{\mathscr{L}_i \cdot \text{Prior}_i}\right)$.

To sample the posterior distributions, we ran 10 parallel chains, each for 6000 iterations (including a 600 burn-in iteration) requiring approximately 50 hours of computation. This extent of MCMC analysis is computationally tractable only due to the ROM's efficiency. Relying on the FOM would be prohibitive; the sequential nature of each chain, requiring 6000 model evaluations, would translate to an estimated 50 years of computation time per chain, in addition to the general immense total computational effort required for all 10×6000 FOM evaluations.

The convergence of the MCMC inversion was monitored using diagnostics presented in Figure S8. We employed the potential scale reduction factor, $\hat{R}$, which compares variance between chains to variance within each chain, values approaching 1 indicate convergence to a common target distribution (Vehtari et al., 2021). We also calculated the Effective Sample Size (ESS) to quantify the number of independent samples in the correlated MCMC output, crucial for reliable posterior inference (Gelman et al., 1995). Figure S8 indicates that $\hat{R}$ values for model parameters generally fell below 1.01 after approximately 4000 total MCMC evaluations (summed across all 10 chains, not including their burn-in periods). The ESS for

key parameters typically surpassed 200, a level often considered sufficient for robust estimation of posterior means and standard deviations (Gelman et al., 1995), after approximately 8000 total MCMC evaluations. Attaining these levels of convergence and sample independence, corresponding to thousands of individual model evaluations, which is not feasible with FOM simulations.

The MCMC analysis, informed by the observed Cascadia SSE recurrence intervals and seismic moments, yields posterior probability distributions for the fault parameters W and σ_W, visualized in Figure 12. The posterior distribution for the width of the low effective normal stress zone, W, is characterized by a mean of 44.7 km and a standard deviation of 16.2 km. For the magnitude of the low effective normal stress, σ_W, the inferred posterior has a mean of 3.8 MPa and a standard deviation of 1.4 MPa. These values represent the constrained estimates and associated uncertainties for these parameters, conditional on the observational data and the physics assumed by our model. A key physical assumption in this framework is the choice of the along-strike width $W_s = 60$ km, which is used to convert the 2D model potency to a 3D seismic moment. This choice, required to compare our 2D model with 3D observations, directly scales the inverted parameters. The conversion from 2D to 3D is inherently non-unique. Although we assume a simple rectangular patch, other approaches such as elliptical shapes, have been proposed (Liu and Rice, 2009). We find $W_s = 60$ km to be a reasonable value that produces SSE characteristics consistent with observations.

Analysis of the relationships within the posterior samples (Figure 12) further illuminates the control of these parameters on SSE characteristics. The SSE recurrence interval (T_c) exhibits a strong dependence on σ_W, with a PCC between their posterior samples of 0.86, whereas its correlation with W is considerably weaker (PCC = 0.2). Conversely, the seismic moment (M_0) is primarily correlated with W (PCC = 0.75), and shows a more moderate correlation with σ_W (PCC = 0.45). These findings suggest that, within our model framework for Cascadia-like SSEs, the magnitude of effective normal stress predominantly governs the timing of SSEs, although the spatial extent of this low-stress region is the primary factor controlling their magnitude.

5 Discussion

5.1 Implications for the state of stress and pore fluid pressure in Cascadia

Our findings support the critical role of low effective normal stress, e.g., maintained by high pore fluid pressure (e.g. Behr and Bürgmann, 2021), in enabling and modulating SSEs at the CSZ. Our MCMC inversion, constrained by observed Cascadia SSE characteristics, suggests that these events are consistent with low effective normal stress of 3.8 ± 1.4 MPa. This aligns well with conditions hypothesized for SSE generation, which necessitate near-lithostatic pore fluid pressures at the depths where Cascadia SSEs typically occur. Such low effective normal stress is consistent with the findings of Audet and Kim (2016), who highlighted near-lithostatic pore-fluid pressure as a dominant control in SSE environments, often evidenced by seismic observations like high V_p/V_s ratios in Cascadia's episodic tremor and slow slip zone. Based on 3D dynamic rupture simulations Madden et al. (2022) proposed that pore fluid pressure likely averages near 97% of lithostatic pressure in the Sumatra megathrust. Our results give even higher pore fluid pressure of $99.6 \pm 0.17\%$ of lithostatic pressure assuming constant crust density of 2.8×10^3 kg/m^3. We are also in quantitative agreement with numerical models by Perez-Silva et al. (2023), which require effective normal stresses in the range of 1-5 MPa to generate SSEs on rate-strengthening faults, a scenario compatible with the transitional stability regime investigated in our study.

The inferred upper depth limit of the SSE generation zone from our MCMC inversion, 30.44 ± 2.8 km, or 175.3 ± 16.2 km along-dip provides quantitative constraints on the transition from locked seismogenic behavior to aseismic creep along the Cascadia megathrust. This depth is consistent with observations and models suggesting that SSEs in Cascadia initiate down-dip of the primary locked seismogenic zone. For instance, Audet and Kim (2016) noted that deep non-volcanic tremors, which are often correlated with

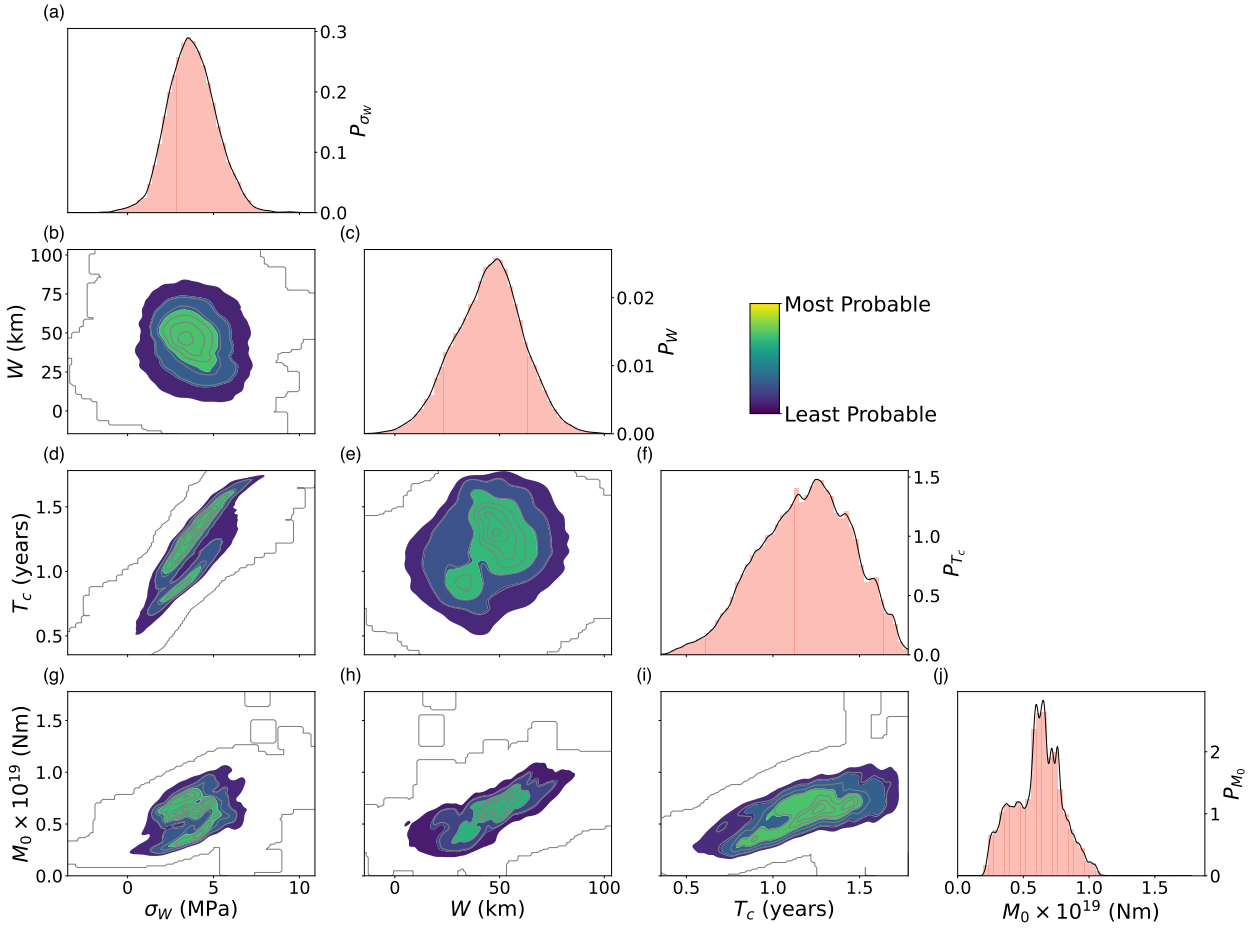

Figure 12: Posterior probability distributions and correlations for fault parameters characteristic properties of SSEs. The parameters shown are: the magnitude of low effective normal stress (σ_W), the width of the low effective normal stress zone (W), the recurrence interval (T_c), and the seismic moment (M_0). Diagonal plots (panels (a), (c), (f), (j)) display the marginal posterior probability distributions for each parameter. The y-axis of each marginal plot indicates the probability density. Off-diagonal plots (panels (b), (d), (e), (g), (h), (i)) illustrate the bivariate joint posterior distributions for pairs of parameters. For example, panel (b) shows the joint posterior of σ_W and W.

SSEs, generally occur at depths of 30 to 45 km , some distance down-dip of the main seismogenic zone. Michel et al. (2019), who inverted geodetic observations using secular linear motion for interseismic locking and for SSE slip distribution, also describe the zone of SSEs and tremors in their Cascadia models as lying inland from the coastline, clearly disconnected from and down-dip of the locked portion of the megathrust by a shallow creeping section. Their modelling indicates this transition zone, characterized by stationary fault creep, spans between approximately 100 km and 150 km away from the trench. The down-dip limit of this shallow creeping section at around 150 km along dip, as identified by Michel et al. (2019), falls within the uncertainty bounds of our MCMC inversion result for the updip limit of the SSE zone, although our specific model setup does not explicitly impose a creeping zone between the locked and SSE-prone sections.

Although we can validate the MCMC-inverted width of the SSE zone (W) against non-volcanic tremor locations and geodetic slip-inversion data, we cannot observationally validate the magnitude of the low effective normal stress in the SSEs zone. The derived value of 3.8 ± 1.4 should be interpreted with caution, because we keep other model parameters, such as a, b, D_c, and the slab geometry, fixed throughout our simulation stack. However, these parameters are not necessarily well constrained. For example, both the $a-b$ profile derived from gabbro friction experiments (He et al., 2006) and the extent to which lab-derived parameters hold in real faults have uncertainties (Marone, 1998).

To better understand the sensitivity of our inverted effective normal stress to other modelling parameters, we perform an additional set of 17 FOM simulations (11 new and 6 from the W vs σ_W LOOCV set), fixing $W = 44.5$ km and varying σ_W and the $a-b$ profile (Figure 13a). We explore the σ_W vs $a-b$ parameter space using the same reduced-order modelling scheme, with the results shown in Figure 13. For smaller values of $\overline{a-b}$ (the mean $a-b$ within the low effective normal stress region), only a narrow range of effective normal stress values ($4.22 - 4.68$ MPa) produces SSEs characteristics consistent with those observed in northern Cascadia. For $\overline{a-b}$ larger than the value derived from gabbro friction experiments He et al. (2006), a wide range of increasing effective normal stress can reproduce the northern Cascadia SSEs characteristics. This trend of increasing $\overline{a-b}$, which broadens the range of allowable σ_W, may have a theoretical limit that we do not explore in this study. Increasing $\overline{a-b}$ results in a shallower transition from the VW to the VS portions of the slab, meaning that a smaller portion of the SSEs zone is VW. With a sufficiently small VW portion of the SSE zone, the simulation will produce steady creep rather than SSE cycles. In the next section, we discuss the effect of including more of the VW portion of the fault in the SSE zone and its importance to the SSE's characteristics.

5.2 Secondary controls on the recurrence time of SSEs

Our parameter exploration, enabled by the ROM, allows us to investigate parameter regimes beyond those typically accessible via computationally intensive direct FOM studies, expanding upon the linear trends reported by Liu and Rice (2009) and highlighting more complex, second-order dependencies on both W and σ_W (Figs 5, 11, 12). Our MCMC inversion (Figure 12) uncovers the posterior distributions of W and σ_W, constrained by characteristic SSE observations from the northern CSZ.

Given that the extent of our parameter exploration allows W/h^* to vary by a factor of 6 due to changes in σ_W (holding W constant), versus a factor of 1.6 due to changes in W (holding σ_W constant), the general gradient of T_c across the sampled parameter space appears predominantly aligned with the σ_W axis. This observation is supported by the MCMC inversion, where T_c exhibits a PCC of 0.83 with σ_W, compared to only 0.2 with W. Notably, some regions exhibit high gradients in recurrence time with a strong dependence on W. For instance, at $\sigma_W = 2.75$ MPa, T_c jumps from 0.85 years to 1.1 years as W increases from 38 km to 43 km. This represents a 30% increase in recurrence time but only a 7% increase in the normalized fault width W/h^*.

