

From Daily Deformation to Millennial Mechanics: Insights from Subduction Zone Earthquake Cycle Models

Thesis by
Tobias Köhne

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy in Geophysics

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2025
Defended September 3, 2024

© 2025

Tobias Köhne

ORCID: 0000-0002-8400-7255

All rights reserved except where otherwise noted

ACKNOWLEDGEMENTS

I want to thank my research advisor, Mark Simons, without whom I wouldn't have started my PhD at Caltech, and maybe none at all. His ability to see the big picture of research fields and figure out the topics that could be most disruptive is inspiring. He was always available when I felt lost and had to meet multiple times a week, and trusted me to make progress when I was radio silent for months, hunkered down over thousands of lines of code. Apart from research, I am grateful for his indulgence of my work outside of geodesy, whether it was other academic projects, volunteering, or skiing.

An additional pillar of Mark's support for me was bringing Rishav Mallick into our group. I cannot thank Rishav enough for the countless hours he helped me, from down in the weeds of debugging individual lines of code to questions about our research and beyond. His patience with me when I didn't understand new concepts was exemplary, and the breadth of his knowledge is simply impressive.

As fantastic collaborators on my research projects, I would furthermore like to thank Lijun Zhu for his work on the Altar code, Théa Ragon for providing me with 3D fault meshes of Japan, and Bryan Riel for his help with timeseries analysis software.

I would also like to thank Konstantin Batygin, who offered me the chance to dip my toes into planetary science early on in my time at Caltech — something I'd always wanted to try, and which culminated in my first first-author journal article. He and his group were wonderfully welcoming, which made my transition from engineering to science both easier and funnier.

I want to thank my other committee members for their interest in and support for my research over these last couple of years: Zachary Ross, who introduced me to machine learning and had the additional role of being my academic advisor; Jean-Philippe Avouac, who provided new geodetic perspectives on my research; and Nadia Lapusta, without whom I would have been completely lost trying to understand what friction and earthquakes actually are.

It's been incredible to have been supported by the Seismolab's administrative and technical staff, and in particular Rosemary Miller, Donna Mireles, Kim Baker-Gatchalian, Monica Nolasco, and Scott Dungan. I want to thank them for always trying to make sure that graduate student life is as easy as possible, a daunting task they excelled at.

Even though it's been many years already, I cannot omit thanking Srinivas Bettadpur and Christian Karpfinger for their role in kick-starting my academic career. Christian Karpfinger made undergraduate math at the Technical University of Munich easy and entertaining, and thereby many other subjects more accessible. In the US, it was Srinivas Bettadpur at UT Austin who gave me the chance to prove myself as a graduate researcher and introduced me to the scientific community.

A huge thank-you goes out to my family, who are an incredible source of love, support, and joy. In fact, it was my parents, Barbara and Wolfgang, who made sure I got started with learning English at a very young age, and who convinced me to apply to Caltech for a PhD when I was doubting myself. I'm grateful for the many visits and countless hours on video calls with my parents and my sisters Nora and Esther, which made the vast distances between us feel small. Not to be outdone, I also want to thank my wonderful friends from back home, especially Florian Kleinig, Karl König, Lara von Rhein, and Max Lehmann, who never fully let me go with their visits, calls, joint travels, or video game sessions.

Arguably the most significant single event during my time at Caltech was when Madison Dunitz walked into my office and literally changed my life. She showed me what love means, and I am so thankful to her for being the sweetest, smartest, funniest, and most caring, supporting, and inspirational girlfriend anyone could hope for. Her being in my life has made every day more fulfilling.

I want to thank all the people in the Graduate Student Council who I worked with for so long for their time and effort to improve Caltech. They helped me make sense of how Caltech works and thereby enabled me to feel like I belong. I am thankful to Joe Kirschvink and all the incredible TAs who organized the many Ge 136 trips I went on to satisfy my regional geotourism curiosity. Similarly, I want to thank the Alpine Club for always making sure I wouldn't go skiing, hiking, running, or camping by myself. My time in California would have been much more boring if it weren't for the adventures I went on with Angus Gruen, Ben Idini, Alex Thomson, or Xenia Boyes. And finally, last but not least, I want to thank the other graduate students in Seismolab, and in particular my office mates from South Mudd 358: Celeste Labeledz, Ethan Williams, James Atterholt, and Chenxi Kong. They made the office feel like the place I wanted to work at, and showed me what being a graduate student can look like before I knew how I wanted mine to look. It is all of these friends I made across campus, groups, and activities, who made Caltech feel less like work and more like community.

ABSTRACT

Subduction zones have hosted all the largest five earthquakes in the last one hundred years, including the 2011 M_w 9.1 Tohoku-oki earthquake on 11 March 2011, one of the largest natural disasters in history. While the general mechanism of these thrust-style earthquakes is well described by stress accumulation due to the locking between the incoming and the overriding tectonic plates, many questions remain as to the size and longevity of the asperities which host the coseismic rupture, the rheological models best describing the rock in and around the fault zone, and the effects of stress shadows and interactions between different asperities on the same plate interface. These questions are addressed by using large earthquakes as natural experiments, which we can observe using geodetic, seismic, and other techniques. However, ambiguities in the modeling results point to the inherent problem of non-uniqueness when interpreting surface observations of single events to infer complex processes at depth.

This dissertation presents a new framework to study subduction zones and their rheological properties by extending both the time period modeled and the observations considered to *all* phases of the seismic cycle, on a fault interface that experiences earthquakes at *multiple points in space and time*. The motivating concept is that the recovery of rheological parameters could be greatly improved when considering that the constitutive laws for fault material must be able to reproduce all phases of the earthquake cycle, since it is the same physical material. Here, we (1) develop a timeseries analysis software that enables the efficient processing of large geodetic networks with long timeseries, allowing us to extract the relevant subduction-zone related signal in the observations, (2) formulate a forward model that, based on ancillary historical and seismic datasets as well as a candidate rheological model, simulates surface motion over multiple earthquake cycles, and (3) use a probabilistic inverse method to estimate the best-fitting rheological parameters given the postprocessed surface deformation timeseries and model uncertainties.

We validate the timeseries analysis software on the transient volcanic deformation of Long Valley Caldera, California, USA, before extracting the megathrust component of the surface observations on Northern Honshu Island, Japan. We then estimate the rheological properties of the Northern Japanese subduction zone using our inversion method, simultaneously producing time-varying estimates of kinematic coupling, slip deficit, and surface deformation. Our model predictions match the pre- and

postseismic displacement timeseries of the 2011 Tohoku-oki earthquake well. On the steadily creeping part of the plate interface, we infer rate-dependent frictional parameters generally increasing with depth, but with second-order along-strike variation. Finally, we discuss the potential impact of our cycle-spanning, probabilistic inversion method on the field of subduction zone studies, and present possible avenues for further improvements to our framework.

PUBLISHED CONTENT AND CONTRIBUTIONS

Köhne, T. and Batygin, K. (2020). “On the Dynamical Origins of Retrograde Jupiter Trojans and Their Connection to High-Inclination TNOs.” In: *Celestial Mechanics and Dynamical Astronomy* 132.9, p. 44. DOI: 10.1007/s10569-020-09985-1.
T.K. performed all the analyses and wrote the original manuscript.

Stephenson, O. L., Köhne, T., Zhan, E., Cahill, B. E., Yun, S.-H., Ross, Z. E., and Simons, M. (2022). “Deep Learning-Based Damage Mapping With InSAR Coherence Time Series.” In: *IEEE Transactions on Geoscience and Remote Sensing* 60, pp. 1–17. DOI: 10.1109/TGRS.2021.3084209.
T.K. participated in the development of the method, wrote significant parts of the published code, and contributed to the analysis and interpretation of the results.

Köhne, T., Riel, B., and Simons, M. (2023). “Decomposition and Inference of Sources through Spatiotemporal Analysis of Network Signals: The DISSTANS Python Package.” In: *Computers & Geosciences* 170, p. 105247. DOI: 10.1016/j.cageo.2022.105247.
T.K. participated in the conception of the software package, and wrote both the code and the original manuscript.

Köhne, T., Mallick, R., and Simons, M. (n.d.). “Probabilistic Estimation of Rheological Properties in Subduction Zones using Sequences of Earthquakes and Aseismic Slip.” In revision.
T.K. participated in the conception of the framework, and wrote both the code and the original manuscript.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	v
Published Content and Contributions	vii
Table of Contents	vii
List of Illustrations	x
List of Tables	xxix
List of Select Abbreviations	xxx
Chapter I: Satellite Geodesy for Subduction Zone Modeling	1
1.1 Motivation	1
1.2 Subduction Zone Models	2
1.3 Observing Surface Deformation	4
1.4 Slip Deficit and Rheological Models	9
Chapter II: Decomposition and Inference of Sources through Spatiotemporal Analysis of Network Signals: The DISSTANS Python Package	13
2.1 Introduction	13
2.2 Background	14
2.3 Code Overview	19
2.4 Validation	22
2.5 Discussion	30
2.6 Conclusion	31
Chapter III: Probabilistic Framework to Estimate Rheological Properties in Subduction Zones using Sequences of Earthquakes and Aseismic Slip	33
3.1 Introduction	33
3.2 Fault Geometry	36
3.3 Rheological Models	37
3.4 Forward Model	40
3.5 Inverse Model	44
3.6 Summary	46
Chapter IV: Framework Performance and Sensitivities from Synthetic 2D Models	47
4.1 Introduction	47
4.2 Model Setup	47
4.3 Inversion Results	51
4.4 Implications	65
4.5 Conclusion	68
Chapter V: Estimating the Frictional Parameters of the Northern Japanese Megathrust using Earthquake Cycle Simulations	70
5.1 Introduction	70
5.2 Model Setup	71

5.3 Inversion Results	84
5.4 Discussion	100
5.5 Conclusion	116
Chapter VI: Conclusion	118
6.1 Summary	118
6.2 Potential for Future Work	119
Bibliography	124
Appendix A: DISSTANS Implementation Details	141
A.1 Structure	141
A.2 Models	143
A.3 Solver Functions	148
A.4 Data Formats	151
A.5 Synthetic Data	152
A.6 Step Detector	152
A.7 Example Workflow	153
A.8 Visualization	156
Appendix B: DISSTANS Supplemental Material	158
B.1 Synthetic Network	158
B.2 Influence of Number of Stations	163
B.3 Long Valley Caldera: Secular Velocity Comparison	165
B.4 Long Valley Caldera: Seasonal Signals	172
Appendix C: Exploration of the Northern Japan 3D Mesh	174
C.1 Finite Fault Models Defining Asperities	174
C.2 Differences Between Mesh Discretizations	175
C.3 Maximum Fit Quality of the Tohoku-oki Reference Model using Kinematic Regression	182

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
1.1 Three simplified models of subduction zone plate motion, described in Section 1.2. From left to right: Backslip model (BSM), Elastic Subducting Plate Model (ESPM), and the ESPM with added asperities in 3D. From top to bottom: the interseismic, coseismic, and cumulative relative motion after an entire earthquake cycle, respectively. Light brown lines denote the free surface (interface between a tectonic plate and the ocean or atmosphere). Arrows at an interface show the direction of the relative motion of the two sides. Black lines and regions denote no relative motion, i.e., regions where slip deficit can accumulate in the interseismic period. Blue and red lines and regions denote slow creep and fast earthquake slip, respectively. Over longer timescales, the BSM includes steady-state subduction, which would result in a net uplift of the overriding plate if modeled by block motion (gray dashed line and arrows).	4
1.2 Interseismic velocity field for Northern Honshu Island, Japan, relative to the Eurasian tectonic plate. Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). Yellow to pink lines show the plate interface at 20 km depth intervals (Hayes et al., 2018), including the plate boundary at the trench (Bird, 2003). Arrows denote the horizontal motion of analyzed GNSS stations, colored by the horizontal velocity magnitude.	6
1.3 Northern Honshu coseismic displacements due to the 2011 M_w 9.1 Tohoku-oki earthquake. Arrows denote the horizontal component, and their color the vertical component. Additional map elements as in Fig. 1.2.	7
1.4 Northern Honshu postseismic displacements due to the 2011 M_w 9.1 Tohoku-oki earthquake, cumulative until 2024. Arrows denote the horizontal component, and their color the vertical component. Additional map elements as in Fig. 1.2.	8