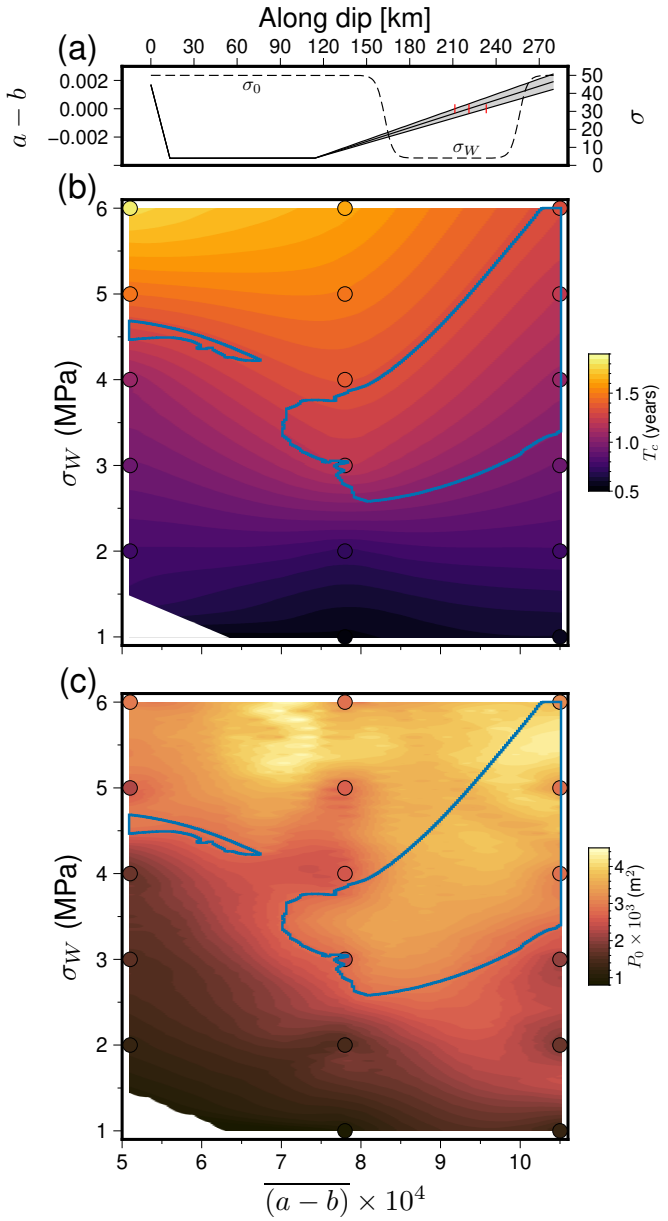

Figure 13: ROM predictions for characteristic properties of SSEs with W fixed at 45.5 km and varying the along-dip $a-b$ profile. (a) Along-fault profiles of $a-b$ derived from gabbro friction experiments (He et al., 2006), where we vary the slope of the $a-b$ curve for along-dip distances larger than 114 km $\pm 10\%$. With $a-b=kd+c$ for $d>114$ km, the distance along-dip, we vary the slope to be $0.9k$, $1.0k$, and $1.1k$, where k is the slope derived for the gabbro friction experiments. Black lines show FOM results, and gray areas indicate ROM evaluations. The red horizontal bar marks the VW-VS transition. The dashed line indicates the initial effective normal stress distribution along the fault (right y-axis), illustrating the fixed $W=45.5$ km width of the low effective normal stress zone. We show the profiles of effective normal stress for $\sigma_W=4$ MPa. (b) Recurrence time T_c as a function of $\overline{(a-b)}10^4$ (x-axis), the mean $a-b$ in the low effective normal stress zone W, and the magnitude σ_W (y-axis) of the low effective normal stress zone. (c) Corrected potency P_0 as a function of $\overline{(a-b)}10^4$ (x-axis) and σ_W (y-axis). Blue polygons in (b) and (c) delineate the parameter space producing SSEs with northern Cascadia-like characteristics.

5.3 Forward modelling method

A crucial aspect underpinning our ROM development is the verified accuracy of the FOMs that generated the training data. Our FOM simulations, performed using the volumetric discontinuous Galerkin finite element code tandem (Section 2, Uphoff et al., 2022), successfully reproduce the key relationships between normalized fault width (W/h^*), SSE recurrence interval, and mean SSE slip previously established by Liu and Rice (2009) using a boundary element method (BEM) (Figure 5). This agreement verifies our physical model setup prior to the development of our ROM.

While BEM offers computational advantages for fault-dominated problems by reducing dimensionality compared to our tandem-based FOMs, for more complex model setups BEM simulations may still be computationally expensive. For example, Tainpakdipat et al. (2025) report that simulations for their SSEs models demand between 20 and 250 CPU hours. Therefore, BEM may be similarly limited as volumetric codes when calculating more than 10,000 model inferences, as done with our ROM for exploring the parameter spaces and estimating uncertainties. Furthermore, classical BEM techniques typically assume a homogeneous or layered elastic medium and simplified model and fault geometries (Barbot, 2019; Lapusta and Liu, 2009; Lapusta et al., 2000; Li and Liu, 2016a, 2017; Liu and Rice, 2005; Rice and Gu, 1983; Segall and Bradley, 2012). Despite recent developments to incorporate more complex models (Mallick and Sathiakumar, 2024; Mallick et al., 2022, e.g.,), efficiently handling widespread and complex off-fault material variations, or complex subsurface geometries and topography, remains a challenge for BEM compared to volumetric codes.

It is important to note that the ROM framework presented in this study is agnostic to the specific numerical method of the FOM. Consequently, our ROM scheme could be readily applied to training data generated from BEM-based simulations or from volumetric simulations that accommodate complex off-fault material properties.

Nie and Barbot (2021) explore SSEs in a 2D anti-plane strain setting, varying both the Dieterich-Ruina-Rice number (R_u), which is proportional to W/h^* by a constant factor close to unity and $R_b = (b-a)/b$, a parameter that controls the ratio of dynamic to static stress drops (Gabriel et al., 2012). In our model, a larger width of the low effective normal stress zone incorporates a greater portion of the VW section of the fault (Figure 1), thereby increasing the apparent R_b of the SSE-producing zone. Although Nie and Barbot (2021) did not directly investigate the change in recurrence time as a function of R_b, their observed data shows rapid changes in the peak slip rate and the system limit cycle style which can consequently alter the recurrence times of events. The picture emerging from our parametric space exploration reveals three distinct dependencies of recurrence time on the normalized fault length: for $R_b < 0.175$, recurrence time increases steeply over a relatively short interval of W/h^* (from 7.5 to 12.5); in contrast, for $R_b > 0.21$, the recurrence interval becomes linearly dependent on the normalized fault length, and a transition zone between these behaviors is identified for R_b values in the range $0.175 \leq R_b \leq 0.21$ (Figure S9). These detailed dependencies of recurrence time on W/h^* agrees with Nie and Barbot (2021) in that SSEs characteristics have at least secondary dependence on R_b.

5.4 Comparison with previous work

The application of ROM techniques to accelerate computationally intensive simulations is gaining traction in earthquake science (Hobson and May, 2025a; Kaveh et al., 2024; Ragu Ramalingam et al., 2025; Rekoske et al., 2023, 2025) and beyond (e.g., Degen et al., 2023; Hobson and May, 2025b; Quiaro et al., 2025). Our ROM methodology, which combines a spline-based latent space representation with RBF interpolation for POD coefficients builds a ROM for each individual slow slip event expanding the approach by Rekoske et al. (2025). A key difference in our approach is the two-step nature of the model order reduction. The first step involves the transformation of the complex SSE cycle data into an efficient, low-dimensional

spline-based latent representation. Key advantage of this latent space is the representation of simulations as fixed-length vectors, which addresses the challenge of variable timestep outputs and differing durations across simulations. This is a prerequisite for the following matrix-based ROM analyses and could also be utilized by other machine learning methods, such as many neural networks, which require uniform-length inputs (Lecun et al., 1998). Another aspect of our ROM framework the per-cycle ROM construction helps effectively managing and reducing the dimensionality of the complex time-history data that are characteristic of SSE cycle simulations.

Kaveh et al. (2024) employ a POD-based ROM to forecast extreme events in a rate-and-state friction fault model that produces SSEs. They focus on identifying precursor states to SSEs by building their ROM from simulation snapshots of the inter-event periods only, and over one simulation settings (no change in initial parametrization), thereby capturing the system's characteristics while excluding the SSEs themselves. This allows them to define an optimization problem within the reduced-order space to find extreme events precursors. In contrast, our ROM is designed to efficiently simulate the entire SSE cycle, including the SSEs.

Physics-Informed Neural Networks (PINNs, Fukushima et al., 2023; Okazaki et al., 2022) represent a class of deep learning models that are trained to solve PDEs by directly incorporating the equations, along with initial and boundary conditions, into the neural network's loss function. Recently, PINNs have emerged as a promising candidate for a data-driven approach to solving and inverting fault and rate-and-state friction equations in both laboratory (Borate et al., 2024) and numerical settings (Fukushima et al., 2025; Rucker and Erickson, 2024). Fukushima et al. (2025) employed PINNs for the direct inversion of spatially distributed frictional parameters $(a, a-b, L)$ from geodetic observations. Parameters are determined by minimizing a composite loss that includes both data misfit and PDEs residuals, the underlying physics thus serves as an inherent regularization constraint during the learning process itself. In such an intrusive framework, the learning phase is coupled with the inversion for specific fault frictional properties. Conversely, our ROM functions as a non-intrusive, data-driven surrogate for the underlying physics, i.e., is agnostic toward the kind of forward model used. The ROM approximates the complex input-output relationships of the system from a dataset of FOM simulations, without requiring explicit knowledge or direct utilization of the governing equations during its construction phase. This process yields an efficient forward model, which can subsequently be integrated into established inversion frameworks. As a result, our ROM approach is well-suited for uncertainty quantification, as showcased by our MCMC analysis. This is a capability not as directly featured in a PINNs parameter estimation framework.

5.5 Limitations

A primary consideration for the scalability of our ROM approach is the offline cost associated with generating the FOM simulations required for training. In this study, with a two-dimensional parameter space ($dim(\mathscr{P}) = 2$, for W and σ_W), the 76 FOM simulations, though computationally intensive, provide a sufficient basis for constructing an accurate ROM. However, the number of FOM evaluations needed to adequately sample the parameter space and train a robust ROM can, in theory, increase exponentially, with the number of parameters being surrogated in the parameter space $\mathscr{P}$. While the speedup achieved during the online phase is substantial, the initial investment in FOM simulations for higher dimensional parameter spaces could become a limiting factor.

A related consideration is the ROM sensitivity to the underlying spatial resolution of the FOM. Different model scenarios require different spatial resolution to resolve the characteristic nucleation size, h^*, and the process zone size, Λ (Rubin and Ampuero, 2005). Using a fixed spatial resolution over the entire parameter space that resolves the minimum values of h^* and Λ across all explored parameters is inefficient, with some models being over-resolved in space. An alternative approach is to use an adaptive spatial resolution which varies for each model scenario (i.e., each point in parameter space) to accurately resolve the length scale

associated with the particular parameter vector. We can view our ROM as a point-wise interpolation over the parameter space and, as such, the number of observation points m is not directly related to the underlying physical model resolution. This is illustrated by a ROM with one single observation point in the SSE zone yielding almost the same results at that location as the corresponding observation point in the full ROM with 346 observation points (Figure S4). Since the ROM-chosen m is independent of the FOM model resolution, the FOM grid may have varying spatial model resolutions while keeping the number of ROM observation points fixed. In addition, in cases where only a portion of the fault is of interest, as in our SSE simulations, one could place observation points only in this region of interest. Because the online inference computational cost is linearly dependent on the number of observation points m (Figure S5), the expected speed up in ROM inference time for the Cascadia SSE ROM presented in this study by limiting the obs points to the SSEs zone is $\sim$40%.

Another limitation observed in our results is the systematic overestimation of potency by the ROM (Figure 8(c),(d)). Although we demonstrated that a simple linear post-processing correction, derived from the LOOCV results, can effectively mitigate this bias for the test set (Figure S6), this is not an ideal solution. The source of this systematic discrepancy may lie in the RBFs interpolation of POD coefficients, in the accumulation of minor errors through the multi-step ROM construction process or potentially in the ROM lacking the ability to perfectly capture the transition between consecutive SSE cycles when they are concatenated in time. Ideally, the ROM should predict potency accurately without requiring such a correction. This suggests an area for future refinement, perhaps through the exploration of alternative interpolation schemes for the POD coefficients, or by investigating strategies like introducing small overlaps or more sophisticated blending techniques between individual SSE cycle ROMs when reconstructing longer time series.

Furthermore, as with most data-driven surrogate models, at its core, our ROM is an interpolatory method. Its accuracy can be trusted only within the convex hull of the training parameter sets (ξ) and it should not be used for extrapolation beyond the sampled parameter range. The quality of the ROM predictions is also contingent on the density and distribution of the FOM training samples. While our iterative refinement strategy aimed to address regions of high error, ensuring comprehensive coverage for complex, high-gradient parameter responses remains a challenge.

Finally, the physical limitations inherent in the FOMs themselves will propagate to the ROM. Our current FOMs, for instance, are 2D and do not capture 3D effects such as along-strike variations in fault properties (Brudzinski and Allen, 2007; Li and Liu, 2017) or in SSE kinematics (Li and Gabriel, 2024; Takagi et al., 2019). Similarly, more complex rheologies (Gao and Wang, 2017) or fluid-transport mechanisms (Ozawa et al., 2024; Perez-Silva et al., 2023), if not included in the FOMs, cannot be represented by the ROM. The ROMs are ultimately bounded by the underlying full-order model, the ROM cannot learn, or capture time-dependence and or physics which is not present in the FOM.