1.5	Kinematic coupling coefficient on the Northern Japan plate interface between 35°N and 42°N (colored mesh). Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). White triangles denote the location of the GNSS stations whose velocities were used in the inversion.	10
2.1	Map view of the synthetic network.	23
2.2	Overall model fit to the data at station Jeckle (see Fig. 2.1), including the decomposition into the primary constituents (East component only). 24	24
2.3	Map view of the transient motion of the synthetic network over the entire timespan. The top panel shows the result without spatial regularization, the bottom one with. Markers correspond to the position of a station relative to its initial position, with colors corresponding to time. The white background curves with black outlines are the true synthetic transient, which is clearly better matched by the solution incorporating spatial awareness.	25
2.4	Modeled horizontal transient displacements of selected stations in the Long Valley Caldera region during the period between 2012 and 2015. The traces begin at the nominal location of each station, with the color of the markers corresponding to the time. Background satellite imagery by Earthstar Geographics & Esri.	27
2.5	Modeled horizontal transient displacement (colored lines) of selected stations (names on the left) from Fig. 2.4, projected along the direction of maximum displacement during the period of 2012–2015. The directions (in grey to the right of the timeseries) are measured counterclockwise from east. CA99’s direction is used for CASA. Black dots are the joint model’s residuals, centered on the transient model.	28
2.6	Modeled seasonal vertical displacement timeseries for station P647 (see Fig. 2.4 for location). Blue and orange lines correspond to the annual and biannual constituents, respectively, and the black line is their sum. The deviation component only includes the annual frequency. 29	29

- 3.1 Workflow schematic, described in detail in Section 3, separated into the forward and inverse parts as shown by the dashed gray line. The forward model (Section 3.4) first uses the fault geometry (Section 3.2), asperity distribution, and earthquake history to create the stress and displacement kernels of the problem. The kernels are then used in the solution of the initial value problem given a rheological model (Section 3.3), yielding the slip history of each simulated fault patch. In the inverse problem (Section 3.5), a Markov-chain Monte Carlo sampler is used to explore the rheological parameters best matching the surface displacement timeseries. 36
- 3.2 Realization of an example earthquake history where both recurrence time T_{rec} and coseismic slip amount Δs are varied over 10 cycles. The cycle length T_{cycle} is set to 100 years, and the observations begin after 950 years (orange shading). The blue dots above vertical lines mark the time (horizontal axis) and slip amount (vertical axis) of each imposed earthquake. The slip amounts for the 10 earthquakes are sampled from a log-normal distribution with scale parameter $\sigma = 0.05$ and mean $\mu = 9.5$ m; the probability density function (PDF) is shown in blue on the left. The recurrence times are sampled from a log-normal distribution with scale parameter $\sigma = 0.05$ and mean $\mu = 100$ a; the PDF is shown in green in the insert at the top. 44
- 4.1 Fault setup for our simulations, based on the Elastic Subducting Plate Model (ESPM, Kanda and Simons, 2010). The downgoing slab is modeled by two interfaces bounding an effective elastic plate of constant thickness H . The upper and lower interfaces experience reverse and normal motion, respectively. The upper interface is furthermore assumed fully coupled (i.e., locked) from the trench down to a depth of D_{shallow} . Along the non-coupled upper interface, shear resistance is described using rheological models (see Section 3.3) with depth-varying parameters with knickpoints at D_{middle} and D_{deep} . Displacement timeseries are evaluated at discrete locations along the free surface and labeled S1–5. The length scales, observer coverage, and dipping angle are loosely based on the Northern Japan subduction zone. 48

- 4.2 Simulated surface displacement timeseries used as input for the reference case. The solid lines are the timeseries $\mathbf{d}^{\text{target}}$, with colors corresponding to observer location. The labels refer to the names of the stations, with the range of trenchward distances given as well. The black dots are the synthetic observations which include a 10 mm standard deviation Gaussian noise. At approx. 14.2 years (vertical gray line), an earthquake occurs and starts a postseismic transient process. The coseismic offset (including the effect of tapered slip) are both removed in this plot, and in the observations used in the inversion. . . . 50
- 4.3 Corner plot of the posterior covariance matrix for the reference case as approximated by the MCMC inversion process for the eight estimated parameters defining the depth-dependent power-law rheology structure (see Fig. 4.4a). D denotes the knickpoints in the depth-variable viscosity structure, and α_{eff} and n denote the effective viscosities and stress exponents at different depths, respectively. The figures on the diagonal represent smoothed 1D histograms of the marginalized posterior probability density functions (PDF) for each parameter. The numerical values of the median as well as the ± 34 th percentile (equivalent to one standard deviation for a normal distribution) are shown above the diagonal. The lower, off-diagonal plots are smoothed 2D histograms of the posterior PDF, with contour lines indicating the 1, 2, and 3 standard deviation ranges from the mean. The orange lines are the values θ_{target} used in generating the target synthetic data. The inset in the top right corner (simplified version of Fig. 4.1, described in detail in Section 4.2) relates the eight estimated parameters to their definition within the fault geometry. The effective viscosity α_{eff} is related to the rheological strength term α_n as $\alpha_{\text{eff}} = \alpha_n^{1/n} \cdot v^{1/n-1}$ (see Section 3.3). . . . 53

- 4.4 Target effective viscosity structure and posterior error for the reference case, converted to viscosity assuming $h = 1$ m. The target viscosity structure used in all cases is given in the leftmost panel (a) for the depth range of creeping fault patches. The gray curve uses the plate convergence rate to convert from the power-law rheological parameters to effective viscosity, and the blue and orange curves use the pre- and postseismic fault patch velocities, respectively. Panels (b–c) show the error in the recovered viscosity structure, as well as the recovered transition depths. Panel (b) corresponds to the preseismic apparent viscosity, and panel (c) to the postseismic one. The blue and orange solid lines show the median error between the logarithm of the recovered (predicted) viscosity profile and the target values, where ± 1 refers to an error of one order of magnitude. The light and medium shaded areas around the solid lines represent the 90th and 50th percentile ranges, respectively. The green shaded areas are horizontal histograms for the estimate of the transition depths. 54
- 4.5 Viscosity recovery in different test cases, explained in detail in Section 4.3.2. Shaded areas, lines, and histograms represent the same as in Fig. 4.4. In the leftmost panel, light blue and green still refer to the pre- and postseismic times, but all other colors now indicate the test case. 56
- 4.6 Logarithmic functional fits to the surface displacement timeseries. The different colored lines are the ensemble mean of the posterior modeled surface displacement timeseries. They almost cover the black line, which is the functional fit calculated from the entire data timespan. The gray solid and dashed lines correspond to fits using 3 and 1 years of data, respectively. 59
- 4.7 Corner plot of the posterior probability density for the rate-dependent rheology inversion. It follows the same notation as Fig. 4.3, with the rate-dependent frictional parameter α_h substituting the power-law parameters α_n and n . No orange lines are present for the frictional parameter plots since there is no straightforward conversion from the “true” effective viscosity used to generate the timeseries into the rate-dependent formulation. 60

4.8 Surface displacement observations (black dots) and fits (pink lines) using rate-dependent parameters. 960 samples are taken from the posterior probability density function, displacement timeseries are calculated, and then plotted over the noisy observations. Within uncertainties, the sampled timeseries all overlap. 61

4.9 Fault setup of the tests with forward models containing a viscoelastic bulk. The downgoing slab is modeled by two interfaces bounding an effective elastic plate of constant thickness H . Above it, the continental crust is modeled as an elastic layer and a “cold nose” touching the slab up to a depth of 50 km. The upper and lower interfaces experience reverse and normal motion, respectively. The upper interface is furthermore assumed fully coupled (i.e., locked) from the trench down to a depth of 35 km (brown line). Between 35–50 km, the interface is creeping following rate-strengthening friction with a constant value of α_h (yellow line). Every other interface obeys a no-slip boundary condition with the viscous bulk (purple lines). The viscosity below the oceanic plate η_o is constant, whereas the one below the continental plate η_c is depth-dependent. Everything else as in Fig. 4.1. 62

4.10 Surface displacement observations (black dots) from a forward model including a shallow frictional interface up to 50 km depth and a viscoelastic bulk. In pink to green lines, the best-fitting surface displacements from forward models only containing a frictional interface up to 150 km depth and no viscoelastic bulk. Each fitted line is a sample from the posterior probability density function of a probabilistic inversion of the tested observations. The preseismic and the horizontal postseismic observations are well-matched by the simpler forward model. 63

4.11 Correlation between the target, synthetic (input, horizontal axis) and the estimated (output, vertical axis) rate-strengthening parameter $\alpha_h = (a - b)\sigma_E$ for the test case where the forward model includes viscoelastic bulk flow, but the inverse model does not. The diagonal line represents perfect 1:1 correlation. The displacement error δ_d for each case is given in the color of the markers. 64

5.1 Visualization of Table 5.1, using the assumed last rupture dates. Colors denote the different asperities as shown in subsequent figures. . . 74

- 5.2 Map view of the 3D subduction zone interface between the incoming Pacific plate from the east and the overriding Eurasian plate in the west. The triangular discretized fault patches are shaded by depth, and the asperities are marked by different colors (introduced in Fig. 5.1). White triangles denote the locations of GNSS stations processed for the inversion. Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). 76
- 5.3 Same as Fig. 5.2, but using a coarser approximation of the interface to reduce the computational cost of the forward model. 77
- 5.4 Coseismic slip distribution for the 2011 Tohoku-oki earthquake, including in adjacent asperities. The asperities included in the inversion framework are shown by the colored outlines (same colors as in Fig. 5.1). Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). Yellow to pink lines show the plate interface at 20 km depth intervals (Hayes et al., 2018), including the plate boundary at the trench (Bird, 2003). White triangles denote the locations of GNSS stations processed for the inversion. . . 78
- 5.5 Visualization of the grid search performed to optimize the DISSTANS postprocessing hyperparameters that produce the input timeseries for the inverse model. All cases assume a base scalar L0 penalty factor (penalty of `spatialfit()`) of 100. The hyperparameters explored are a factor f to increase the penalty for the vertical component, the `eps` (ϵ) and `scale` (s) values for the `ReweightFunction`, and the `spatial_reweight_percentile` (q) of `spatialfit()`. For each hyperparameter combination, the root-mean-squared residuals between the modeled secular velocity field and smooth secular velocity fields (as described in Section 5.2.2) are shown as two triangular halves of the rectangular grid element. The color of each triangular patch corresponds to the value of the RMS residual. The optimal hyperparameter combination minimizing both horizontal and vertical residuals is highlighted by the black rectangle (including the corresponding RMS residual values). 80

- 5.6 Example timeseries decomposition for station 0031 for the timespan 1996–2024 in the East, North, and Up components. The top panel shows the input data (daily position solutions), with outliers removed. Using DISSTANS, the timeseries is split into the constituents modeled by our inversion framework (secular and postseismic motion) and those not modeled (coseismic, seasonal, and other transient signals such as volcanic or slow slip events). To preserve potential unmodeled signal, the residuals are not removed from the postprocessed dataset. 82
- 5.7 Values of four example basis functions in the down-dip and along-strike dimensions. The value of each basis is given as the darkness and transparency of the color, with the colors corresponding to the different bases, and lighter values implying smaller values. At each patch, the sum of all basis functions approximately sums to one, creating a suitable basis to parametrize a field of scalar values. 83
- 5.8 Mean fault slip rate on the Northern Japanese megathrust for the reference model inversion at six different times (two pre- and four postseismic). The discretized interface is colored according to the base-10-logarithm of the ratio between inferred slip rate and the assumed plate convergence rate. The cumulative fault slip timeseries at the patches presented with black outlines are shown in Fig. 5.9. The black cross indicates the location of the Kamaishi repeating earthquake sequence, its relation to the results is discussed in Section 5.4.6. Additional map elements as in Fig. 5.4. 87
- 5.9 Cumulative fault slip since 1996 at select patches on the interface. The vertical axis is displacement [m] and the horizontal axes is time. Lines and shaded areas show the sample mean and ± 3 standard deviations of the forward model output based on the estimated rheological parameters. The range of plate rates on the interface is shown by the shaded gray region for comparison. 88

- 5.10 Observed and modeled timeseries for the reference model inversion at select stations. The vertical axis is displacement [m] and the horizontal axes is time. Trenchward, along-strike, and vertical data components are shown in blue, orange, and green, respectively. Circles show the observations to be fitted, reset to zero after the 2011 Tohoku-oki earthquake. Lines and shaded areas show the sample mean and ± 3 standard deviations of the forward model output based on the estimated rheological parameters in the reference case using the coarse mesh. The dotted lines show the forward model output using the fine mesh using the same estimated rheological parameters. The center bottom panel shows a zoomed-in map view of the station network, where the stations shown in this figure are labeled and marked with black triangles. The Miyagi and Tohoku asperities are shown as orange and green polygons. 89
- 5.11 Interseismic velocity residuals for Northern Honshu Island, Japan. The residuals are defined as the difference between the observed velocities and the sample mean velocities of the forward model output. The arrows show the horizontal residual, colored by the vertical residual. The asperities included in the inversion framework are shown by polygons (same colors as in Fig. 5.1). Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). Yellow to pink lines show the plate interface at 20 km depth intervals (Hayes et al., 2018), including the plate boundary at the trench (Bird, 2003). 90
- 5.12 Interseismic maximum residual shear strain rate (square root of the second invariant of the deviatoric strain rate tensor), calculated from the horizontal velocity residuals shown in Fig. 5.11 using the VISR algorithm (Shen et al., 2015) as implemented in DISSTANS (Köhne et al., 2023). Other map content as in Fig. 5.11. 91
- 5.13 Interseismic observed (filled circles) and modeled (background field) vertical velocities for Northern Honshu. Other map content as in Fig. 5.11. 92