5.6 Future work

A natural future extension of our SEAS ROM approach involves applying it to more complex simulations, such as those incorporating fast earthquakes in addition to SSEs, more extensive frictional parameter variations, or to extend the forward simulations to 3D SSE models, all requiring to approximate more parameters. For example, a transition from 2D to 3D domains (and consequently from a 1D to a 2D fault) would increase the number of coupled rate-and-state friction equations. These added complexities could potentially disrupt the well-behaved nature of the phase-space limit cycle trajectories, possibly leading to chaotic behavior(Barbot, 2019; Cattania, 2019), which would be more challenging to capture with a ROM. Wang (2024) studied the effect on cycle simulations of a 1D fault embedded in a 2D domain and found that an increase in the number of interacting rate-and-state friction points in space can lead to more complicated time-dependence behavior, potentially resulting in more complex phase-space trajectories. Distinguishing between quasi-periodic and truly chaotic behavior is not straightforward, and the former can be misinter-

preted as the latter (Wang, 2024).

The stability of these limit cycles and their potential transition to chaotic behavior can be linked to the friction parameters explored. Viesca (2016b) demonstrated that as the ratio of rate-and-state friction parameters a/b approaches 1, the system tends towards instability and can exhibit chaotic characteristics. Conversely, smaller a/b values are associated with stable, periodic limit cycles. However, Viesca (2016b) also notes that in most seismic cycle models $a/b \leq 0.8$, in this regime, slip instability does not develop chaotically but rather in a universal manner. Barbot (2019) explored the effect of varying the R_b and R_u parameters in in-plane, anti-plane, and 3D SEAS models on slow-slip characteristics. They found that specific combinations of R_b and R_u values can lead to more complex behaviour and deterministic chaos. In 2D models, simulations can produce multi-event cycles that correspond to a limit cycle with multiple loops. In 3D models, some values of R_u can lead to deterministic chaos, which corresponds to an open limit cycle. These earlier results highlight that even modest changes in frictional properties can qualitatively alter SSE cycle behavior, and that our inverted parameter combinations for Cascadia are one realization within a broader, potentially more complex, parameter space.

In the context of our Cascadia-like model, the $a-b$ parameters vary along the fault, transitioning from VW to VS behavior. In the VW sections, our chosen b value of 0.0045 and a values (Figure 1(c)) result in a/b ratios that are generally sufficiently less than 1. This suggests that the SSEs generated in our 2D FOMs should, and do, exhibit stable limit cycles. However, extending the ROM to exploring parameter regimes where a/b is closer to unity or applying it to 3D model would necessitate careful consideration of these potential transitions to more complex, possibly chaotic or quasi-periodic, behaviors. The adaptability of the spline latent space representation and the POD-RBFs framework, which handles the simulation cycle-by-cycle, makes it a promising candidate for future work aiming at capturing such behavior, provided the training FOMs adequately sample these complex regimes of the parameter space.

To more adequately sample the complex regime of the FOM parameter space, future work could improve our sampling method, improve the ROM reconstruction metric, or both. For example, we might improve sampling by adopting more adaptive or sparse strategies to mitigate the challenge of FOM generation in higher-dimensional parameter spaces (Bui-Thanh et al., 2008) and an adaptive FOM sampling may increase efficiency.

The ROM reconstruction metric used in this study validates the characteristic SSE recurrence time interval and potency. We ensure, through iterative refinement, that the ROM characteristics are sufficiently close to the FOM characteristics of the LOOCV. While these scalar metrics are valid for our parametric exploration and MCMC inversion, we also perform a finer-grained comparison between the ROM and FOM simulations. A direct time-series comparison is impractical, as the FOM and ROM simulations use different temporal sampling vectors. However, our latent space (Section 3.1.1) provides a fixed-length representation of both the FOM and ROM simulations. We therefore use this latent space for a direct FOM and ROM simulation comparison (see Appendix C). This approach, analogous to functional data analysis schemes (Ramsay and Silverman, 2005), projects the time trajectories of the simulation slip rate, state variable and cumulative slip onto fixed-length latent vectors $q_{\dot{s}}$, q_{ψ}, and q_s and corresponding ROM counterparts $q^*_{\dot{s}}$, q^*_{ψ}, and q^*_s containing the spline coefficients of the phase-space trajectory, all with a fixed length. Figure 14 presents the cosine similarity and relative Euclidean distance between the FOM and ROM simulations over the LOOCV set. The ROM was not explicitly optimized for these metrics, yet the metrics show good mean reconstruction values are 0.18 ± 0.06, 0.2 ± 0.5 and 0.13 ± 0.05 relative Euclidean distance and $98 \pm 1\%$, $98 \pm 1\%$ and $99 \pm 0.7\%$ cosine similarity for the slip rate, state-variable, and cumulative slip, respectively. The state-variable and cumulative-slip worst performers in these metrics are on the periphery of the parameter space, where the interpolative power is limited. By contrast, the slip rate worst performers are located in the band of rapid recurrence-interval change discussed in Section 5.2. In future work, combining an adaptive parameter-space sampler with this more fine-grained reconstruction metric may improve the absolute reconstruction performance.

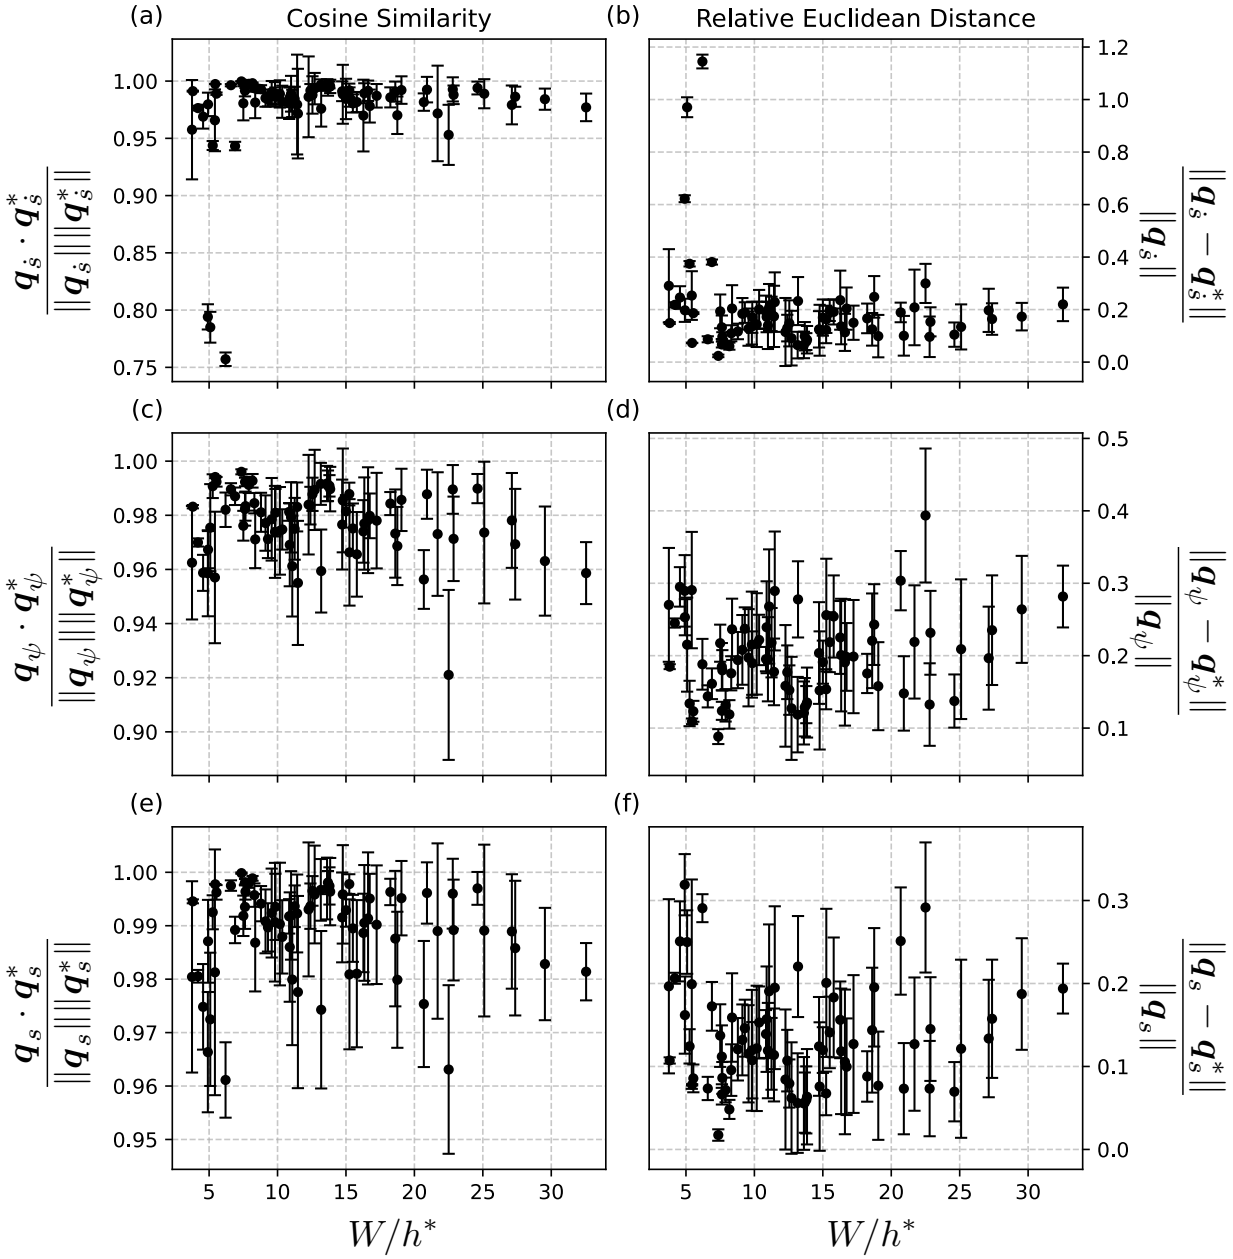

Figure 14: LOOCV preference of the ROM, measured via latent-space vector similarity (see Appendix C), as a function of normalized fault length (W/h^*). Rows correspond to slip rate ($\dot{s}$ (a), (b)), state variable (ψ (c), (d)), and cumulative slip (s (e), (f)). Columns correspond to the cosine similarity metric ((a), (c), (e)), and relative Euclidean distance ((b), (d), (f)).

6 Conclusions

We present a two-step, scientific machine learning reduced-order modelling (ROM) framework that accelerates rate-and-state friction simulations of the slow slip cycle by 3.6×10^5 compared with full-order sequences of earthquakes and aseismic slip (SEAS) models. First, each simulated slow slip event (SSE) is recast into a compact spline-based latent space using a phase-space representation of slip rate and state. Second, proper orthogonal decomposition (POD) combined with radial-basis-function (RBF) interpolation emulates the simulations with varying initial conditions. Our ROMs are validated with leave-one-out cross-validation and comparison to earlier, independent SSE simulations. We use the ROMs to explore complex, non-linear dependencies of northern Cascadia-like SSE characteristics on the width W and magnitude σ_W of a deep low effective normal stress zone. We perform a Bayesian Markov-chain Monte-Carlo inversion, constraining these parameters and their uncertainties to $W = 44.7 \pm 16.2$ km and $\sigma_W = 3.8 \pm 1.44$ MPa under the assumption of known frictional parameters. These values imply near-lithostatic pore fluid pressure ($99.6 \pm 0.17\%$ lithostatic) and place the upper SSE source boundary (i.e., the frictional transition zone) at 30.44 ± 2.8 km depth, which is consistent with geophysical observations. Because the method is non-intrusive and agnostic to the underlying forward model, future work may extend our ROMs to even higher-dimensional parameter spaces, mixed seismic–aseismic cycles, and fully 3-D geometries, offering a practical route to systematic uncertainty quantification throughout the earthquake cycle. By systematically linking megathrust properties to rate-and-state governed slow slip cycle characteristics, our study helps to constrain first- and second-order controls on how plate boundaries slip, providing input for seismic hazard assessment and future 3-D modelling.

7 Data Availability

All codes and data products supporting this study are openly available. The input files required to reproduce the tandem simulations, a static version of the tandem code, and the ROM analysis code are publicly hosted on Zenodo Magen et al. (2025). The complete raw outputs from all simulations are archived on the National Data Platform S3 bucket. The Zenodo repository provides the direct link and instructions for accessing these raw data.

Acknowledgments

We thank Ismael Perez and Alexander Tuna for helping setting up the computational framework on the Natulius National Research Platform cluster at the University of California, San Diego. We thank the Editor Marine Denolle, the Assistant Editor Fern Storey, and two anonymous reviewers for their constructive comments. This work was supported by Schmidt Science, LLC and the Statewide California Earthquake Center (SCEC, project 25313). DAM and A-AG acknowledge additional support by the National Science Foundation (NSF) through grant nos. EAR-2121568 and OAC-2311208. A-AG acknowledges additional support by the National Science Foundation (grant numbers OAC-2139536, EAR-2225286), by Horizon Europe (ChEESE-2P, grant number 101093038, DT-GEO, grant number 101058129, and Geo-INQUIRE, grant number 101058518), and the National Aeronautics and Space Administration (grant no. 80NSSC20K0495). The authors acknowledge computational support from the National Research Platform via NSF awards CNS-1730158, ACI-1540112, ACI-1541349, OAC-1826967, OAC-2112167, CNS-2100237, CNS-2120019. Additional computing resources were provided by the Institute of Geophysics of LMU Munich (Oeser et al., 2006) and by the Texas Advanced Computing Center (TACC) The University of Texas at Austin.