5.14	Postseismic observed and modeled trenchward cumulative displacements for the timeframe 2011–2024 in Northern Honshu. The background field shows the modeled trenchward cumulative displacement. The arrows show the observed cumulative displacement at each of the considered GNSS stations, colored by its trenchward component, using the same color scheme as for the model output. The closer the arrow colors are to the background field color where the arrow originates, the better the fit. Other map content as in Fig. 5.11.	93
5.15	Same as Fig. 5.14 but for the along-strike (North along trench) component.	94
5.16	Same as Fig. 5.14 but for the vertical component (with filled circles instead of arrows).	95
5.17	Postseismic cumulative displacement residuals for the timeframe 2011–2024 in Northern Honshu. The residuals are defined as the difference between the observed cumulative displacement in 2024 and the sample mean displacement of the forward model output. The arrows show the horizontal residual, colored by the vertical residual. Other map content as in Fig. 5.11.	96
5.18	Mean and standard deviation of the base-10-logarithm of the frictional parameter $\alpha_h = (a-b)\sigma_E$ [Pa] of the discretized plate interface. Assumed parameter bounds are $5 \leq \log_{10} \alpha_h \leq 10$. North of the Sanriku, south of the Ibaraki asperities, and below 200 km, the frictional parameter is set to $\alpha_h = 10^{10}$ Pa to mimic steady plate convergence far away from the model domain, and so we do not interpret the apparent variations or inferences in these border regions (red dashed line). Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). The asperities are marked by different colors (introduced in Fig. 5.1). Yellow to pink lines show the plate interface at 20 km depth intervals (Hayes et al., 2018), including the plate boundary at the trench (Bird, 2003). White triangles denote the locations of GNSS stations processed for the inversion.	98
5.19	Same as Fig. 5.18a but for test case (7). The logarithmic uncertainty everywhere within the estimated region is 0.1 or less (not shown; compare Fig. 5.18b).	99

5.20	Excess fault slip on the Northern Japanese megathrust for the reference model inversion at six postseismic times. The discretized interface is colored according to the total fault slip between the logarithmically-spaced timestamps, excluding plate convergence contributions. The patches outlined in black are the same as in Fig. 5.8. Additional map elements as in Fig. 5.4.	102
5.21	Mean excess moment magnitude over the entire creeping fault interface for the interval between the observation timestamps. Does not include contributions of the plate convergence.	103
5.22	Correlation between horizontal surface displacement at the observing stations for an early (2011-05-17–2011-07-17) and late (2019-07-26–2023-10-29) time window. Blue and orange dots and dashed lines show the surface displacements at the observing stations and their best-fit linear slope for the model predictions and observations, respectively. The gray line corresponds to a 1:1 correlation, but note that linear slopes larger or smaller than one only mean the logarithmic time window growth is smaller or larger than required by the postseismic transient, respectively.	103
5.23	Residual between the displacement over a late time frame (2019-07-26–2023-10-29) and the displacement calculated from the best-fit correlation slope (see Fig. 5.22) using the displacement in an early time frame (2011-05-17–2011-07-17). Correlation residuals are shown for the observations (filled circles) and reference model results (background field). Other map content as in Fig. 5.11.	104
5.24	Comparison of the preseismic coupling ($1 - v/v_p$) and accumulated slip deficit ($1 - \int v(t)/v_p dt$) as output from the reference inversion. Trench, slab depths, and GNSS station locations as in Fig. 5.18.	107
5.25	Mean and standard deviation of the base-10-logarithm of the frictional parameter $\alpha_h = (a - b)\sigma_E$ [Pa] of the discretized plate interface for inversion case (9), ignoring all vertical data in the fitting process. Map description and additional elements as in Fig. 5.18.	110
5.26	Difference of the base-10-logarithm of the frictional parameter α_h between reference case (1) and the inversion test case (9), which ignores all vertical data in the fitting process. Additional map elements as in Fig. 5.18	110

5.27	Rate-dependent $a-b$ parameter derived from α_h assuming hydrostatic pressure. Additional map elements as in Fig. 5.18	111
5.28	Condensed corner plot showing the posterior sample distribution of the 20 estimated α_h spline coefficients spanning the interface and yielding the frictional parameter values at each patch as shown in Fig. 5.18. The figures on the diagonal represent smoothed 1D histograms of the marginalized posterior probability density functions (PDF) for each parameter. A muted orange background indicates spline coefficients at the boundary of the model domain (i.e., the northernmost, southernmost, and deepest coefficients). The numerical values of the median as well as the ± 34 th percentile (equivalent to one standard deviation for a normal distribution) are shown above the diagonal. The lower, off-diagonal plots are smoothed 2D histograms of the posterior PDF, with contour lines indicating the 1, 2, and 3 standard deviation ranges from the mean. Most non-boundary coefficients are well recovered and are similar to Normal distributions. Boundary coefficients show one-sided PDFs skewed towards the imposed limits of the parameter range.	113
5.29	Slip rate in the Kamaishi region from the seismic analyses of Uchida et al. (2012, orange markers) and Uchida et al. (2015, green markers), and the SEAS forward model results on a fine mesh (blue line) using frictional parameters of the reference model. The date of the 2011 Tohoku-oki earthquake is shown in by the vertical black line.	116
A.1	Code structure of DISSTANS, explained in detail in Appendix A.1.	142
A.2	Example of BSpline (left) and ISpline (right) spline functions for a single timescale and order at seven different center times. Each curve on the right is the integral of the curve on the left of the same color.	146
A.3	Visualization of the intermediate functions used by AmpPhModulatedSinusoid . The spline basis functions h_j over the considered time interval are in the left panel. Multiplying the spline functions with the cosine and sine of a given period then yields the modulated cosines and sines in the center and right panel, respectively.	147

- A.4 Flowchart of the spatiotemporal L0-regularized solver as described in Riel et al. (2014). Symbols and colors from Fig. A.5. At each station, an L1-regularized least-squares fit is computed, where each parameter has an associated weight. The weight is inversely correlated to the parameter magnitude. Parameters close to zero are iteratively penalized, whereas significant parameters have their penalty gradually reduced to zero. Iterated L1 regularization effectively approximates an L0-regularized solution (see Candès et al., 2008). By combining the weights between stations with a median in an intermediate step, parameters that are significant at other nearby stations as well are promoted, and parameters that are insignificant are demoted. 151
- A.5 Example workflow for using DISSTANS, explained in detail in Appendix A.7. Blue rectangles represent single computational steps, orange rectangles with cut corners represent sub-workflows discussed in more detail elsewhere, and green, rounded rectangles represent datasets at their different stages of processing. The numbered steps in the text correspond to the numbering in the top left corners of the rectangles. 154
- A.6 Preprocessing sub-workflow, following the same symbolic and coloring as Fig. A.5 (step 2), with rose circles representing mathematical operations. First, a running median of the input is calculated, which results in a lowpass filtered timeseries. The variance of the input around the lowpass timeseries is used to detect outliers. Removing them from the input yields the outlier-free input. Without common mode estimation, this is also the final output. To remove the common mode, the difference between the lowpassed input and the outlier-free input is calculated, which yields an outlier-free, high-passed input. The dominant component of this timeseries is the best estimate of the common mode error. Removing this from the outlier-free input yields the outlier-free, common-mode-removed output. . . . 155

- B.1 Convergence of the iterative, spatiotemporal L0-regularized solver (line) for 8416 total spline coefficients. For reference, values for the local L1-regularized (circles) and local L0-regularized (diamonds) solutions are also plotted on the axes. The results are shown both for the individual components (blue and orange, right vertical axis), as well as the overall solution (black, left vertical axis). The total number of non-zero parameters (i.e., the sum over all stations and components of the number of non-zero coefficients) as well as the number of *unique* non-zero parameters (i.e., the number of all splines that are non-zero at least at one station, per component) converges monotonically onto their final values. The latter number specifically demonstrates the effect of spatial sparsity. 158
- B.2 Model parameter correlation matrix for all models and both east-north components. The covariances for spline parameters that are estimated to be close to zero are set to zero as well, and not shown. Tradeoffs between models and within splines are clearly identifiable. 159
- B.3 Timeseries comparison for the two neighboring stations Jeckle and Cylon and three regularization schemes in the two east and north components. Black dots are the synthetic observations, with the grey shading corresponding to three standard deviations of simulated observation uncertainty. The blue line represents the final overall model fit. The fitted models are virtually indistinguishable between regularization schemes, except for the overfitting of some colored noise at station Cylon, which is reduced with the spatial L0 regularization. . . 160

B.4 Scalograms of the transient model for the stations and regularization in Fig. B.3. The horizontal and vertical axes correspond to time and the discrete periods of the splines, respectively. Patches (colored by the spline coefficient's value) in this time-period-space represent a single spline in the dictionary, with their extent in time defined as the active period of the spline (i.e., having non-zero gradient), and their height defined by the relative magnitude of the particular spline compared to all splines active at that time. Using the L1 solver, the transients (two shortterm, one longterm) are sparsely fitted in time, but not in space (i.e., each station's timeseries is fit using different splines). The local L0 regularization does not change this general behavior. Spatial L0 regularization leads to the transients being sparsely fitted in time *and* space (i.e., every station's timeseries is fit with a similar set of splines). Modeling the transients with coefficients sparse in time, space and period is beneficial in the context of identifying signals close to the noise floor that are appearing at multiple stations, since the respective coefficients will be penalized less, allowing for a more physically-consistent decomposition. Conversely, the penalization of coefficients that are only seen at isolated stations makes it easier to identify local shortterm noise processes. 162

- B.5 Comparison of the secular velocity estimates (East component) presented in Section 2.4.1, Fig. 2.2 with other methods. The dark gray dots correspond to the sum of true secular velocity and the generated noise; the light gray dots additionally include the transient constituents. The true secular velocity (orange line) and secular velocity estimated by the spatial L0 solver (blue line) correspond to the same lines in Fig. 2.2. The secular velocity estimated by the local L0 solver (compare Fig. 2.3) is shown with the green line (essentially overlapping the spatial L0 result). The three other lines correspond to other commonly used methods: local, unregularized least squares with (red) and without (purple) steps at the transient center times (assumed to be known); and the MIDAS solution (brown). In our synthetic example, the presence of the transient signal throughout most of the considered timeseries is significant enough to heavily deteriorate the estimated produced by the simple least squares as well as the MIDAS methods. For this station, and in the east component, adding spatial awareness only provides a minimal benefit, although Table B.1 shows that across components and stations, there are significant benefits. 162
- B.6 Sample mean (ϵ , colored lines) and sample standard deviation (σ_ϵ , vertical error bars) for all the sampled subnetworks as a function of the number of stations used in the solution process (n , horizontal axis), and noise level ratio (σ , different colors). The dashed, horizontal grey line corresponds to the maximum allowable error if no transient signal is fitted at all. The dotted grey line is a reference line parallel to $1/\sqrt{n}$ 164

- B.7 Upper panel: Modeled horizontal secular velocities in the study region for the three different solutions DISSTANS (Section 2.4.2), MIDAS (Blewitt et al., 2016), and GAGE (Herring et al., 2016) (in GAGE’s North America-fixed reference frame). The Caldera ring fault (USGS Quaternary Fault Database, Bailey, 1989) is shown in purple. Uncertainties are shown as one-standard-deviation ellipses for each solution. The green rectangle shows the extent greater Long Valley Caldera Region (LVCR). Lower panel: Same as upper, zoomed into the LVCR. While the DISSTANS-derived velocities mostly match the published velocities outside the LVCR, they are significantly different within the LVCR. 167
- B.8 Background velocity fields as calculated by the best-fit Euler pole for the entire study area and the Long Valley Caldera Region in the upper and lower panels, respectively. Uncertainties, fault outlines, and colors are the same as in Fig. B.7. The DISSTANS-derived background velocity field slightly differs from the MIDAS- and GAGE-derived fields, but exhibit the same overall pattern. 168
- B.9 Residual secular velocities for the entire study area and the Long Valley Caldera Region (LVCR) in the upper and lower panels, respectively. Uncertainties, fault outlines, and colors are the same as in Fig. B.7. The DISSTANS solution has smaller residual velocities in the LVCR than the MIDAS and GAGE solutions. (Note the different vector scale length.) 169
- B.10 Unregularized (average) component of the vertical seasonal fitted annual model. The marker size corresponds to the amplitude, and the color corresponds to the time of the sinusoid’s maximum during the year. Background satellite imagery by Earthstar Geographics & Esri. 173
- B.11 Full vertical seasonal model (unregularized and regularized, annual and semi-annual) for the selected stations from Fig. 2.5, in the same order. Black dots are the overall model’s residuals, centered on the seasonal model. 173

C.1	Interseismic velocity differences for Northern Honshu between forward models using the fine and coarse mesh. The arrows show the horizontal difference, colored by the vertical difference. The asperities included in the inversion framework are shown by outlined and filled polygons for the fine and coarse meshes, respectively (same colors as in Fig. 5.1). Other map content as in Fig. 5.11.	177
C.2	Postseismic cumulative displacement difference for the timeframe 2011–2024 in Northern Honshu between forward models using the fine and coarse mesh. The arrows show the horizontal difference, colored by the vertical difference. The asperities included in the inversion framework are shown by outlined and filled polygons for the fine and coarse meshes, respectively (same colors as in Fig. 5.1). Other map content as in Fig. 5.11.	178
C.3	Correlation between the predicted surface displacements at all timesteps for forward models using the coarse and fine mesh, split up into the three data components, and offset for visual clarity. Blue, red, and green dots refer to the East, North, and Up components, and are shaded by distance to the Miyagi asperity, with smaller distances marked by darker colors. The joint best-fit linear correlation slope for all components is shown by the dashed gray line; the solid gray line shows a 1:1 correlation.	179
C.4	Mean fault slip rate on the Northern Japanese megathrust for the forward models using the fine mesh at six different times (two pre- and four postseismic). The discretized interface is colored according to the base-10-logarithm of the ratio between inferred slip rate and the assumed plate convergence rate. Additional map elements as in Fig. 5.8.	180
C.5	Mean excess moment magnitude over the entire creeping fault interface for the interval between the observation timestamps for the forward models using the coarse and fine mesh. Does not include contributions of the plate convergence.	181
C.6	Fault slip rate on the Northern Japanese megathrust as inferred by the kinematic inversion at six different times (two pre- and four postseismic). The discretized interface is colored according to the base-10-logarithm of the ratio between inferred slip rate and the assumed plate convergence rate. Trench, slab depths, and GNSS station locations as in Fig. 5.18.	183

- C.7 (Same as Fig. 5.10 but including the kinematic inversion results.) Observed and modeled timeseries for the reference model and kinematic inversion at select stations. The vertical axis is displacement [m] and the horizontal axes is time. Trenchward, along-strike, and vertical data components are shown in blue, orange, and green, respectively. Circles show the observations to be fitted, reset to zero after the 2011 Tohoku-oki earthquake. Continuous lines and shaded areas show the sample mean and uncertainty of the cycle model output based on the estimated rheological parameters. Dashed lines show the best-fit displacement timeseries from the kinematic inversion. The center bottom panel shows a zoomed-in map view of the station network, where the stations shown in this figure are labeled and marked with black triangles. The Miyagi and Tohoku asperities are shown as orange and green polygons. 185
- C.8 Interseismic velocity residuals from the kinematic inversion for Northern Honshu Island, Japan. The residuals are defined as the difference between the observed velocities and the best-fit velocity from the kinematic inversion. The arrows show the horizontal residual, colored by the vertical residual. The asperities included in the inversion framework are shown by polygons (same colors as in Fig. 5.1). Trench and slab as in Fig. 5.18. 186
- C.9 Postseismic cumulative displacement residuals from the kinematic inversion for the timeframe 2011–2024 in Northern Honshu. The residuals are defined as the difference between the observed cumulative displacement in 2024 and the best-fit displacement from the kinematic inversion. The arrows show the horizontal residual, colored by the vertical residual. Trench, slab, and asperities as in Fig. C.8 187
- C.10 Excess moment magnitude over the entire creeping fault interface for the interval between the observation timestamps for the kinematic and cycle models. Does not include contributions of the plate convergence. 188

LIST OF TABLES

<i>Number</i>	<i>Page</i>
4.1 Rheological parameters α_{eff} , n used to create the synthetic target dataset of surface displacement timeseries, and their respective knick-point depths D . Note that α_{eff} is defined in (3.3) and depends on a reference slip velocity.	49
4.2 Average viscosity and surface displacement errors for different inversion experiments. Metrics defined in eqs. (4.2) and (4.3)	55
5.1 Recurrence interval T_{rec} , characteristic slip magnitude s_0 , and last observed and assumed dates of rupture for the 4 asperities identified on the megathrust offshore Northern Honshu.	74
5.2 Inversion performance of the additional test runs (2–9) relative to the reference case (1) presented in Section 5.3.1. Quality metrics defined in eqs. (5.1) and (5.2).	100
B.1 Root-Mean-Squared-Error between the secular velocity estimates and the true secular velocity, averaged across the entire network, for the methods presented in Fig. B.5, and both data components individually. The spatial L0 solution significantly outperforms the other solutions (including the local L0 solution).	159
B.2 Root-Mean-Squared (RMS) residual magnitudes (in mm/a) between the modeled and background horizontal secular velocities for this study (DISSTANS) as well as the published velocities from MIDAS (Blewitt et al., 2016) and GAGE (Herring et al., 2016). The rows correspond to different subsets of the data over which the RMS is calculated. “Inside LVCR” corresponds to the stations shown in the lower panels of Figs. B.7–B.9, “All” to the ones shown in the upper panels of Figs. B.7–B.9, and “Outside LVCR” to the ones that are in the latter but not in the former. For stations outside the LVCR, the models produce approximately the same residual RMS (approx. $\pm 3\%$), but within the LVCR, DISSTANS reduces the residuals by approx. 25–47%, respectively.	166

LIST OF SELECT ABBREVIATIONS

<i>Abbreviation</i>		<i>Introduced on page</i>
BSM	Backslip Model	2
CPU	Central Processing Unit	31
ESPM	Elastic Subducting Plate Model	3
FEM	Finite Element Method	121
GNSS	Global Navigation Satellite System	5
GPU	Graphics Processing Unit	42
ICA	Independent Component Analysis	16
IVP	Initial Value Problem	42
LVC(R)	Long Valley Caldera (Region)	25
MCMC	Markov-chain Monte Carlo	44
PCA	Principal Component Analysis	16
PDF	Probability Density Function	45
RDF	Rate-dependent Friction	38
RMS	Root-mean-square	79
RSF	Rate-and-state Friction	33
SEAS	Sequences of Earthquakes and Aseismic Slip	35

Chapter 1

SATELLITE GEODESY FOR SUBDUCTION ZONE MODELING

1.1 Motivation

Earthquakes are some of the most powerful natural phenomena witnessed by humanity. Their study is therefore deeply relevant to society, since it not only aims to satisfy the innate human curiosity to explain the world we live in, but also because it can provide tools to enhance societal resilience in the face of potential future ruptures. In the late 19th and early 20th century, large earthquakes such as the 1906 San Francisco one gave final evidence that shear sliding of rocks in the subsurface was the direct cause of the events (Scholz, 2002). However, it was not until the theory of plate tectonics emerged that the driving force generating earthquakes could largely (albeit not exclusively) be attributed to the motion of the plates (for early overviews, see, e.g., Wegener, 1920; Sugimura and Uyeda, 1973).

There are now various techniques in use to investigate subsurface sliding, both on short (e.g., during earthquakes) and long timescales (e.g., steady continental motion). This dissertation focuses on the use of geodetic observations; more specifically, measurements of the movement and deformation of the surface of the Earth. The main power of geodesy for the purposes of earthquake and plate tectonic studies comes from the fact that, at the scales relevant for plate tectonics, the Earth behaves largely elastically, meaning that there are simple, linear relationships between motion below Earth's surface and the observed motion at the surface. There are additional physical processes happening in the Earth's crust apart from plate tectonics (e.g., the hydrological cycle, hot spot volcanism, glacial rebound, or anthropogenic sources), and so differentiating between the source mechanisms all interacting with each other to create a single, combined expression at the surface is always a challenge.

This introduction is designed to present concepts common to the rest of the dissertation and to motivate the projects carried out over the course of my PhD. The following chapters then present progress on two different scientific fronts: one improving the way data from large-scale geodetic observing networks can be analyzed (Chapter 2), and one developing a new approach to modeling earthquake cycles in a physically meaningful way (Chapters 3–5). The conclusions in Chapter 6 will tie

the individual findings of this dissertation back together and provide an outlook on possible future work.

1.2 Subduction Zone Models

Subduction zones occur where one tectonic plate slides (*subducts*) underneath another one (called the *overriding* plate). All five of the world's largest earthquakes since 1900 (i.e., moment magnitude $M_w \geq 9.0$) have occurred in subduction zones, namely the 1952 M_w 9.0 Kamchatka, the 2004 M_w 9.1 Sumatra, the 2011 M_w 9.1 Tohoku-oki, the 1964 M_w 9.2 Alaska, and the 1960 M_w 9.5 Chile earthquakes (U. S. Geological Survey, 2017). Because of the convergent thrust motion between the plates in subduction zones, and their ability to host the largest earthquakes ever recorded, these regions are also referred to as *megathrusts*. In most cases, the subducting plate is an oceanic crust, and the overriding plate is continental crust. Significant exceptions to this are, e.g., where the Indian subcontinent subducts under Asia along the Himalayan fault, or where the African plate underthrusts the Aegean Sea.

An important concept throughout this dissertation is that of the existence of earthquake *cycles*. From observations all over the world, including at subduction zones, we know that earthquakes on plate interfaces (i.e., the boundary between two plates) happen repeatedly, although not regular enough in space or time to characterize anything more than an average recurrence time (e.g., Griffin et al., 2020; Philiposian and Meltzner, 2020). The quasi-cyclic nature of earthquakes is reasonably explained by the *elastic rebound* theory. This concept boils down earthquake recurrence to a simple sawtooth-like stress pattern: between ruptures, two tectonic plates are stuck ("locked" or "fully coupled") but are getting steadily more "loaded" due to large-scale tectonic plate motion. Once the accumulated load reaches a critical value, an earthquake occurs, releasing all the pent-up stress, and slides the plates by the same amount as they would have if they were never stuck in the first place (i.e., the earthquake releases the accumulated *slip deficit*). The process then repeats, leading to the definition of an earthquake cycle, described by a *coseismic* phase (the earthquake itself) and the *interseismic* phase (the time between earthquakes). To specify the periods right after or right before an earthquake, geophysicists use the term *post-* or *preseismic*, respectively.

Models of increasing complexity explaining the relative motion in a subduction zone are shown in Fig. 1.1. The *backslip* model (BSM) proposed by Savage (1983) as-

sumes that all sliding between two plates happens only in the area that hosts the earthquakes (Fig. 1.1, left), and that the overriding plate experiences negligible deformation over the seismic cycle. During a rupture, the relative motion below the interface is downward and towards the land (as the incoming plate subducts), meaning that the relative motion above the interface is upward and towards the sea. Between ruptures, accordingly, the directions are reversed in order to have net-zero deformation after an entire cycle. The backslip model is still popular today, since it predicts the general motion during all phases of the earthquake cycle using minimal model complexity. It does, however, seem counterintuitive, as we ascribe interseismic normal motion (opposite to the sense of thrusting) in an area of convergence that should only slip during earthquakes. The reason the model still performs so well is that most observations at the time (and still today, in most areas), are not sensitive enough to decipher whether the apparent “backslip” motion is indeed coming from the earthquake source region, or somewhere else close by.

The density, quality, and length of the record of surface observations has steadily improved over the last decades (Blewitt et al., 2018), to the point where we are increasingly sensitive to the location of slip in our model, and where we desire a model that is not only accurate for one recurrence time interval, but multiple. The latter is an issue since in the backslip model, the incoming plate never actually subducts, it just goes back and forth in the shallow region. Kanda and Simons (2010) therefore proposed the Elastic Subducting Plate Model (ESPM, Fig. 1.1, center), adding a second interface just below the first one. This formulation eliminates the need for the backslip to occur on the upper interface, as now the entire lower interface as well as the down-dip extension of the upper interface always slide (*creep*) at the plate convergence rate. The locked area in the shallow, upper interface only slides during an earthquake.