References

Shrirang Abhyankar, Jed Brown, Emil M. Constantinescu, Debojyoti Ghosh, Barry F. Smith, and Hong Zhang. PETSc/TS: A modern scalable ODE/DAE solver library. *arXiv e-preprints*, 2018.

Eiichiro Araki, Demian M. Saffer, Achim J. Kopf, Laura M. Wallace, Toshinori Kimura, Yuya Machida, Satoshi Ide, Earl Davis, IODP Expedition 365 shipboard scientists, Sean Toczko, Stephanie Carr, Chihiro Kinoshita, Reiji Kobayashi, and Alexander Rösner. Recurring and triggered slow-slip events near the trench at the Nankai Trough subduction megathrust. *Science*, 356(6343):1157–1160, 2017. doi: 10.1126/science.aan3120.

Pascal Audet and YoungHee Kim. Teleseismic constraints on the geological environment of deep episodic slow earthquakes in subduction zone forearcs: A review. *Tectonophysics*, 670:1–15, 2016. ISSN 0040-1951. doi: https://doi.org/10.1016/j.tecto.2016.01.005.

Pascal Audet, Michael G Bostock, Nikolas I Christensen, and Simon M Peacock. Seismic evidence for overpressured subducted oceanic crust and megathrust fault sealing. *Nature*, 457(7225):76–78, January 2009.

C. Audouze, F. De Vuyst, and P. B. Nair. Reduced-order modeling of parameterized PDEs using time–space-parameter principal component analysis. *International Journal for Numerical Methods in Engineering*, 80(8):1025–1057, 2009. doi: https://doi.org/10.1002/nme.2540.

Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. *Efficient Management of Parallelism in Object Oriented Numerical Software Libraries*. Birkhäuser Press, 1997.

Satish Balay, Shrirang Abhyankar, Mark F. Adams, Steven Benson, Jed Brown, Peter Brune, Kris Buschelman, Emil Constantinescu, Lisandro Dalcin, Alp Dener, Victor Eijkhout, Jacob Faibussowitsch, William D. Gropp, Václav Hapla, Tobin Isaac, Pierre Jolivet, Dmitry Karpeev, Dinesh Kaushik, Matthew G. Knepley, Fande Kong, Scott Kruger, Dave A. May, Lois Curfman McInnes, Richard Tran Mills, Lawrence Mitchell, Todd Munson, Jose E. Roman, Karl Rupp, Patrick Sanan, Jason Sarich, Barry F. Smith, Hansol Suh, Stefano Zampini, Hong Zhang, Hong Zhang, and Junchao Zhang. PETSc/TAO users manual. Technical Report ANL-21/39 - Revision 3.23, Argonne National Laboratory, 2025a.

Satish Balay, Shrirang Abhyankar, Mark F. Adams, Steven Benson, Jed Brown, Peter Brune, Kris Buschelman, Emil M. Constantinescu, Lisandro Dalcin, Alp Dener, Victor Eijkhout, Jacob Faibussowitsch, William D. Gropp, Václav Hapla, Tobin Isaac, Pierre Jolivet, Dmitry Karpeev, Dinesh Kaushik, Matthew G. Knepley, Fande Kong, Scott Kruger, Dave A. May, Lois Curfman McInnes, Richard Tran Mills, Lawrence Mitchell, Todd Munson, Jose E. Roman, Karl Rupp, Patrick Sanan, Jason Sarich, Barry F. Smith, Stefano Zampini, Hong Zhang, Hong Zhang, and Junchao Zhang. PETSc Web page. `https://petsc.org/`, 2025b.

Sylvain Barbot. Slow-slip, slow earthquakes, period-two cycles, full and partial ruptures, and deterministic chaos in a single asperity fault. *Tectonophysics*, 768:228171, 2019. ISSN 0040-1951. doi: https://doi.org/10.1016/j.tecto.2019.228171.

Sylvain Barbot, Nadia Lapusta, and Jean-Philippe Avouac. Under the hood of the earthquake machine: Toward predictive modeling of the seismic cycle. *Science*, 336(6082):707–710, May 2012.

Noel M. Bartlow, Shin'ichi Miyazaki, Andrew M. Bradley, and Paul Segall. Space-time correlation of slip and tremor during the 2009 Cascadia slow slip event. *Geophysical Research Letters*, 38(18), 2011. doi: https://doi.org/10.1029/2011GL048714.

Whitney M. Behr and Roland Bürgmann. What's down there? The structures, materials and environment of deep-seated slow slip and tremor. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 379(2193):20200218, 2021. doi: 10.1098/rsta.2020.0218.

Yehuda Ben-Zion and Lupei Zhu. Potency-magnitude scaling relations for southern California earthquakes with $1.0 < M_L < 7.0$. *Geophysical Journal International*, 148(3):F1–F5, 03 2002. ISSN 0956-540X. doi: 10.1046/j.1365-246X.2002.01637.x.

G Berkooz, P Holmes, and J L Lumley. The Proper Orthogonal Decomposition in the analysis of turbulent flows. *Annual Review of Fluid Mechanics*, 25:539–575, 1993. ISSN 1545-4479. doi: https://doi.org/10.1146/annurev.fl.25.010193.002543.

J. Biemiller, A.-A. Gabriel, D. A. May, and L. Staisch. Subduction zone geometry modulates the megathrust earthquake cycle: Magnitude, recurrence, and variability. *Journal of Geophysical Research: Solid Earth*, 129(8):e2024JB029191, 2024. doi: https://doi.org/10.1029/2024JB029191.

Prabhav Borate, Jacques Rivière, Samson Marty, Chris Marone, Daniel Kifer, and Parisa Shokouhi. Physics informed neural network can retrieve rate and state friction parameters from acoustic monitoring of laboratory stick-slip experiments. *Scientific Reports*, 14(1):24624, Oct 2024. ISSN 2045-2322. doi: 10.1038/s41598-024-75826-y.

Michael R. Brudzinski and Richard M. Allen. Segmentation in episodic tremor and slip all along Cascadia. *Geology*, 35(10):907–910, 10 2007. ISSN 0091-7613. doi: 10.1130/G23740A.1.

T Bui-Thanh, Murali Damodaran, and Karen Willcox. Proper orthogonal decomposition extensions for parametric applications in compressible aerodynamics. In *21st AIAA Applied Aerodynamics Conference*, Reston, Virigina, June 2003. American Institute of Aeronautics and Astronautics.

T. Bui-Thanh, K. Willcox, and O. Ghattas. Model reduction for large-scale systems with high-dimensional parametric input space. *SIAM Journal on Scientific Computing*, 30(6):3270–3288, 2008. doi: 10.1137/070694855.

Roland Bürgmann. The geophysics, geology and mechanics of slow fault slip. *Earth and Planetary Science Letters*, 495:112–134, 2018. ISSN 0012-821X. doi: https://doi.org/10.1016/j.epsl.2018.04.062.

C. Cattania. Complex earthquake sequences on simple faults. *Geophysical Research Letters*, 46(17-18):10384–10393, 2019. doi: https://doi.org/10.1029/2019GL083628.

Federico Ciardo and Robert C Viesca. Non-linear stability analysis of slip in a single-degree-of-freedom elastic system with frictional evolution laws spanning aging to slip. *arXiv preprint arXiv:2407.16846*, 2024.

Giuseppe Costantino, Sophie Giffard-Roisin, Mathilde Radiguet, Mauro Dalla Mura, David Marsan, and Anne Socquet. Multi-station deep learning on geodetic time series detects slow slip events in Cascadia. *Communications Earth & Environment*, 4(1):435, Nov 2023. doi: 10.1038/s43247-023-01107-7.

Denise Degen, Daniel Caviedes Voullième, Susanne Buiter, Harrie-Jan Hendricks Franssen, Harry Vereecken, Ana González-Nicolás, and Florian Wellmann. Perspectives of physics-based machine learning strategies for geoscientific applications governed by partial differential equations. *Geoscientific Model Development*, 16(24):7375–7409, 2023.

James H. Dieterich. Modeling of rock friction: 1. Experimental results and constitutive equations. *Journal of Geophysical Research: Solid Earth*, 84(B5):2161–2168, 1979. doi: https://doi.org/10.1029/JB084iB05p02161.

John R Dormand and Peter J Prince. A family of embedded Runge-Kutta formulae. *Journal of computational and applied mathematics*, 6(1):19–26, 1980.

Herb Dragert, Kelin Wang, and Thomas S. James. A silent slip event on the deeper Cascadia subduction interface. *Science*, 292(5521):1525–1528, 2001. doi: 10.1126/science.1060152.

Philippe Druault, Philippe Guibert, and Franck Alizon. Use of proper orthogonal decomposition for time interpolation from PIV data. *Experiments in Fluids*, 39(6):1009–1023, December 2005.

Brittany A. Erickson, Junle Jiang, Valère Lambert, Sylvain D. Barbot, Mohamed Abdelmeguid, Martin Almquist, Jean-Paul Ampuero, Ryosuke Ando, Camilla Cattania, Alexandre Chen, Luca Dal Zilio, Shuai Deng, Eric M. Dunham, Ahmed E. Elbanna, Alice-Agnes Gabriel, Tobias W. Harvey, Yihe Huang, Yoshihiro Kaneko, Jeremy E. Kozdon, Nadia Lapusta, Duo Li, Meng Li, Chao Liang, Yajing Liu, So Ozawa, Andrea Perez-Silva, Casper Pranger, Paul Segall, Yudong Sun, Prithvi Thakur, Carsten Uphoff, Ylona van Dinther, and Yuyun Yang. Incorporating full elastodynamic effects and dipping fault geometries in community code verification exercises for simulations of earthquake sequences and aseismic slip (SEAS). *Bulletin of the Seismological Society of America*, 113(2):499–523, 01 2023. ISSN 0037-1106. doi: 10.1785/0120220066.

Rikuto Fukushima, Masayuki Kano, and Kazuro Hirahara. Physics-informed neural networks for fault slip monitoring: Simulation, frictional parameter estimation, and prediction on slow slip events in a spring-slider system. *Journal of Geophysical Research: Solid Earth*, 128(12):e2023JB027384, 2023. doi: https://doi.org/10.1029/2023JB027384.

Rikuto Fukushima, Masayuki Kano, Kazuro Hirahara, Makiko Ohtani, Kyungjae Im, and Jean-Philippe Avouac. Physics-informed deep learning for estimating the spatial distribution of frictional parameters in slow slip regions. *Journal of Geophysical Research: Solid Earth*, 130(5):e2024JB030256, 2025. doi: https://doi.org/10.1029/2024JB030256.

A.-A. Gabriel, J.-P. Ampuero, L. A. Dalguer, and P. M. Mai. The transition of dynamic rupture styles in elastic media under velocity-weakening friction. *Journal of Geophysical Research: Solid Earth*, 117(B9), 2012. doi: https://doi.org/10.1029/2012JB009468.

Xiang Gao and Kelin Wang. Rheological separation of the megathrust seismogenic zone and episodic tremor and slip. *Nature*, 543(7645):416–419, Mar 2017. ISSN 1476-4687. doi: 10.1038/nature21389.

Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*. Chapman and Hall/CRC, 1995.

J. Gomberg, A. Wech, K. Creager, K. Obara, and D. Agnew. Reconsidering earthquake scaling. *Geophysical Research Letters*, 43(12):6243–6251, 2016. doi: https://doi.org/10.1002/2016GL069967.

W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 04 1970. ISSN 0006-3444. doi: 10.1093/biomet/57.1.97.

Gavin P. Hayes, Ginevra L. Moore, Daniel E. Portner, Mike Hearne, Hanna Flamme, Maria Furtney, and Gregory M. Smoczyk. Slab2, a comprehensive subduction zone geometry model. *Science*, 362(6410):58–61, 2018. doi: 10.1126/science.aat4723.

Changrong He, Wenming Yao, Zeli Wang, and Yongsheng Zhou. Strength and stability of frictional sliding of gabbro gouge at elevated temperatures. *Tectonophysics*, 427(1):217–229, 2006. ISSN 0040-1951. doi: https://doi.org/10.1016/j.tecto.2006.05.023.

Gabrielle M. Hobson and Dave A. May. Sensitivity analysis of the thermal structure within subduction zones using reduced-order modeling. *Geochemistry, Geophysics, Geosystems*, 26(5):e2024GC011937, 2025a. doi: https://doi.org/10.1029/2024GC011937.

Gabrielle M. Hobson and Dave A. May. Sensitivity analysis of the thermal structure within subduction zones using reduced-order modeling. *Geochemistry, Geophysics, Geosystems*, 26(5):e2024GC011937, 2025b. doi: https://doi.org/10.1029/2024GC011937.

Claudia Hulbert, Romain Jolivet, Blandine Gardonio, Paul A. Johnson, Christopher X. Ren, and Bertrand Rouet-Leduc. Tremor waveform extraction and automatic location with neural network interpretation. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–9, 2022. doi: 10.1109/TGRS.2022.3156125.

Satoshi Ide and Gregory C. Beroza. Slow earthquake scaling reconsidered as a boundary between distinct modes of rupture propagation. *Proceedings of the National Academy of Sciences*, 120(32):e2222102120, 2023. doi: 10.1073/pnas.2222102120.