A further continuation of both the BSM and the ESPM is to transition to a three-dimensional world, where locked regions (“asperities”) are interspersed on an otherwise steadily slipping surface. Fig. 1.1 (right) shows this concept for the ESPM case. The asperities are regions on the plate interface where slip only occurs during earthquakes. The degree to which the area between asperities is slipping or not during the interseismic phase then depends on the *stress state* of the fault. The stress state, in turn, depends on the proximity to a fully-locked asperity, and the mechanical response behavior of the fault material to stress as defined by *rheological models*, to which we come back in Section 1.4. For the rest of this dissertation, we will only

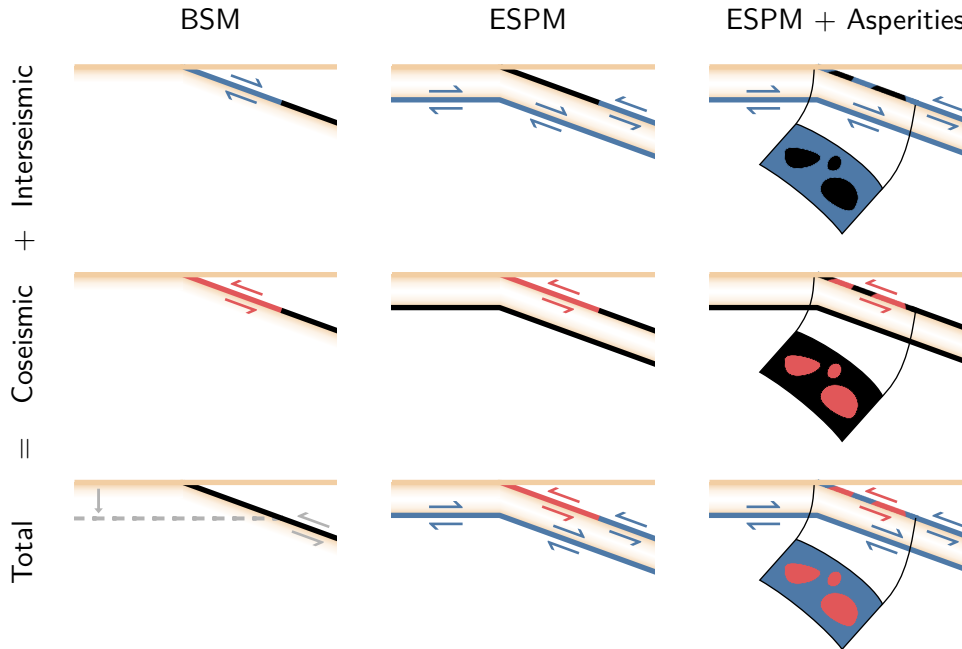


Figure 1.1: Three simplified models of subduction zone plate motion, described in Section 1.2. From left to right: Backslip model (BSM), Elastic Subducting Plate Model (ESPM), and the ESPM with added asperities in 3D. From top to bottom: the interseismic, coseismic, and cumulative relative motion after an entire earthquake cycle, respectively. Light brown lines denote the free surface (interface between a tectonic plate and the ocean or atmosphere). Arrows at an interface show the direction of the relative motion of the two sides. Black lines and regions denote no relative motion, i.e., regions where slip deficit can accumulate in the interseismic period. Blue and red lines and regions denote slow creep and fast earthquake slip, respectively. Over longer timescales, the BSM includes steady-state subduction, which would result in a net uplift of the overriding plate if modeled by block motion (gray dashed line and arrows).

use the ESPM and its asperity extension.

1.3 Observing Surface Deformation

The types of data that can be used for the study of subduction zones are far-ranging. For example, seismographs detect and locate earthquakes, which in turn directly define the plate interface assuming that ruptures occur on or near this interface; research vessels can drill into the seafloor, and sometimes even into the interface itself at shallow depths, to study the materials at the interface; and geochemists can analyze volcanic material sourced from the subducting plate (e.g., Bürgmann and Dresen, 2008; Lin et al., 2013; Uchida and Bürgmann, 2021). To study the large-scale, subduction-cycle-style behavior of megathrusts hundreds of kilometers long, however, arguably the best choice is to monitor the motion (or displacement) of the

surface of the plates.

Displacement of the surface and associated deformation (since tectonic plates are not fully rigid) can be monitored by a variety of techniques, including leveling surveys, radar satellites, or using a Global Navigation Satellite System (GNSS) like the United States' Global Positioning System (GPS) or Europe's Galileo (e.g., Herring et al., 2016). For the purposes of earthquake cycle modeling, networks of GNSS stations dispersed on land provide the best combination of spatial density, temporal resolution, and length of the data record (e.g., Simons et al., 2011; Perfettini and Avouac, 2014; Loveless and Meade, 2016). These stations are directly tied to stable bedrock, and record their GNSS-derived location at daily (or shorter) timescales. Together, they provide a joint view of the motion of the covered area over time. Geodetic networks can also be readily expanded onto the seafloor using acoustic or pressure-based techniques, e.g., offshore Japan, delivering data from right above potential earthquake source regions (e.g., Iinuma et al., 2016; Yokota et al., 2018).

Figs. 1.2 and 1.3 show the large-scale motion during the interseismic and coseismic phases for the representative example of Northern Japan, respectively. There, the Pacific plate (also defined using GNSS observations, Bird, 2003; Altamimi et al., 2017), coming in approximately from the East with a velocity of around 80 mm/a, subducts underneath the Eurasian plate, creating the Northern Japanese megathrust (host of the 2011 M_w 9.1 earthquake). During the interseismic period (Fig. 1.2), the Pacific plate collides with the Eurasian plate, leading to westward motion on land. This motion is reversed during a subduction zone earthquake (Fig. 1.3), where the land rebounds towards the east. The variability of the coseismic motion along-strike (in this case, North to South) is a direct consequence of the limited area that ruptured (i.e., the size of the asperity). Finally, Fig. 1.4 shows the postseismic motion typical of large earthquakes, where the plate interface continues to creep steadily (i.e., independent of aftershocks) over months to years in the same direction as the main earthquake itself.

To obtain velocity fields such as the one in Fig. 1.2, some data processing is necessary. As alluded to in Section 1.1, the part of the surface observations related to the subduction zone is being overlain by other processes. For example, the hydrological cycle leads to a downward motion during wet periods of the year, as water masses weigh down on the surface, and an upward motion during the dry periods, when the continental crust rebounds. In areas of volcanic activity, magma can rise into or flow out of the crust, leading to concentric expansion or compression. Small

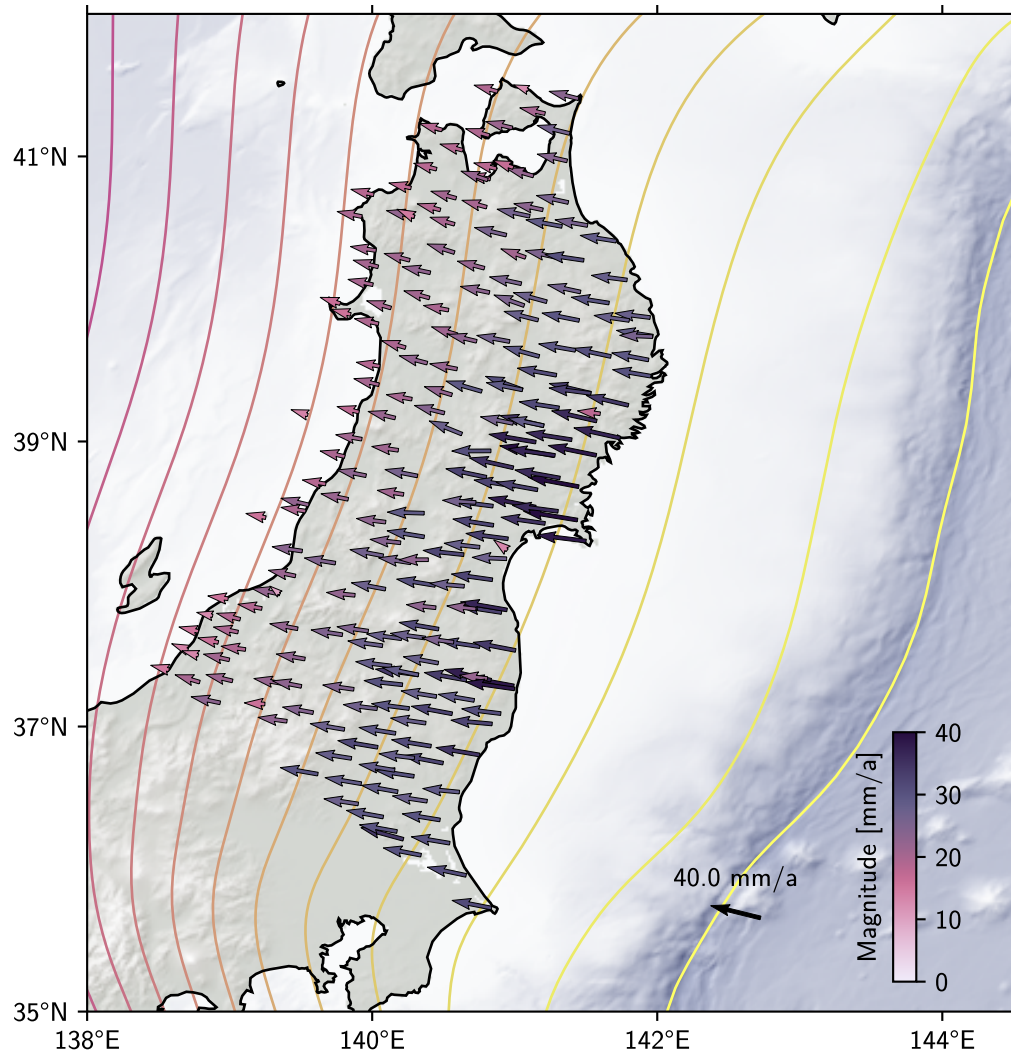


Figure 1.2: Interseismic velocity field for Northern Honshu Island, Japan, relative to the Eurasian tectonic plate. Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). Yellow to pink lines show the plate interface at 20 km depth intervals (Hayes et al., 2018), including the plate boundary at the trench (Bird, 2003). Arrows denote the horizontal motion of analyzed GNSS stations, colored by the horizontal velocity magnitude.

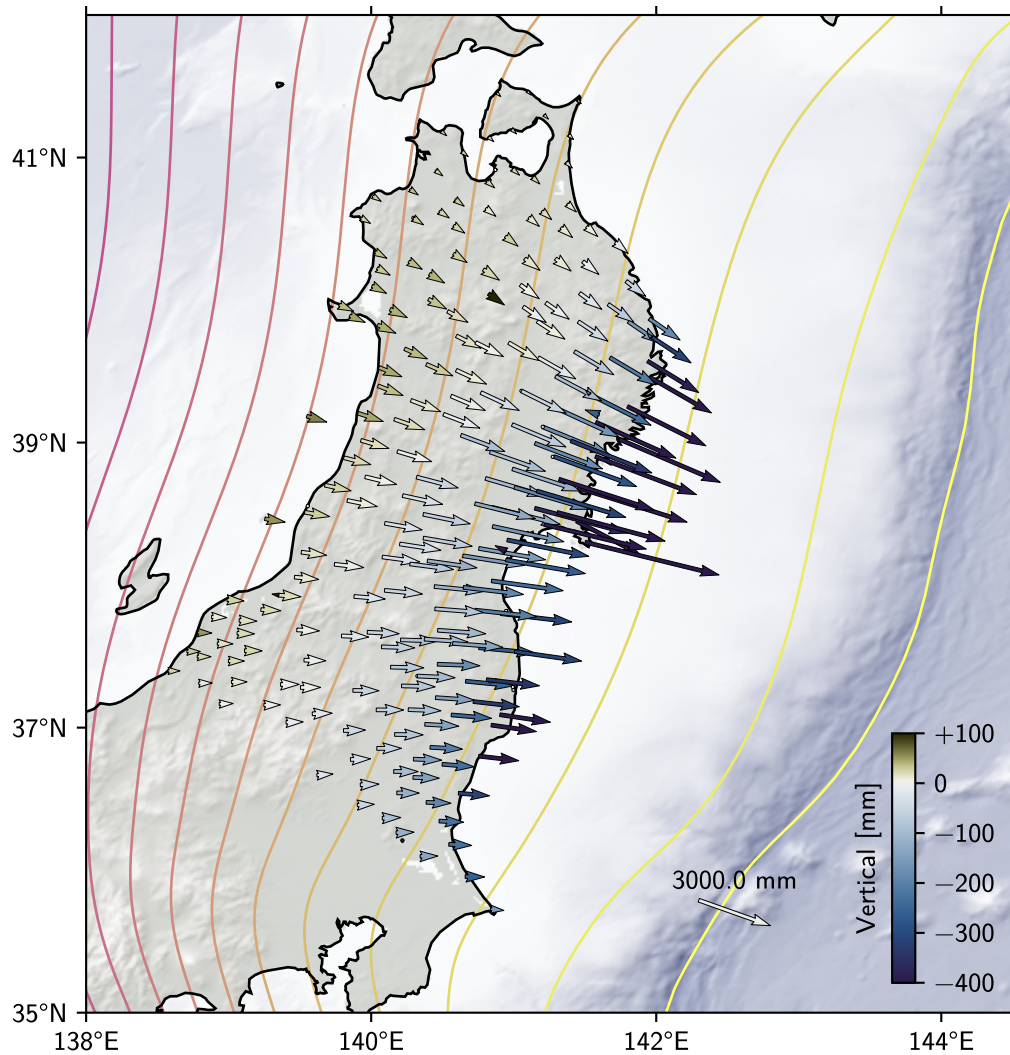


Figure 1.3: Northern Honshu coseismic displacements due to the 2011 M_w 9.1 Tohoku earthquake. Arrows denote the horizontal component, and their color the vertical component. Additional map elements as in Fig. 1.2.

earthquakes outside the subduction zone can have regional effects in all directions, depending on the location, size, and type of the earthquake. Furthermore, the Pacific and Eurasian plates are not just moving relative to each other, but also in an absolute way on the surface on the Earth, leading to continental-scale translation and rotation. Lastly, non-physical effects such as changes in the instrumentation can also lead to apparent, but not real, motion of an observing station.