Junle Jiang, Brittany A. Erickson, Valère R. Lambert, Jean-Paul Ampuero, Ryosuke Ando, Sylvain D. Barbot, Camilla Cattania, Luca Dal Zilio, Benchun Duan, Eric M. Dunham, Alice-Agnes Gabriel, Nadia Lapusta, Duo Li, Meng Li, Dunyu Liu, Yajing Liu, So Ozawa, Casper Pranger, and Ylona van Dinther. Community-driven code comparisons for three-dimensional dynamic modeling of sequences of earthquakes and aseismic slip. *Journal of Geophysical Research: Solid Earth*, 127(3):e2021JB023519, 2022. doi: https://doi.org/10.1029/2021JB023519.

Naoyuki Kato. Seismic cycle on a strike-slip fault with rate-and state-dependent strength in an elastic layer overlying a viscoelastic half-space. *Earth, Planets and Space*, 54(11):1077–1083, 2002.

Hojjat Kaveh, Jean Philippe Avouac, and Andrew M Stuart. Spatiotemporal forecast of extreme events in a chaotic model of slow slip events. *Geophysical Journal International*, 240(2):870–885, 11 2024. ISSN 1365-246X. doi: 10.1093/gji/ggae417.

Ron Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. IJCAI'95, page 1137–1143, San Francisco, CA, USA, 1995. Morgan Kaufmann Publishers Inc. ISBN 1558603638.

Nadia Lapusta and Yi Liu. Three-dimensional boundary integral modeling of spontaneous earthquake sequences and aseismic slip. *Journal of Geophysical Research: Solid Earth*, 114(B9), 2009. doi: https://doi.org/10.1029/2008JB005934.

Nadia Lapusta, James R. Rice, Yehuda Ben-Zion, and Gutuan Zheng. Elastodynamic analysis for slow tectonic loading with spontaneous rupture episodes on faults with rate- and state-dependent friction. *Journal of Geophysical Research: Solid Earth*, 105(B10):23765–23789, 2000. doi: https://doi.org/10.1029/2000JB900250.

Damiana Lazzaro and Laura B. Montefusco. Radial basis functions for the multivariate interpolation of large scattered data sets. *Journal of Computational and Applied Mathematics*, 140(1):521–536, 2002. ISSN 0377-0427. doi: https://doi.org/10.1016/S0377-0427(01)00485-X.

Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.

Duo Li and Alice-Agnes Gabriel. Linking 3D long-term slow-slip cycle models with rupture dynamics: The nucleation of the 2014 Mw 7.3 Guerrero, Mexico earthquake. *AGU Advances*, 5(2):e2023AV000979, 2024. doi: https://doi.org/10.1029/2023AV000979.

Duo Li and Yajing Liu. Spatiotemporal evolution of slow slip events in a nonplanar fault model for northern Cascadia subduction zone. *Journal of Geophysical Research: Solid Earth*, 121(9):6828–6845, 2016a. doi: https://doi.org/10.1002/2016JB012857.

Duo Li and Yajing Liu. Spatiotemporal evolution of slow slip events in a nonplanar fault model for northern Cascadia subduction zone. *Journal of Geophysical Research: Solid Earth*, 121(9):6828–6845, 2016b. doi: https://doi.org/10.1002/2016JB012857.

Duo Li and Yajing Liu. Modeling slow-slip segmentation in Cascadia subduction zone constrained by tremor locations and gravity anomalies. *Journal of Geophysical Research: Solid Earth*, 122(4):3138–3157, 2017. doi: https://doi.org/10.1002/2016JB013778.

Jiun-Ting Lin, Amanda Thomas, Loïc Bachelot, Douglas Toomey, Jake Searcy, and Diego Melgar. Detection of hidden low-frequency earthquakes in Southern Vancouver Island with deep learning. *Seismica*, 2(4), Oct. 2024. doi: 10.26443/seismica.v2i4.1134.

Alan T. Linde, Michael T. Gladwin, Malcolm J. S. Johnston, Ross L. Gwyther, and Roger G. Bilham. A slow earthquake sequence on the San Andreas fault. *Nature*, 383(6595):65–68, Sep 1996. ISSN 1476-4687. doi: 10.1038/383065a0.

Peter Lindstrom. Fixed-rate compressed floating-point arrays. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2674–2683, 2014.

Yajing Liu and James R Rice. Aseismic slip transients emerge spontaneously in three-dimensional rate and state modeling of subduction earthquake sequences. *Journal of Geophysical Research: Solid Earth*, 110 (B8), 2005.

Yajing Liu and James R Rice. Spontaneous and triggered aseismic deformation transients in a subduction fault model. *Journal of Geophysical Research: Solid Earth*, 112(B9), 2007.

Yajing Liu and James R Rice. Slow slip predictions based on granite and gabbro friction data compared to GPS measurements in northern Cascadia. *Journal of Geophysical Research: Solid Earth*, 114(B9), 2009.

Yingdi Luo and Jean-Paul Ampuero. Stability of faults with heterogeneous friction properties and effective normal stress. *Tectonophysics*, 733:257–272, 2018. ISSN 0040-1951. doi: https://doi.org/10.1016/j.tecto.2017.11.006.

Elizabeth H. Madden, Thomas Ulrich, and Alice-Agnes Gabriel. The state of pore fluid pressure and 3-d megathrust earthquake dynamics. *Journal of Geophysical Research: Solid Earth*, 127(4):e2021JB023382, 2022. doi: https://doi.org/10.1029/2021JB023382.

Yohai Magen, Dave A May, and Alice-Agnes Gabriel. Reduced-order modelling of Cascadia's slow slip cycles dataset and code. https://zenodo.org/records/17575959, 2025. URL `https://zenodo.org/records/17575959`.

Rishav Mallick and Sharadha Sathiakumar. A fast numerical routine to model viscoelastic earthquake cycles in subduction zones. *ESS Open Archive*, December 2024.

Rishav Mallick, Valère Lambert, and Brendan Meade. On the choice and implications of rheologies that maintain kinematic and dynamic consistency over the entire earthquake cycle. *Journal of Geophysical Research: Solid Earth*, 127(9):e2022JB024683, 2022. doi: https://doi.org/10.1029/2022JB024683.

Chris Marone. Laboratory-derived friction laws and their application to seismic faulting. *Annual Review of Earth and Planetary Sciences*, 26(Volume 26, 1998):643–696, 1998. ISSN 1545-4495. doi: https://doi.org/10.1146/annurev.earth.26.1.643.

Takanori Matsuzawa, Hitoshi Hirose, Bunichiro Shibazaki, and Kazushige Obara. Modeling short- and long-term slow slip events in the seismic cycles of large subduction earthquakes. *Journal of Geophysical Research: Solid Earth*, 115(B12), 2010. doi: https://doi.org/10.1029/2010JB007566.

Sylvain Michel, Adriano Gualandi, and Jean-Philippe Avouac. Interseismic coupling and slow slip events on the Cascadia megathrust. *Pure and Applied Geophysics*, 176(9):3867–3891, September 2019.

Jannes Münchmeyer, Sophie Giffard-Roisin, Marielle Malfante, William Frank, Piero Poli, David Marsan, and Anne Socquet. Deep learning detects uncataloged low-frequency earthquakes across regions. *Seismica*, 3(1), May 2024. doi: 10.26443/seismica.v3i1.1185.

Shiying Nie and Sylvain Barbot. Seismogenic and tremorgenic slow slip near the stability transition of frictional sliding. *Earth and Planetary Science Letters*, 569:117037, 2021. ISSN 0012-821X. doi: https://doi.org/10.1016/j.epsl.2021.117037.

Kazushige Obara and Aitaro Kato. Connecting slow earthquakes to huge earthquakes. *Science*, 353(6296): 253–257, 2016. doi: 10.1126/science.aaf1512.

Kazushige Obara, Hitoshi Hirose, Fumio Yamamizu, and Keiji Kasahara. Episodic slow slip events accompanied by non-volcanic tremors in southwest Japan subduction zone. *Geophysical Research Letters*, 31 (23), 2004. doi: https://doi.org/10.1029/2004GL020848.

Jens Oeser, Hans-Peter Bunge, and Marcus Mohr. Cluster design in the earth sciences tethys. In *International conference on high performance computing and communications*, pages 31–40. Springer, 2006.

Tomohisa Okazaki, Takeo Ito, Kazuro Hirahara, and Naonori Ueda. Physics-informed deep learning approach for modeling crustal deformation. *Nature Communications*, 13(1):7092, November 2022.

Shinzaburo Ozawa, Hisashi Suito, and Mikio Tobita. Occurrence of quasi-periodic slow-slip off the east coast of the Boso Peninsula, Central Japan. *Earth Planets Space*, 59(12):1241–1245, December 2007.

So Ozawa, Yuyun Yang, and Eric M. Dunham. Fault-valve instability: A mechanism for slow slip events. *Journal of Geophysical Research: Solid Earth*, 129(10):e2024JB029165, 2024. doi: https://doi.org/10.1029/2024JB029165.

Zhigang Peng and Joan Gomberg. An integrated perspective of the continuum between earthquakes and slow-slip phenomena. *Nature Geoscience*, 3(9):599–607, Sep 2010. ISSN 1752-0908. doi: 10.1038/ngeo940.

Andrea Perez-Silva, Yoshihiro Kaneko, Martha Savage, Laura Wallace, and Emily Warren-Smith. Characteristics of slow slip events explained by rate-strengthening faults subject to periodic pore fluid pressure changes. *Journal of Geophysical Research: Solid Earth*, 128(6):e2022JB026332, 2023. doi: https://doi.org/10.1029/2022JB026332.

Casper Pranger, Patrick Sanan, Dave A. May, Laetitia Le Pourhiet, and Alice-Agnes Gabriel. Rate and state friction as a spatially regularized transient viscous flow law. *Journal of Geophysical Research: Solid Earth*, 127(6):e2021JB023511, 2022. doi: https://doi.org/10.1029/2021JB023511.

Alejandro Quiaro, Dawei Liu, and Mauricio D Sacchi. Nonintrusive reduced basis approximation to the solution of the Helmholtz equation: The magnetotellurics case. *Geophysics*, 90(3):WA323–WA337, 2025.

N. Ragu Ramalingam, K. Johnson, M. Pagani, and M. L. V. Martina. Advancing nearshore and onshore tsunami hazard approximation with machine learning surrogates. *Natural Hazards and Earth System Sciences*, 25(5):1655–1679, 2025. doi: 10.5194/nhess-25-1655-2025.

J. O. Ramsay and B. W. Silverman. *Functional Data Analysis*. Springer Series in Statistics. Springer New York, New York, NY, 2nd ed. 2005. edition, 2005. ISBN 9786610612420.

K Ranjith and J.R Rice. Stability of quasi-static slip in a single degree of freedom elastic system with rate and state dependent friction. *Journal of the Mechanics and Physics of Solids*, 47(6):1207–1218, 1999. ISSN 0022-5096. doi: https://doi.org/10.1016/S0022-5096(98)00113-6.

John M Rekoske, Alice-Agnes Gabriel, and Dave A May. Instantaneous physics-based ground motion maps using reduced-order modeling. *Journal of Geophysical Research: Solid Earth*, 128(8), August 2023.

John M Rekoske, Dave A May, and Alice-Agnes Gabriel. Reduced-order modelling for complex three-dimensional seismic wave propagation. *Geophysical Journal International*, 241(1):526–548, 02 2025. ISSN 1365-246X. doi: 10.1093/gji/ggaf049.

J R Rice and Y Ben-Zion. Slip complexity in earthquake fault models. *Proceedings of the National Academy of Sciences*, 93(9):3811–3818, 1996. doi: 10.1073/pnas.93.9.3811.

James R. Rice and Ji-cheng Gu. Earthquake aftereffects and triggered seismic phenomena. *Pure and Applied Geophysics*, 121(2):187–219, Apr 1983. ISSN 1420-9136. doi: 10.1007/BF02590135.

James R Rice and Simon T Tse. Dynamic motion of a single degree of freedom system following a rate and state dependent friction law. *J. Geophys. Res.*, 91(B1):521–530, January 1986.

HF Ried. The elastic-rebound theory of earthquakes. *Univ. California Publ. Bull. Dept. Geol.*, 6:413–444, 1911.

Shmuel Rippa. An algorithm for selecting a good value for the parameter c in radial basis function interpolation. *Advances in Computational Mathematics*, 11(2):193–210, Nov 1999. ISSN 1572-9044. doi: 10.1023/A:1018975909870.

Baptiste Rousset, Roland Bürgmann, and Michel Campillo. Slow slip events in the roots of the San Andreas fault. *Science Advances*, 5(2):eaav3274, 2019. doi: 10.1126/sciadv.aav3274.

A. M. Rubin and J.-P. Ampuero. Earthquake nucleation on (aging) rate and state faults. *Journal of Geophysical Research: Solid Earth*, 110(B11), 2005. doi: https://doi.org/10.1029/2005JB003686.

Allan M Rubin. Episodic slow slip events and rate-and-state friction. *Journal of Geophysical Research: Solid Earth*, 113(B11), 2008.

Cody Rucker and Brittany A. Erickson. Physics-informed deep learning of rate-and-state fault friction. *Computer Methods in Applied Mechanics and Engineering*, 430:117211, 2024. ISSN 0045-7825. doi: https://doi.org/10.1016/j.cma.2024.117211.