When only a handful of stations are available, or there are few confounding effects present, this analysis can easily be done manually. As networks of geodetic stations get denser with time, and their data record longer, however, semi-automated tools that can efficiently process large datasets become more important. At the same time,

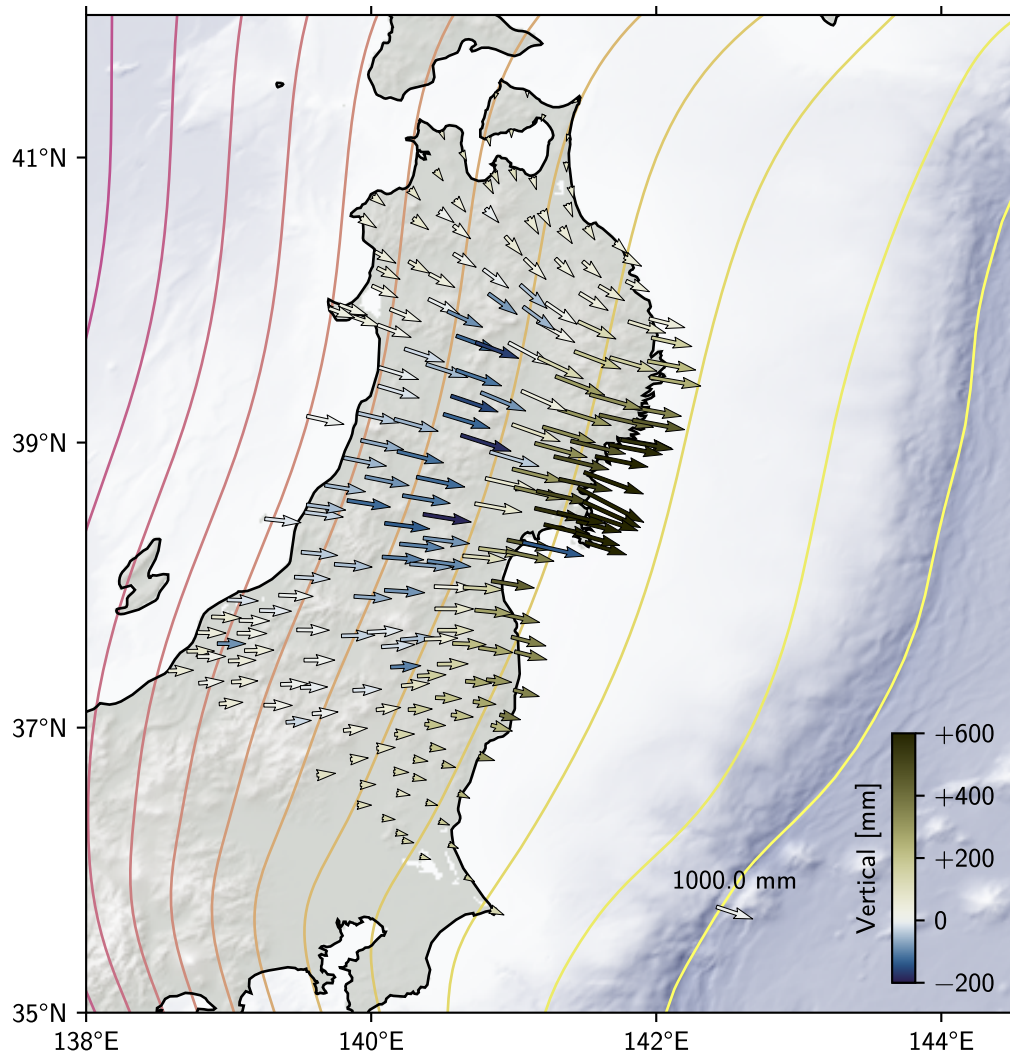


Figure 1.4: Northern Honshu postseismic displacements due to the 2011 M_w 9.1 Tohoku-oki earthquake, cumulative until 2024. Arrows denote the horizontal component, and their color the vertical component. Additional map elements as in Fig. 1.2.

the suite of methods used to process timeseries has grown from simple linear and sinusoidal fits to longterm (*secular*) plate motion and *seasonal* variations, respectively, to include *transient* (i.e., bounded in time) modeling using principal component analysis, multi-timescale logarithmic functional forms, and spline-based dictionaries, to name a few. The combination of large datasets with a high processing standard desired by the scientific community without any general software available (as of 2020) to deal with “solved problems” of timeseries analysis furthermore creates an unnecessary barrier of entry for new researchers.

In this context, Chapter 2 presents the DISSTANS Python software package that aims to address the diversified needs of the community. It includes both standard and

state-of-the-art processing tools in a framework that is easy to learn and to extend. DISSTANS furthermore features parallelization for the most demanding tasks to decrease analysis runtimes for large networks significantly and an expansive set of visualization routines.

1.4 Slip Deficit and Rheological Models

Just as the surface patterns of coseismic displacement are an indicator of the finite extent of the rupture source at depth, one can use the non-uniformity in the inter-seismic phase to determine the spatially-variable sliding rate on the interface. By comparing the estimated slip rate with the expected one from large-scale plate motion, the degree of *coupling* can be calculated. In our simplified model of earthquake cycles, areas that are fully uncoupled (and therefore sliding, coupling coefficient of 0) are not accumulating any stress, and as a consequence are not expected to host large earthquakes. On the other hand, fully coupled regions (asperities, coupling coefficient of 1) are assumed to host large subduction zone earthquakes once they reach a critical state.

Fig. 1.5 shows a map of inferred coupling of the Northern Japanese megathrust based on the data shown in Section 1.2 and following the inversion scheme of Lindsey et al. (2021). Contrary to Section 1.3 and Fig. 1.2, however, we have here assumed that the overriding plate is the Okhotsk plate, rather than the Eurasian one, to be consistent with other published coupling maps, including Lindsey et al. (2021). The coupling estimate is based on the sliding rate of the interface, and is therefore showing the *kinematic* coupling. This definition hints at the fact that we are only characterizing the current *velocity* of sliding — even though the metric that directly affects our earthquake size estimate is how much slip has occurred since the last rupture (i.e., the time integral of the slip velocity). It may be an area that has always slid at the same speed, in which case the kinematic coupling is a good proxy for the *slip deficit*. However, if an area slipped significantly during any time in the past (e.g., due to induced afterslip after an earthquake, or slow slip events), the slip deficit would be much smaller. In this case, the area in question would not be currently slipping because it is in the *stress shadow* of a nearby asperity, keeping it stuck.

Another limitation of the kinematic view of the evolution of plate interfaces is that it cannot yield predictions of future slip or slip rates. To forecast the sliding on plate interfaces, we need to invoke *rheological models*, i.e., physical laws that describe the mechanical responses of materials to stress changes. For the purposes of

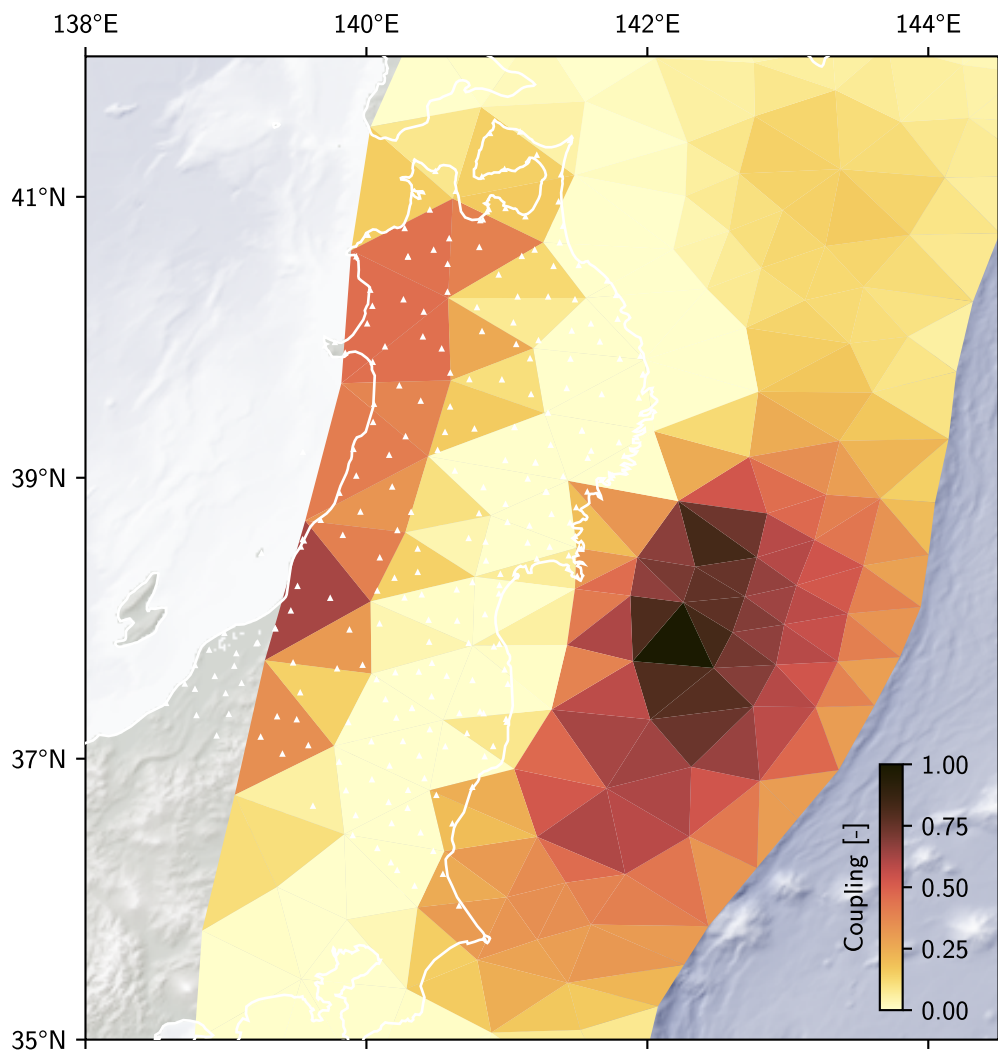


Figure 1.5: Kinematic coupling coefficient on the Northern Japan plate interface between 35°N and 42°N (colored mesh). Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). White triangles denote the location of the GNSS stations whose velocities were used in the inversion.

this dissertation, rheological models (also called *constitutive laws*) simply define the (potentially nonlinear) relationship between stress and strain (or equivalently, stress and fault slip rate). At plate boundaries, strain in the volume at and around the plate interface can be thought of as relative shearing or sliding between two plates, allowing us to directly tie slip on the fault to stress. Therefore, if we assign a rheological model to a plate interface, and we know the stress it is exposed to (i.e., the pushing of the incoming plate), then we can calculate how the slip on the fault will evolve over time.