Andy Ruina. Slip instability and state variable friction laws. *Journal of Geophysical Research: Solid Earth*, 88(B12):10359–10370, 1983. doi: https://doi.org/10.1029/JB088iB12p10359.

S. Ruiz, M. Metois, A. Fuenzalida, J. Ruiz, F. Leyton, R. Grandin, C. Vigny, R. Madariaga, and J. Campos. Intense foreshocks and a slow slip event preceded the 2014 Iquique Mw 8.1 earthquake. *Science*, 345 (6201):1165–1169, 2014. doi: 10.1126/science.1256074.

D. A. Schmidt and H. Gao. Source parameters and time-dependent slip distributions of slow slip events on the Cascadia subduction zone from 1998 to 2008. *Journal of Geophysical Research: Solid Earth*, 115 (B4), 2010. doi: https://doi.org/10.1029/2008JB006045.

Christopher H. Scholz. Earthquakes and friction laws. *Nature*, 391(6662):37–42, Jan 1998. ISSN 1476-4687. doi: 10.1038/34097.

Susan Y. Schwartz and Juliana M. Rokosky. Slow slip events and seismic tremor at circum-Pacific subduction zones. *Reviews of Geophysics*, 45(3), 2007. doi: https://doi.org/10.1029/2006RG000208.

Paul Segall and Andrew M. Bradley. Slow-slip evolves into megathrust earthquakes in 2D numerical simulations. *Geophysical Research Letters*, 39(18), 2012. doi: https://doi.org/10.1029/2012GL052811.

Paul Segall, Allan M. Rubin, Andrew M. Bradley, and James R. Rice. Dilatant strengthening as a mechanism for slow slip events. *Journal of Geophysical Research: Solid Earth*, 115(B12), 2010. doi: https://doi.org/10.1029/2010JB007449.

John Suppe. Fluid overpressures and strength of the sedimentary upper crust. *Journal of Structural Geology*, 69:481–492, 2014. ISSN 0191-8141. doi: https://doi.org/10.1016/j.jsg.2014.07.009.

Napat Tainpakdipat, Mohamed Abdelmeguid, Chunhui Zhao, Kamyar Azizzadenesheli, and Ahmed Elbanna. Fourier Neural Operators for accelerating earthquake dynamic rupture simulations. *ESS Open Archive*, April 2025. doi: 10.22541/essoar.174526096.60853381/v1.

Ryota Takagi, Naoki Uchida, and Kazushige Obara. Along-strike variation and migration of long-term slow slip events in the Western Nankai subduction zone, Japan. *Journal of Geophysical Research: Solid Earth*, 124(4):3853–3880, 2019. doi: https://doi.org/10.1029/2018JB016738.

Dingwen Tao, Sheng Di, Zizhong Chen, and Franck Cappello. Significantly improving lossy compression for scientific data sets based on multidimensional prediction and error-controlled quantization. In *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 1129–1139. IEEE, 2017.

Carsten Uphoff, Dave A May, and Alice-Agnes Gabriel. A discontinuous Galerkin method for sequences of earthquakes and aseismic slip on multiple faults using unstructured curvilinear grids. *Geophysical Journal International*, 233(1):586–626, 11 2022. doi: 10.1093/gji/ggac467.

Aki Vehtari, Andrew Gelman, Daniel Simpson, Bob Carpenter, and Paul-Christian Bürkner. Rank-normalization, folding, and localization: An improved r for assessing convergence of MCMC (with discussion). *Bayesian Analysis*, 16(2), June 2021.

Robert C Viesca. Self-similar slip instability on interfaces with rate- and state-dependent friction. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 472(2192):20160254, August 2016a.

Robert C. Viesca. Stable and unstable development of an interfacial sliding instability. *Physical Review E*, 93:060202, Jun 2016b. doi: 10.1103/PhysRevE.93.060202.

Laura M. Wallace, John Beavan, Stephen Bannister, and Charles Williams. Simultaneous long-term and short-term slow slip events at the Hikurangi subduction margin, New Zealand: Implications for processes that control slow slip event occurrence, duration, and migration. *Journal of Geophysical Research: Solid Earth*, 117(B11), 2012. doi: https://doi.org/10.1029/2012JB009489.

Shiqi Wang. Toward quantitative characterization of simulated earthquake-cycle complexities. *Scientific Reports*, 14(1):16811, Jul 2024. ISSN 2045-2322. doi: 10.1038/s41598-024-67685-4.

D. Xiao, F. Fang, C. Pain, and G. Hu. Non-intrusive reduced-order modelling of the Navier–Stokes equations based on RBF interpolation. *International Journal for Numerical Methods in Fluids*, 79(11):580–595, 2015. doi: https://doi.org/10.1002/fld.4066.

A Mathematical symbols and definitions

Table 2: Table of mathematical symbols and definitions used throughout the study

Symbol	Definition	Dimension
τ	Fault shear stress	$\mathbb{R}^{D-1}$
$\dot{s}$	Slip rate	$\mathbb{R}^{D-1}$
ψ	State-variable in rate-and-state friction	Scalar
a, b	Empirical friction parameters	Scalar
$\dot{s}_0$	Reference slip rate	Scalar
L	Characteristic slip distance	Scalar
f_0	Reference friction coefficient	Scalar
θ	State-variable in standard aging law	Scalar
$\dot{s}_p$	Plate convergence velocity (11.7 cm/year)	$\mathbb{R}^{D-1}$
h^*	Characteristic nucleation size	Scalar
Λ	Process zone size	Scalar
d_{poly}	Polynomial degree of basis functions	Scalar
σ_W	Effective normal stress	Scalar
W_l	Up-dip extent of the low effective normal stress region	Scalar
W_r	Down-dip extent of the low effective normal stress region	Scalar
W	Width of low effective normal stress seismogenic region	Scalar
ξ	Parameter pair	$\mathbb{R}^k$
$\mathscr{P}$	2D parameter space	$\mathbb{R}^{N\times b}$
Q	Data set of simulation outputs	$\mathbb{R}^n \times \mathbb{R}^{m\times n} \times \mathbb{R}^{m\times n} \times \mathbb{R}^{m\times n}$
t	FOM simulation time steps vector	$\mathbb{R}^n$

Continued on next page

Symbol	Definition	Dimension
$\dot{S}$	FOM simulation slip rate outputs	$\mathbb{R}^{m\times n}$
Ψ	FOM simulation state-variable outputs	$\mathbb{R}^{m\times n}$
S	FOM simulation cumulative slip outputs	$\mathbb{R}^{m\times n}$
$\mathrm{ROM}_i(\cdot)$	Reduced-order model for cycle i	Function
ξ^*	Parameter pair input for the ROM ($\xi^* \notin \mathscr{P}$)	$\mathbb{R}^k$
Q^*	ROM simulation output set	$\mathbb{R}^n \times \mathbb{R}^{m\times n} \times \mathbb{R}^{m\times n} \times \mathbb{R}^{m\times n}$
q	Latent space representation of Q	$\mathbb{R}^l$
$G(\cdot)$	Simulation to spline latent space projector	Function
$G^{-1}(\cdot,\cdot)$	Spline latent space to simulation reconstruction	Function
$H(\cdot)$	Trajectory of the phase-space parametric curve $\log_{10}(\dot{s}(t), \psi(t))$	Function
$\phi(t)$	Progression variable along trajectory in phase-space	Scalar
$\mathscr{B}(\cdot,\cdot)$	B-spline transformation	Function
$\mathscr{B}^{-1}(\cdot,\cdot,\cdot)$	Inverse B-spline transformation	Function
$\mathbf{k}_x$	B-spline knot vector placed along x	$\mathbb{R}^M$
$\mathbf{c}_f$	B-spline coefficient vector evaluated along $f(x)$	$\mathbb{R}^M$
$\mathbf{D}$	Matrix of latent vector representations	$\mathbb{R}^{N\times l}$
$\mathbf{u}_k$	Basis vector of the image of $\mathbf{D}$	$\mathbb{R}^l$
α_k^j	POD coefficient	Scalar
$\varphi(\cdot)$	RBF kernel	Function
T_c	SSEs recurrence interval	Scalar
P_0	SSEs potency	Scalar
W_s	along strike fault length	Scalar
N	number of FOM evaluation	Scalar
i	FOM parameter index (e.g. ξ_i)	Scalar
n	length of time steps vector	Scalar
m	number of fault observation points	Scalar
j	index for specific observation point (e.g. Q_j)	Scalar
p	number of SSEs cycles	Scalar
k	index for specific SSE cycle (e.g. Q^k)	Scalar

B Governing Equations

The mechanical response to slip along the fault in our FOM physics-based models is solved with a quasi-dynamic approach, which requires satisfying the conservation of momentum equation (neglecting inertial terms):

$$-\frac{\partial \sigma_{ij}(u)}{\partial x_j} = f_i, \tag{34}$$

where σ_{ij} is the stress tensor, u is the displacement vector, and f_i is the body force vector. The stress tensor follows a linear elastic constitutive relationship for an isotropic material given by

$$\sigma_{ij}(u) = \Big(\lambda \delta_{ij}\delta_{kl} + \mu(\delta_{ik}\delta_{jl} + \delta_{il}\delta_{jk})\Big)\varepsilon_{kl}(u), \quad \varepsilon_{ij}(u) = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right), \tag{35}$$

where ε_{ij} is the strain tensor, λ and μ are the Lamé parameters, and δ is the Kronecker delta.

We consider three types of boundary conditions to define the Cascadia-like model setting (see Figure 1): (i) Free surface, (ii) Dirichlet boundary condition, and (iii) rate-and-state governed fault. The free surface boundary condition is applied to the top and bottom boundaries of the domain by prescribing that

$$\sigma_{ij}(u)n_j = 0, \tag{36}$$

where n_j is outward pointing unit normal vector from the free surface. Dirichlet boundary conditions are applied to the left and right boundaries and the deep portion of the central boundary. On the left and right boundaries, the displacement is prescribed by

$$u = \pm 0.5 \dot{s}_p t. \tag{37}$$

On the central boundary (fault) at depths exceeding 280 km along-dip, the displacement is prescribed via

$$u = \dot{s}_p t, \tag{38}$$

with $\dot{s}_p = (3.64, 0.64)$ cm/year being the long-term Cascadia convergence rate and t is time.

The shallow 280 km of the central boundary represents the segment of the fault governed by a rate-and-state friction law (see Section 2). On this fault interface, let $\hat{n}$ be the unit normal vector and $\hat{n}_\perp$ be the unit tangent vector in the slip direction, $\hat{n}_\perp = (\cos(10^\circ), \sin(10^\circ))$. The slip on the fault s is related to displacement via

$$u = \hat{n}_\perp s. \tag{39}$$

Given a displacement u that satisfies the boundary conditions, the traction vector on the fault is $T = \sigma(u)\hat{n}_\perp$ and the scalar normal stress is $\sigma_n = T \cdot \hat{n}$. The quasi-dynamic assumption relates the elastic slip parallel traction $T \cdot \hat{n}_\perp$ to the fault shear stress and a radiation damping term via

$$T \cdot \hat{n}_\perp = \sigma_n \tau(\dot{s}, \psi) + \eta \dot{s} \tag{40}$$

where $\tau(\dot{s}, \psi)$ is the rate-and-state dependent friction (Eq. (2)), $\dot{s}$ is the fault slip rate, ψ is the fault state-variable , and η is the damping coefficient (e.g. a viscosity).

C ROM latent space reconstruction evaluation

We evaluate the presented ROM for Cascadia's SSEs by comparing the FOM and ROM SSEs' recurrence times and potencies (See Section 3.2). While this metric is sufficient for the parameter exploration and MCMC inversion presented, for other tasks, a finer-grained comparison of the ROM and FOM simulations may be required. A direct comparison of the FOM outputs $Q(\xi) = \left(\mathbf{t}, \dot{\mathbf{S}}, \Psi, \mathbf{S}\right)$ and ROM outputs $Q(\xi^*) = \left(\mathbf{t}^*, \dot{\mathbf{S}}^{k*}, \Psi^{k*}, \mathbf{S}^{k*}\right)$ is not possible as the FOM simulations is evaluated on $\mathbf{t} \in \mathbb{R}^n$ time steps, and the ROM over $\mathbf{t}^* \in \mathbb{R}^{n*}$ where $n \neq n^*$. Here, we show how to use the simulation latent space to obtain a finer metric for evaluating ROM reconstruction.

The simulation latent space (Section 3.1.1) defines a function transforming the simulation output to a fixed-length latent state vector $G(Q) = q \in \mathbb{R}^l$. This fixed-length latent vector contains the spline knots and coefficients of the simulation m observation point slip rate ($\dot{s}$), state-variable (ψ), and cumulative slip (s). We can similarly define a latent representation for each of our simulation parameters

$$q_{\dot{s}} = (c_{\dot{s}1} \dots c_{\dot{s}m}) \tag{41a}$$
$$q_{\psi} = (c_{\psi 1} \dots c_{\psi m}) \tag{41b}$$
$$q_{s} = (c_{s1} \dots c_{sm}) \tag{41c}$$

where $c_{\dot{s}i}, c_{\psi i}$, and c_{si} are the spline coefficient samples along the simulation limit cycle of observation point i.