Rheological models play an important role in our scientific understanding of subduction zones. For seismic hazard studies, a plate boundary with a “stronger” type of rock (i.e., one that deforms less when subjected to a constant force) could host larger earthquakes than one with softer material, since the fault could accumulate stress (and slip) for longer before releasing it in an earthquake. For geomorphologists, a weaker rock will be less able to form the mountain ranges that we see behind subduction zones such as the Andean megathrust. And on a grand scale, even the existence of subduction zones as a phenomenon in the first place requires specific stiffness contrasts (Stern, 2002). Simply drilling into the crust and studying the dug-up material in laboratories is, however, not sufficient, since upscaling values derived from samples that fit into a building to the hundreds of kilometers that span a subduction zone is not straightforward (Handy, 1994; Yamashita et al., 2015; Fagereng and Beall, 2021; Bercovici et al., 2023). Indeed, what we experience at the surface of the Earth is more likely to be driven by an average, *effective* rheology of the plate interface. Geodetic modeling studies are therefore an important complement to laboratory work when studying the rheological properties of plate interfaces. By using observations from large earthquakes as a natural experiment, many studies have been able to fit parameters of constitutive laws (e.g., for Japan, Sun et al., 2014; Hu et al., 2016; Freed et al., 2017; Agata et al., 2019; Muto et al., 2019; Fukuda and Johnson, 2021) However, current modeling approaches have found ambiguous success since many parameter combinations and even *types* of rheological models explain the available data similarly well, a problem referred to as *non-uniqueness*.

Chapters 3–5 present a framework that aims to provide a new perspective on the non-uniqueness of recovered subduction zone rheological profiles. Based on prior work by Kanda et al. (2013), our setup radically simplifies the model complexity compared to the abovementioned studies, introduces new model constraints based on ancillary datasets, and uses more observations in the inversion process. As such, it

is the first inversion study for rheological parameters of subduction zones using both pre- and postseismic surface displacement timeseries — timeseries post-processed with the DISSTANS software presented in Chapter 2. By simulating entire earthquake cycles and assuming a rheological model, the framework furthermore tracks the stress state of the subduction interface, avoiding the kinematic coupling caveats when producing slip deficit maps. Insights from this new class of cycle-based models may be of great benefit to the study of subduction zones and the assessment of megathrust seismic hazard.

*Chapter 2***DECOMPOSITION AND INFERENCE OF SOURCES
THROUGH SPATIOTEMPORAL ANALYSIS OF NETWORK
SIGNALS: THE DISSTANS PYTHON PACKAGE**

Köhne, T., Riel, B., and Simons, M. (2023). “Decomposition and Inference of Sources through Spatiotemporal Analysis of Network Signals: The DISSTANS Python Package.” In: *Computers & Geosciences* 170, p. 105247. DOI: 10.1016/j.cageo.2022.105247.

2.1 Introduction

Networks of Global Navigation Satellite System (GNSS) stations enable the direct observation of surface displacement down to millimeter accuracy (e.g., Blewitt, 2015). Originally using only the Global Positioning System (GPS) and consisting of only a handful of stations, modern quasi-permanent deployments sometimes incorporate more than 1,000 receivers and take advantage of other GNSS constellations such as the European Galileo or Chinese BeiDou systems. Analyzing network position timeseries requires awareness of the many processes that affect the observations, both desired and confounding, and an ability to distinguish between them. While dominant constituents like the secular motion of a particular station can usually be inferred by simple linear regression, quantifying less prominent constituents (e.g., displacements due to low-magnitude slow fault slip events or small-volume magma chamber pressurization) requires a better understanding of the contributing processes.

Here, we present the Decomposition and Inference of Sources through Spatiotemporal Analysis of Network Signals (DISSTANS) Python package to facilitate the temporal and spatial decomposition of GNSS timeseries. The code is written in a generic, fully object-oriented fashion with minimal assumptions as to study location, data units, and sampling frequency. Different data loading methods are implemented that interface with common existing timeseries file formats, but are also easily adapted to new formats. All downstream processing is independent of the original format and origin. To make the code as usable and accessible as possible, it is open-source and extensively commented. The repository includes tagged versions, verbose commit messages, and full documentation. The documentation

features tutorials based on synthetic and real timeseries data, a subset of which are presented here. DISSTANS already contains many common processing workflows. These workflows are usable with just a few lines of code, and more are in the planned development roadmap. DISSTANS is parallelized for the most demanding tasks — most notably the model fitting component. We also provide extensive plotting options and graphical user interfaces, simplifying interactions with the data.

Section 2.2 of this report introduces some key structural decisions and presents a brief literature review of previous work, placing this study in the broader scientific context. Section 2.3 provides an overview of the code design, with Appendix A detailing the lower-level implementation. To validate our processing, Section 2.4 contains the analysis of a synthetic network of GNSS stations, as well as results from a real-world application using data from the Long Valley Caldera region, California, USA. Section 2.5 discusses key design choices. Finally, we end in Section 2.6 with a brief summary and some possible future avenues for extensions to DISSTANS.

2.2 Background

The list of scientific questions that can be addressed with GNSS networks is long, and the list of approaches that can be used is even longer (e.g., Herring et al., 2016; Blewitt et al., 2018; Herring et al., 2018; Bertiger et al., 2020). For studies of plate motion, surface deformation, and related fields, the key data are displacement timeseries, i.e., the relative movement over time of a receiver with respect to a defined reference frame. To obtain these timeseries, processing centers start from raw receiver observables (time, pseudoranges, and phases) and take into account a large number of physical processes (e.g., tropospheric and ionospheric travel time delays, gravitational effects, relativistic effects) to produce the best estimate of true receiver position for any given sampling interval (e.g., daily or hourly); see, for example, Misra and Enge (2010) or Blewitt (2015).

With these displacement timeseries, we can now interrogate the timeseries: Is the entire signal explained by rigid plate motion (e.g., Altamimi et al., 2017)? What are the causes for shortterm or longterm transients (e.g., Houston et al., 2011)? How can we use inter-, co- and postseismic station velocities to constrain fault locking (e.g., Meade and Hager, 2005)? Similarly, one might also want to identify and characterize noise processes (e.g., power-law noise, Langbein, 2020). All of these questions, however, require the *decomposition* of the timeseries into components that are the direct effect of specific physical processes (e.g., hydrological seasonal loading,

earthquake offsets and transients, plate motion), and a residual component which is the result of noise processes, processing artifacts, and imperfect modeling.

In this study, we focus on this intermediate step, and refer to it as simply *timeseries decomposition*. Therefore, we will refer to the displacement timeseries as produced by GNSS network processing centers as the *raw* or *input* timeseries, and the decomposition process will aim to isolate timeseries *constituents*.

2.2.1 Approaches to Timeseries Decomposition

We categorize timeseries decomposition tools using three main criteria.

Process-aware vs. process-agnostic

This first criterion aims to distinguish approaches that either make a priori assumptions about the physical processes affecting the data (expecting a certain structure in the data), or alternatively, assume the least possible. For example, fitting a model containing a complete set of basis functions to a timeseries is, in its most generic form, process-agnostic (e.g., Riel et al., 2014), but fitting a logarithmic decay function to a postseismic transient effectively assumes a specific tectonic process (e.g., Hsu et al., 2009).

Process-agnostic approaches will usually achieve the “best” fit to the observations — at least as measured by the magnitude of the residuals, since that is the principle optimization criterion for such methods. However, over-reliance on the data and its residuals makes these methods susceptible to “overfitting”; i.e., interpreting noise as signal. Process-agnostic methods also have difficulties determining trade-offs between different source processes, for example in the case of Principal Component Analysis (PCA) when signals manifest in multiple principal components, or a single principal component mixes different signals. In contrast, process-aware approaches might ignore parts of the observation if they either (a) do not have an appropriate way of describing the observation (e.g., an unexpected transient), or (b) try to fit a signal with an inappropriate model (e.g., mapping postseismic deformation into the coseismic one); as these approaches naturally prefer a decomposition that follows the assumed underlying functional forms.

Parametric vs. non-parametric

This second criterion assesses whether one estimates parameters (coefficients) for predetermined models to decompose the timeseries. The models can be as complex

as desired (high dimensionality, non-linear). Examples for non-parametric decompositions are linear time-invariant filters used in signal processing (e.g., bandpass or lowpass filters) or basis rejections like Principal Component Analysis (PCA), Independent Component Analysis (ICA), or Singular Spectrum Analysis (SSA).

Note that this criterion ignores the impact of hyperparameters (e.g., regularization penalties, frequency windows). With non-parametric approaches, the assumptions and hyperparameters are minimal compared to model-based methods, thereby simplifying the problem setup immensely. Furthermore, reducing the influence of hyperparameters translates into a reduction of possible sources of errors. On the other hand, parametric approaches enable a straightforward implementation of the formal covariances between model parameters, and by extension, uncertainties in the predicted timeseries. These approaches can also deal naturally with data gaps (i.e., without the need for data imputation). Crucially, a parametric approach is necessary for process-aware studies, because non-parametric approaches have no inherent knowledge about how to group different source processes into components (see above).

Station- vs. network-level

An additional criterion acknowledges the role that spatial information can play in the analysis process. For example, if the same models are fit to every timeseries in a network, regardless of where the stations are located, then the decomposition code is not aware of the spatial context. These local, station-level solutions are therefore independent from another. If one recognizes, however, that geophysical signals usually have a spatially coherent signature (assuming sufficiently dense networks), then we can and should incorporate that understanding. For example, PCA makes use of the fact that all stations in the network can potentially see the same source signal (even though the network geometry is neglected). Taking advantage of potential spatial structure can be beneficial, although the complexity of the resulting code and additional computational costs are not negligible.

2.2.2 Review of Existing Tools

Considering the diversity of possible approaches, the selection of a certain approach (or the design of a hybrid approach) depends on one's goals and the available data. Additional factors include the ease of software implementation, or possibilities to extend the methods to include ancillary datasets (e.g., rainfall, earthquake catalogs,

atmospheric pressure). We review selected published work in the field of timeseries decomposition in the context of process-agnostic vs. process-aware, parametric vs. non-parametric, and degree of spatial awareness.

Before high-quality station timeseries became ubiquitous, the QOCA software (Dong et al., 1998) could be used to combine “quasi-observations” (lightly-processed input data from GPS, Electronic Distance Measurements, Satellite Laser Ranging, or Very Long Baseline Interferometry) using a Kalman filter approach. QOCA includes the module `analyze_tseri` to estimate linear, episodic, and stochastic motion of the different stations individually in a least-squares-based, process-aware, and parametric framework.

With an increasing number of GNSS stations, more GNSS constellations, and more precise understanding of the physical processes affecting GNSS positioning solutions, GNSS networks became common for monitoring surface deformation. Today, the analysis tools developed to produce GNSS displacement timeseries routinely also include simple timeseries decomposition functionality. For example, the current iterations of JPL’s GipsyX/RTGx (Bertiger et al., 2020) and MIT’s GAMIT/GLOBK (Herring et al., 2018) software both contain methods to estimate position, velocity, seasonal variations, offsets¹, and postseismic deformations². These Kalman-filter-based methods are parametric, process-aware, but in contrast to QOCA, not spatially aware.

For regions where complex geophysical processes are at play (such as near a volcano or in subduction zones), more complex analysis is necessary to distinguish between different processes. A common example is the impact of an unmodeled transient period on the estimated secular plate velocity. In the following, we present a (non-exhaustive, unordered) small selection of tools that start from raw GNSS displacement timeseries to analyze stations exhibiting more complex behavior.

The Network Inversion Filter (NIF), first proposed by Segall and Matthews (1997) and subsequently expanded upon by a variety of studies (e.g., McGuire and Segall, 2003; Bekaert et al., 2016), estimates slip rates on predetermined fault structures from (GNSS or other) observations using a Kalman filter. It is therefore process-aware, and because slip on the modeled faults affect multiple stations, which are jointly used to estimate the slip coefficients, it is also spatially aware. The NIF estimates slip and therefore transient displacement constituents non-parametrically, but the hyperparameters specifying the fault geometry and the characteristics of fault slip in time and space (e.g., smoothness) play an important role.

The Median Interannual Difference Adjusted for Skewness (MIDAS, Blewitt et al., 2016) algorithm explicitly recognizes the importance of unmodeled steps and short-term transient deformation in the raw timeseries. Not being a traditional regression scheme, it uses the median of velocities computed from data pairs separated by one year, providing a degree of insensitivity to offsets, small data gaps, and annual seasonal signals if the timeseries is sufficiently long. This process-aware, station-level method is mostly defined by its hyperparameters, although other parameters such as known maintenance and earthquake offset times are used. It is therefore a powerful, largely automated method to estimate secular plate velocities, that does not attempt to extract non-annual seasonal, transient, or decaying signals. MIDAS is at the core of UNR’s Nevada Geodetic Laboratory openly-accessible global GNSS timeseries repository (Blewitt et al., 2018).