Similarly, stage 7 of the ROM scheme (Figure 2) is to calculate the approximated latent vector q^* corresponding to ξ^* (Eq. (27)), from which we can extract the approximated spline coefficient along the observation-phase-space trajectory.

$$q^*_{\dot{s}} = (c^*_{\dot{s}1} \dots c^*_{\dot{s}m}) \tag{42a}$$
$$q^*_{\psi} = (c^*_{\psi 1} \dots c^*_{\psi m}) \tag{42b}$$
$$q^*_{s} = (c^*_{s1} \dots c^*_{sm}) \tag{42c}$$

where all $q_{\dot{s}}, q_{\psi}, q_s, q^*_{\dot{s}}, q^*_{\psi}$, and q_s are vectors of length mK_1.

With those fixed-length vectors representing the simulation, we can use any standard vector comparison metric. In section 5.6, we show the leave-one-out cross-validation result with the relative Euclidean distance and cosine similarity defined as

$$\mathrm{Er}(q_{par}, q^*_{par}) = \frac{\|(q_{par} - q^*_{par})\|}{\|q_{par}\|} \tag{43a}$$
$$\cos\_\mathrm{sim}(q_{par}, q^*_{par}) \frac{q_{par} \cdot q^*_{par}}{\|q_{par}\| \|q^*_{par}\|} \tag{43b}$$

with q_{par}, q^*_{par} being the latent representation of any of the simulation slip rate, state-variable, and cumulative slip.

D B-spline knot placement strategy

The accuracy and efficiency of our spline-based latent space representation depends on the strategic placement of B-spline knots. An optimal knot distribution allocates more knots to regions of high functional complexity, increases representation fidelity with a minimal number of spline knots. Given the multi-scale nature of the SSE cycle simulations data, we developed a two-step, adaptive knot placement strategy for our two distinct spline mappings: (1) mapping the phase progression ϕ from the simulation time t using $\mathbf{k}_t \in \mathbb{R}^{K_0}$ knots, and (2) mapping the physical variables ($\dot{s}$, s, ψ) from the phase progression variable ϕ using $\mathbf{k}_\phi \in \mathbb{R}^{K_1}$ knots.

For the first mapping, which connects simulation time t to the phase progression ϕ, the primary challenge is the highly non-uniform distribution of time steps from the full-order model's adaptive time-stepping scheme. As shown in Figure 15, the relationship $\phi(t)$ is characterized by a long, low-gradient inter-event period, preceded and followed by an abrupt change during the SSEs. To model this, we use a combination of quantile-based and uniform knot placement along time. The adaptive time stepping used in the FOM

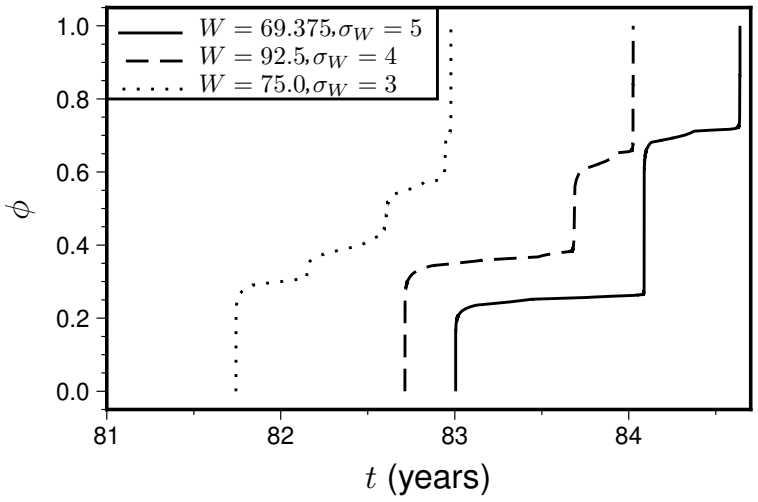

Figure 15: The relationship between simulation time t and phase-space progression ϕ. The plot illustrates the evolution of the phase-space progression variable ϕ as a function of time t for a single representative SSE cycle from three different simulations.

generates a high density of time steps during the rapid slip of an SSE, placing knots according to the quantiles of the time steps vector ($\mathbf{t}$) allocates more knots to the SSE period itself. On the other hand, uniformly placed knots ensure coverage in the inter-event period, where time steps can be sparse. This combined approach is crucial for accurately resolving the sharp onset and evolution of the slow slip event while efficiently representing the long, quasi-static interseismic period with fewer knots. Based on trial and error for the best reconstruction fidelity, a ratio of 0.7 is chosen between quantile and uniform placement, which gives $0.7K_0$ quantile knots and $0.3K_0$ uniform knots along t.

For the second mapping, which represents the trajectory in phase space, a similar approach is taken for placing knots along the ϕ. However, the mapping from ϕ to the fault variables is much smoother and lacks sharp, abrupt gradients (Figure 3, Figure 6, and Section 4.2). Consequently, a majority of the knots are placed uniformly. A ratio of 0.2 is chosen between quantile and uniform placement, which gives $0.2K_1$ quantile knots and $0.8K_1$ uniform knots along ϕ.

submitted to *Geophys. J. Int.*

Supporting Information for "Reduced-order modelling of Cascadia's slow slip cycles"

Yohai Magen[1], Dave A. May[1], Alice-Agnes Gabriel[1,2]

[1]*Institute of Geophysics and Planetary Physics, Scripps Institution of Oceanography University of California, San Diego, C*

[2]*Department of Earth and Environmental Sciences, Ludwig-Maximilians-Universität München, Munich, Germany*

(i) Text S1

(ii) Figures S1 to S9

1 INTRODUCTION

This document provides supplementary material for the main manuscript. It consists of additional figures that offer further details and supporting evidence for the methodologies, numerical simulations, and analyses presented in the primary article.

2 TEXT S1: RADIAL BASIS FUNCTION KERNEL SELECTION

The construction of the reduced-order model (ROM) involves interpolating Proper Orthogonal Decomposition (POD) coefficients using Radial Basis Functions (RBFs), as detailed in Section 3.2.2 of the main text. The choice of the RBF kernel is a critical hyperparameter that influences the accuracy of this interpolation. The following text and Figure S1 compare the performance of linear and cubic RBF kernels for this task. Our selection of linear RBFs over higher-order alternatives was validated by comparing the ROM accuracy when constructed with both linear ($\varphi(r) = r$) and cubic ($\varphi(r) = r^3$) RBFs. This comparison was conducted for both the LOOCV procedure and the three specific parametric sets used in section 3.2.2: $\boldsymbol{\xi}_1 = (W = 37.375\,\text{km}, \sigma_W = 5\,\text{MPa})$, $\boldsymbol{\xi}_2 = (W = 60.5\,\text{km}, \sigma_W = 4\,\text{MPa})$, and $\boldsymbol{\xi}_3 = (W = 43.0\,\text{km}, \sigma_W = 3\,\text{MPa})$. In general, a higher-order RBF can provide a tighter fit to the provided data points, but this comes with an increased risk of overfitting, especially if the data sampling is not sufficiently dense to support the higher complexity. The comparison, illustrated in Figure S1, confirms that in the case of our ROM scheme, the higher-order cubic RBFs result in a poorer

fit for the LOOCV tests. Our parametric space sampling, which was designed with the linear RBFs in mind, appears to be too sparse for the effective application of higher-order RBFs. This relative sparsity leads to more outliers in the predicted time series data and a worse fit to the recurrence time and potency values derived from the FOM simulations when using cubic RBFs. Consequently, the linear kernel was adopted for all ROMs presented in this study.

3 SUPPLEMENTARY FIGURES

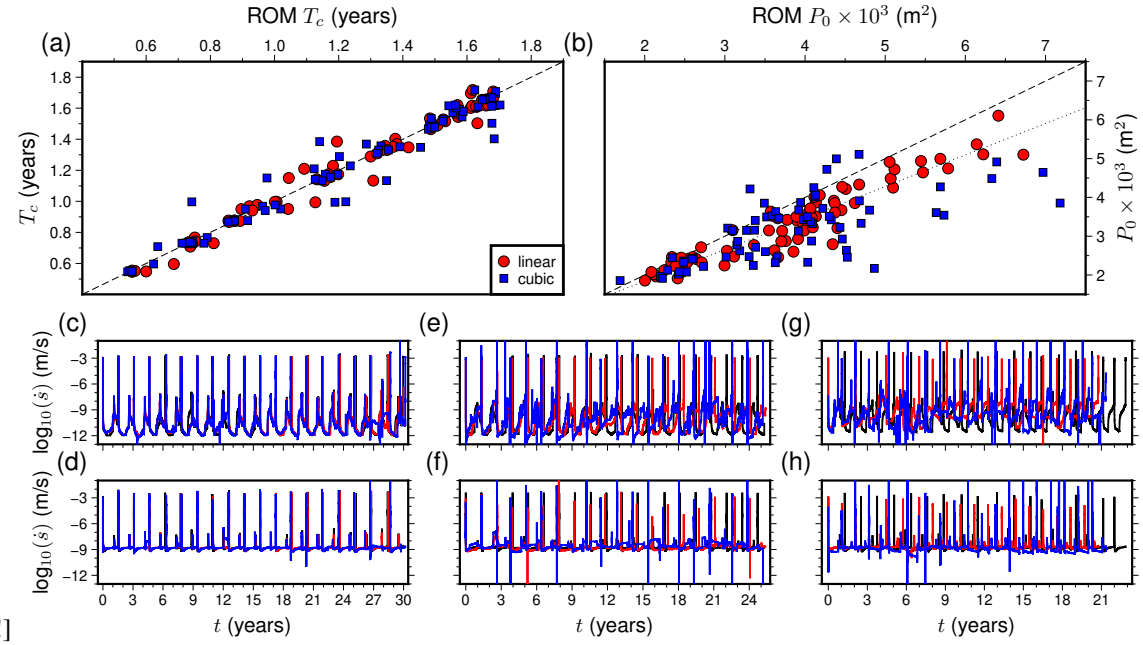

[ht!]

Figure S1. Comparison of the LOOCV results with linear (red) and cubic (blue) RBF interpolation kernels. (a) Comparison of FOM results versus ROM-predicted calculated recurrence time (T_c). The dashed line indicates a 1:1 ratio. (b) Comparison of FOM results versus ROM-predicted potency (P_0). The dashed line indicates a 1:1 ratio and the dotted line shows the best linear fit between the linear-RBFs ROM and FOM datasets. (c)-(h) shows the same results as from Figure **??** (c)-(h) with FOM shown in black, ROM reconstruction using a linear RBF kernel (red) and ROM reconstruction using a cubic RBF kernel (blue).

[ht!]

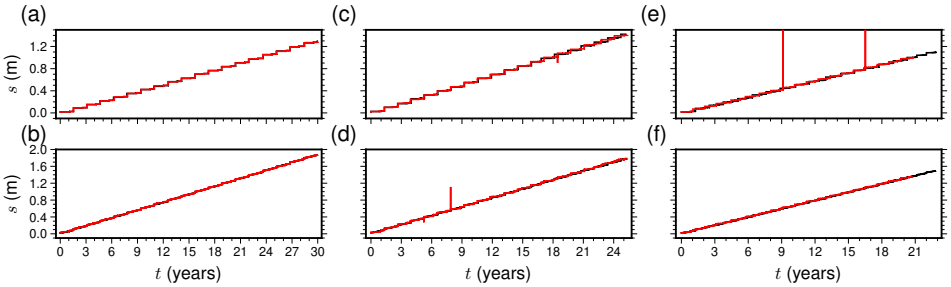

Figure S2. LOOCV comparison of cumulative slip predictions from the ROM (red curves) versus those from the FOM (black curves). Results are displayed for three different parameter sets at two along-dip observation points: 195 km (panels a, c, e) and 220 km (panels b, d, f). The parameter sets are: Panels (a,b): $\boldsymbol{P}_1$ using $W = 37.375\,\text{km}$ and $\sigma_W = 5\,\text{MPa}$. Panels (c,d): $\boldsymbol{P}_2$ using $W = 60.5\,\text{km}$ and $\sigma_W = 4\,\text{MPa}$. Panels (e,f): $\boldsymbol{P}_3$ using $W = 43.0\,\text{km}$ and $\sigma_W = 3\,\text{MPa}$.

[ht!]

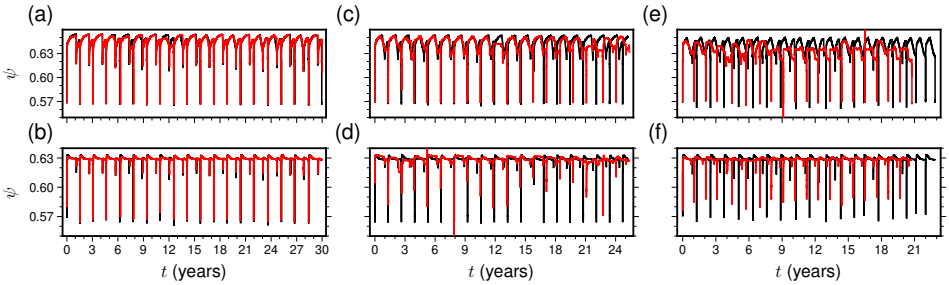

Figure S3. LOOCV comparison of state-variable predictions from the ROM (red curves) versus those from the FOM (black curves). Results are displayed for three different parameter sets at two along-dip observation points: 195 km (panels a, c, e) and 220 km (panels b, d, f). The parameter sets are: Panels (a,b): $\boldsymbol{P}_1$ using $W = 47.375\,\text{km}$ and $\sigma_W = 5\,\text{MPa}$. Panels (c,d): $\boldsymbol{P}_2$ using $W = 60.5\,\text{km}$ and $\sigma_W = 4\,\text{MPa}$. Panels (e,f): $\boldsymbol{P}_3$ using $W = 43.0\,\text{km}$ and $\sigma_W = 3\,\text{MPa}$.

[ht!]

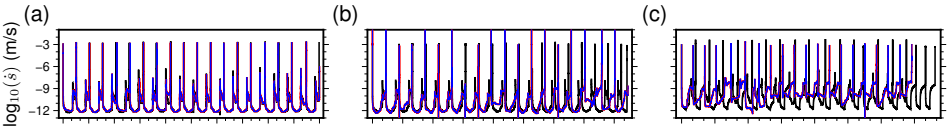

Figure S4. LOOCV comparison of slip-rate predictions from the ROM with 346 observation points (red curves) ROM with one observation point (blue curve) versus those from the FOM (black curves). Results are displayed for three different parameter sets at along-dip observation points: 195 km. The parameter sets are: Panels (a): $\boldsymbol{P}_1$ using $W = 47.375$ km and $\sigma_W = 5$ MPa. Panels (b): $\boldsymbol{P}_2$ using $W = 60.5$ km and $\sigma_W = 4$ MPa. Panels (c): $\boldsymbol{P}_3$ using $W = 43.0$ km and $\sigma_W = 3$ MPa.

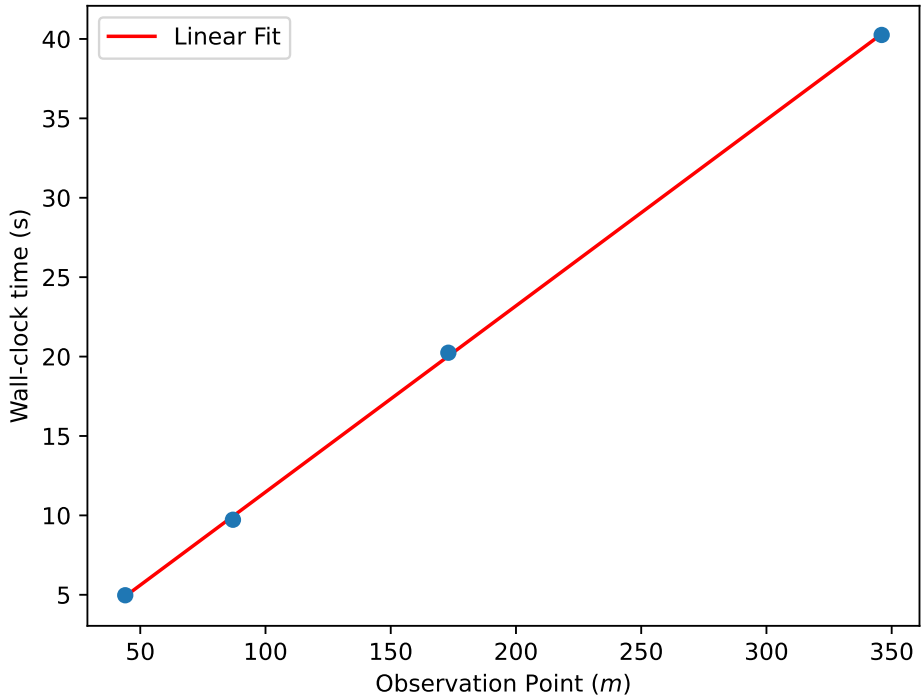

[ht!]

Figure S5. Average wall-clock time per observation point, averaged over 20 ROM inferences, as a function of the number of ROM observation points m. The ROM inference was performed on an AMD EPYC 7662 64-Core cluster using 20 cores and one worker per model inference.

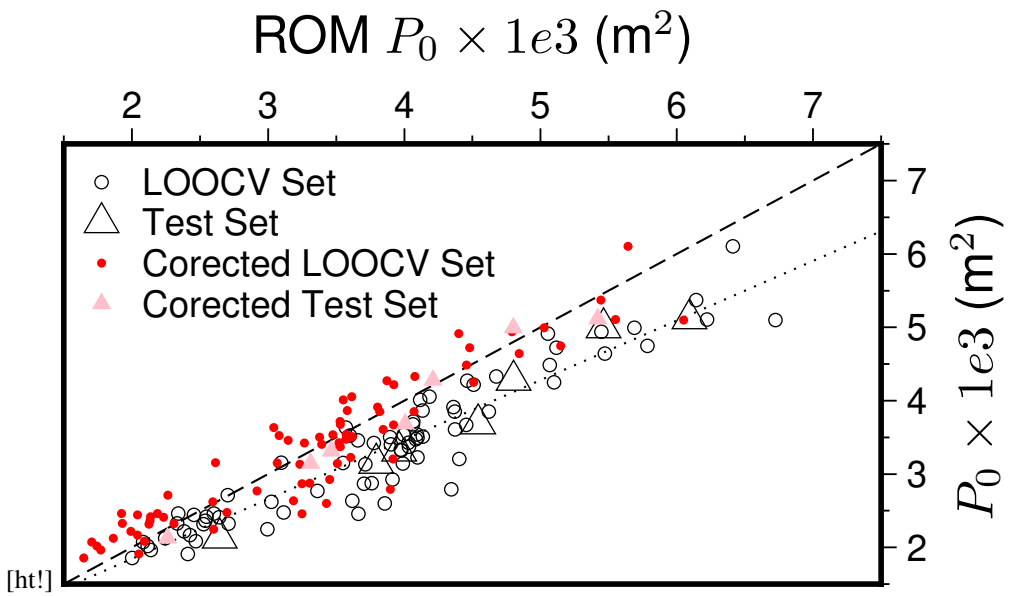

[ht!]

Figure S6. Effect of the linear post-processing correction on ROM-predicted SSEs potency (P_0) when compared against Full Order Model (FOM) values. Results are shown for both the LOOCV dataset (circles), which were used to calculate the linear correction, and the independent test set (triangles). Dashed line indicates a 1:1 correspondence and doted line best linear fit for the LOOCV not-corrected set.

[ht!]

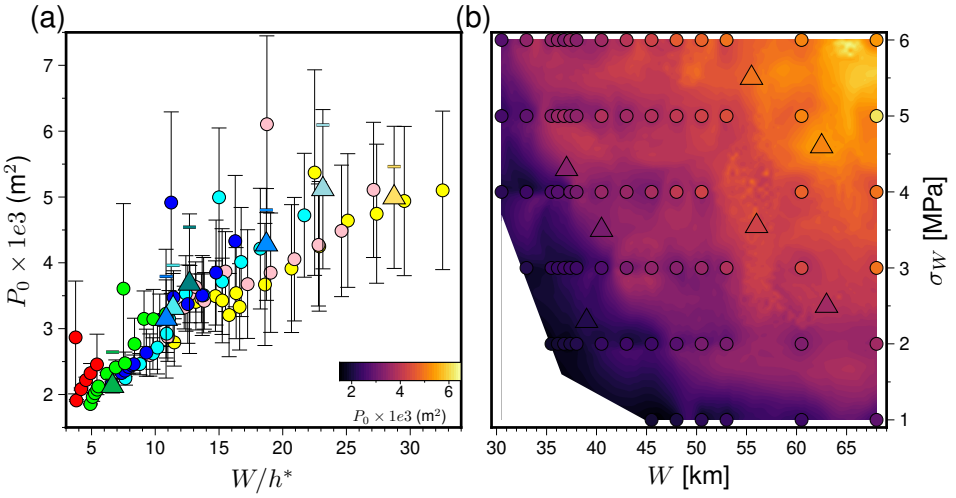

Figure S7. ROM predictions for uncorrected SSEs potency (P_0). (a) Uncorrected SSE potency (P_0) as a function of the normalized fault width (W/h^*). (b) Uncorrected SSE potency (P_0) as a function of the width (W, x-axis) and magnitude (σ_W, y-axis) of the low effective normal stress zone. These are the raw ROM potency predictions before the linear correction detailed in Section 4.3 and shown in Figure 11.

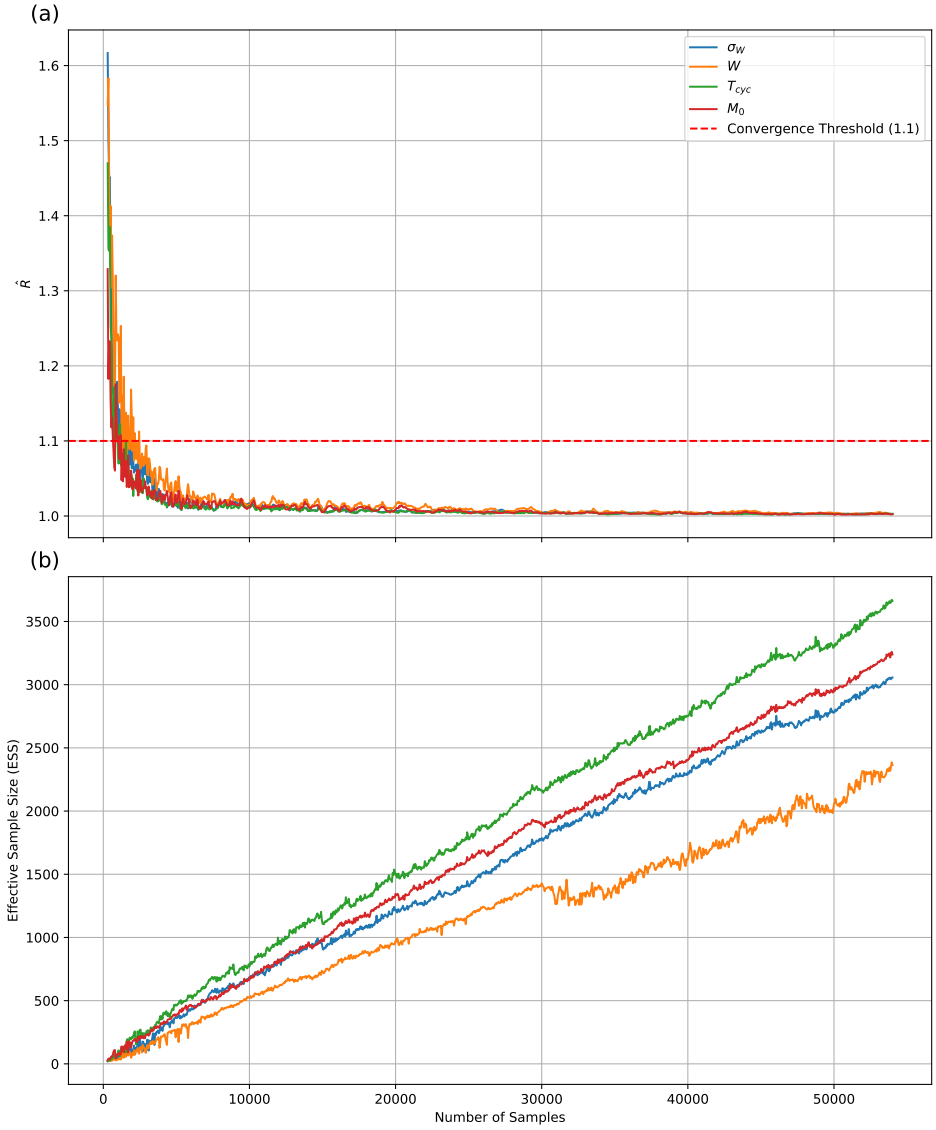

[ht!]

Figure S8. MCMC convergence diagnostics for the inversion of fault parameters W and σ_W. (a) normalized $\hat{R}$ diagnostics as a function of the total number of MCMC evaluations across 10 parallel chains (excluding burn-in). Values below 1.01 (dashed line) indicate convergence to a common target distribution. (b) Effective Sample Size (ESS) for key parameters as a function of total MCMC evaluations. An ESS above 200 is typically considered sufficient for robust posterior inference of mean and standard deviation.

[ht!]

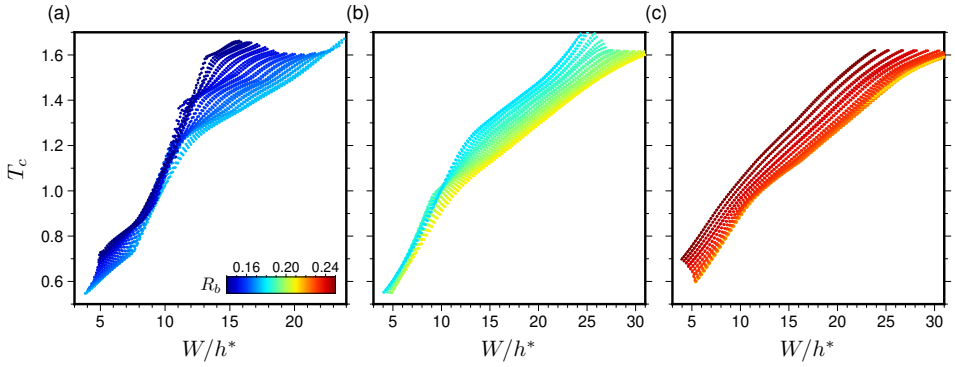

Figure S9. Dependence of SSEs recurrence time (T_c) on the normalized fault width (W/h^*), categorized by different ranges of the frictional parameter $R_b = (b - a)/b$. These results are from the Reduced Order Model (ROM) parametric space exploration. (a) Simulations with $R_b < 0.6$. (b) Simulations where $0.6 \leq R_b \leq 0.9$. (c) Simulations with $R_b > 0.9$.