The Señales y Análisis de Ruido Interactivo (Interactive Signal and Noise Analysis, SARI, Santamaría-Gómez, 2019) software performs process-aware, parametric, station-level regression focusing on an interactive user interface. Least-squares or Kalman filtering is used to fit polynomial, sinusoidal, exponential, logarithmic, and step models, allowing for a detailed timeseries decomposition. It also contains useful additional functionality such as automatic discontinuity detection, periodogram visualization, and noise characterization.

The Greedy Automatic Signal Decomposition (GrAtSiD, Bedford and Bevis, 2018) algorithm is an iterative, station-level method that focuses on detecting and modeling transient signals in the timeseries. At each iteration, a least-squares regression is performed that includes a linear trend, sinusoidal oscillations, predefined steps, as well as a selection of sparse, transient functions (“multitransients”). Only multitransients that significantly improve the data fit are kept for the next iteration, until a convergence criterion is reached. GrAtSiD can therefore be classified as a parametric approach, that is partly process-aware (for the non-multitransient parts of the regression) and partly process-agnostic (since the multitransients can have a variety of shapes and are not tied to a particular physical source).

MIDAS, SARI, and GrAtSiD are limited to station-level model fit solutions, and do not incorporate spatial awareness.

An example of a non-parametric, process-agnostic, and spatially-aware method to decompose timeseries is the variational Bayesian Independent Component Analysis (vbICA, Gualandi et al., 2016), a modern iteration of basis reprojection algorithms particularly suitable for GNSS networks. Its key distinction from traditional

PCA/ICA is to recognize that probability density functions for individual components are generally not normally distributed by nature, and alleviates this problem by using mixtures of Gaussians. vbICA therefore allows for a more accurate signal separation, as well as a formal way to incorporate component uncertainties.

Riel et al. (2014) proposed a method that builds on parametric, process-aware regularized regression. Their approach adds a process-agnostic set of B-Spline functions to model transients in a spatially-aware framework. DISSTANS builds on this framework, which we describe in more detail in Sections 2.3.2 and 2.3.3.

Here, we have just described a subset of available tools, focusing on publicly available, complete software packages that provide a reasonable level of portability. There are many other studies that have implemented or adapted codes and methods for specific study regions or purposes; an analysis and comparison of which would be beyond the scope of this work.

2.3 Code Overview

DISSTANS aims to build on advancements and best practices of previous work, combining them into a single package that adheres to standards of free, extensible, shareable, and scalable software. At its core, it models timeseries as the linear combination of constituent functions, and estimates the functions' coefficients using least-squares. DISSTANS also includes a suite of pre- and postprocessing tools. In this section, we present key properties and design choices made in the DISSTANS package (Section 2.3.1), as well as two core functionalities (spline-based transient modeling and spatial regularization, see Sections 2.3.2–2.3.3). More implementation details can be found in Appendix A.

2.3.1 Goals

Commonly-used workflows included. To allow researchers to focus more on science and less on programming, timeseries decomposition software should include easy-to-use versions of well-established timeseries decomposition workflows. Such software can then be used for generic pre- and postprocessing, as well as serve as a base on which new analysis methods can be developed. An additional benefit of comprehensive software is the lowering of the entry barrier for researchers new to the field. DISSTANS therefore provides a vast array of such workflows, ranging from data cleaning methods and PCA/ICA decomposition to simple least-squares fitting with standard models and residual analysis.

Incorporating process knowledge. Where knowledge about physical processes affecting GNSS timeseries is present (e.g., an inflating magma chamber), such information can theoretically improve model fitting. It is therefore desirable for timeseries decomposition methods to both include models that best represent known physical processes, as well as methods that are flexible enough to account for unmodeled, unknown processes. DISSTANS allows for such a distinction by offering a range of process-aware, as well as process-agnostic models (see Section 2.3.2 and Appendix A.2).

Spatial awareness. With GNSS networks becoming more widespread — and more importantly, denser — we should explicitly recognize that nearby stations subject to common geophysical processes may behave similarly. If we only consider each station individually, we may miss the opportunity to identify constituents that manifest themselves around the noise floor. However, if many stations experience the same signal (with different magnitudes), a joint estimation can theoretically enhance our ability to detect them. Such a method would thereby lower the effective signal-to-noise ratio necessary for constituent extraction. DISSTANS allows one to take advantage of the available spatial information by building on the spatiotemporal transient fitting algorithm developed by Riel et al. (2014) (also see Section 2.3.3 and Appendix A.3).

Scalability. In order to scale well with both the number of stations, as well as the length of the observation record, it is useful to parallelize the computationally demanding parts. DISSTANS includes an option to parallelize the station-level, least-squares solutions, as well as the evaluation of the predicted model timeseries including the full model covariance matrix.

Uncertainty estimation. Given the possible complexities of displacement timeseries, a proper interpretation of signal decomposition results can only be made if the trade-offs between and within models and east-north-up components can be quantified. The full, formal model covariances can be estimated and propagated in the DISSTANS workflow.

Step detection. One omnipresent challenge when analyzing timeseries is the detection and subsequent estimation (or equivalently, removal) of steps in the data. Improper step removal can significantly affect secular plate velocities as well as the character of GNSS noise (e.g., Blewitt et al., 2016; Santamaría-Gómez and Ray, 2021), but there is no fully-automated algorithm that would remove the need for manual inspection (e.g., Gazeaux et al., 2013). DISSTANS contains semi-

automated tools that aid modeling all relevant offsets: a step detector (similar to the one in GipsyX, Bertiger et al., 2020, also see Appendix A.6), a visualization GUI to inspect the data (see Appendix A.8), and loading functions for maintenance records in multiple formats. DISSTANS also features both an empirical (following Blewitt et al., 2016) and an elastic-half-space-based method to determine whether or not to allow coseismic offsets to be estimated at any given station and time.

Portability and extendability. As new GNSS networks are built, and output formats of data processing centers change, one must be able to easily incorporate these changes. DISSTANS separates the data loading tasks from all other analysis steps, such that the former can easily be updated without affecting the latter. Furthermore, to enable the development and integration of new approaches, DISSTANS is written as a modular, extendable framework (in contrast to single-use collections of scripts, see Appendices A.1 and A.4).

2.3.2 Spline-Based Transient Modeling

To optionally model transient signals in the displacement timeseries of unknown functional shape, DISSTANS includes spline-based models. *B-splines* in particular are piecewise-polynomial functions that, when constructed in a specific manner, form a full basis for any polynomial function of a given degree over the basis' support. As introduced by Hetland et al. (2012) for geophysical applications, sets of repeated, uniform, integrated B-splines (see Fig. A.2 for a visualization) of various timescales and center times can be used to approximate any given unknown transient signal of similar timescales. The ability to approximate arbitrary functions in a process-agnostic framework makes sets of splines useful for timeseries decomposition where standard functions (polynomials, sinusoids, exponential functions, etc.) cannot capture the full breadth of the observations (e.g., aseismic slow slip or volcanic expansion events). A more detailed mathematical description of the available spline-based models in DISSTANS can be found in Appendix A.2.

2.3.3 Local and Spatial Regularization

Sets (or “dictionaries”) created by shifting and scaling a single uniform B-spline are not linearly independent (see Hetland et al., 2012), and therefore do not form a “proper” basis in the mathematical sense. It follows that any signal decomposition using such sets is non-unique, and thus requires regularizing the solution. The most commonly used regularization is based on the L2 (Euclidean) vector norm $\|\cdot\|_2^2$, promoting solutions with smaller overall magnitudes. However, in the context of fitting

transient signals that may or may not be present in the timeseries, we prefer a regularization scheme that yields sparse solutions, i.e., spline coefficients should be driven to zero if there is no sufficient evidence in the data to warrant usage of any given spline in the overall model fit. L1-norm regularization is such a sparsity-promoting regularization scheme: it penalizes the absolute magnitudes $\|\cdot\|_1$ of the estimated parameters, driving many parameters close to zero. L0-norm regularization goes further by penalizing the *existence* $\|\cdot\|_0$ of a parameter, thereby either driving parameters to zero, or not penalizing a parameter at all (Candès et al., 2008). This type of regularization is therefore more suited for physical processes which occur sporadically, are not ubiquitous, and have an “arbitrary”, but significant, magnitude. All three regularization schemes are implemented in DISSTANS (see Appendix A.3).

Riel et al. (2014) combined the potential of using dictionary of splines with the benefits of L0 regularization. Using the algorithm introduced by Candès et al. (2008), they proposed a method to extend the regularization from a timeseries at a single station (henceforth referred to as *local* L0 regularization) to all the timeseries in a network of stations (*spatial* L0 regularization). Their approach yields spline-based fits whose estimated model coefficients are sparse in time (i.e., for a single timeseries at one station) *and* space: transient signals common to multiple stations are decomposed using the same spline functions. An additional benefit of a spatially-coherent set of splines is that it is harder for the solver to fit local noise processes with splines that would only be relevant at isolated stations and times. DISSTANS builds on the method of Riel et al. (2014) (for which the relevant source code was never published), extending it in various ways (most notably, adding parallelization and improving the numerical stability). More details on the implementation of the spatial L0 regularization can be found in Appendix A.3.

2.4 Validation

We present two validation datasets and results. The first, in Section 2.4.1, is a synthetic dataset of 16 stations exhibiting some commonly seen patterns in GNSS network timeseries. Using this synthetic network, we demonstrate key capabilities of this code in estimating spatially-coherent complex signals, all while being able to compare fitted models to the true underlying timeseries. The second dataset, in Section 2.4.2, is a collection of GNSS stations in the Long Valley Caldera region in California, USA. Here, the main goal is to recover the transient caldera inflation signal, and discuss some subtleties in the analysis when dealing with imperfect, real-world data.




Figure 2.1: Map view of the synthetic network.

2.4.1 Synthetic Dataset

The code for this analysis, as well as additional discussion, can be found in Tutorial 3 of the online documentation.

A key feature of DISSTANS is its ability to use spatial coherence as an additional source of information and constraint. In general, signals like earthquakes, slow fault slip events, or seasonal loading are spatially correlated, as the processes affecting each station have the same underlying sources. By using this knowledge in combination with the enforcement of sparsity, we ensure that the estimated models are consistent between stations. Processes that only affect a single station are considered noise for the purposes of this study (e.g., antenna maintenance or strongly localized displacements).

Setup

The synthetic dataset is comprised of 16 stations randomly positioned on an elongated, rectangular grid (see Fig. 2.1). Each two-component station is affected by a secular, linear trend, one annual seasonal signal, an earthquake (with both co- and postseismic components), two shortterm slow slip events, one longterm transient, common mode error, and measurement error (correlated between the components). The linear trend, coseismic and postseismic constituents are all equal in direction and magnitude, whereas the seasonal constituent is random at each station. The three transients are all equal in onset time, duration, and direction, but differ in magnitude. Furthermore, one station (“Cylon”) experiences significant power-law noise, and a different station is affected by an unmodeled maintenance step. Both signals represent site-specific noise processes that the spatial coherence constraint aims to suppress. Lastly, the amplitudes of the three transients decrease exponentially towards the east.

The analysis follows a simplified version of the example workflow presented in Appendix A.7. Because the data is synthetic, no quality metrics need to be applied, nor is step detection necessary. We add the following constituents to our inversion:




Figure 2.2: Overall model fit to the data at station Jeckle (see Fig. 2.1), including the decomposition into the primary constituents (East component only).

polynomial, sinusoidal, step, and logarithmic (to recover everything except the transients; unregularized); and a set of splines. To recover the transient episodes, the spatially-regularized splines contain timescales between a month and multiple years, amounting to hundreds of individual splines (see Fig. B.4). The fitting converges smoothly onto the final solution (see Fig. B.1). In the following, we compare the results obtained using local and spatial L0 regularization to highlight the benefits of promoting spatial coherence.

Results

Fig. 2.2 shows the East component of a representative station. The inferred model fits the synthetic data well. We find a small tradeoff between the secular and transient constituents, although we note that in real world applications, such a conclusion is frequently difficult. (A visualization of the full model parameter correlation matrix can be found in Fig. B.2.) Fig. B.5 and Table B.1 present comparison results of our spatial L0 solution with the L0 solution without spatial regularization (see below) and other commonly used methods (simple least-squares, MIDAS), showing that the spatially-aware L0 solution clearly outperforms other methods.

Fig. 2.3 shows the improvement from local to spatial L0 regularization in map view for all stations: the transient components are smoother (therefore fitting less noise) and more closely follow the true signal (shown in the background). Importantly, the homogenous displacement field is obtained without degrading the fit to the data (compare Fig. B.3). This is enabled by the spatial solver's identification of the set

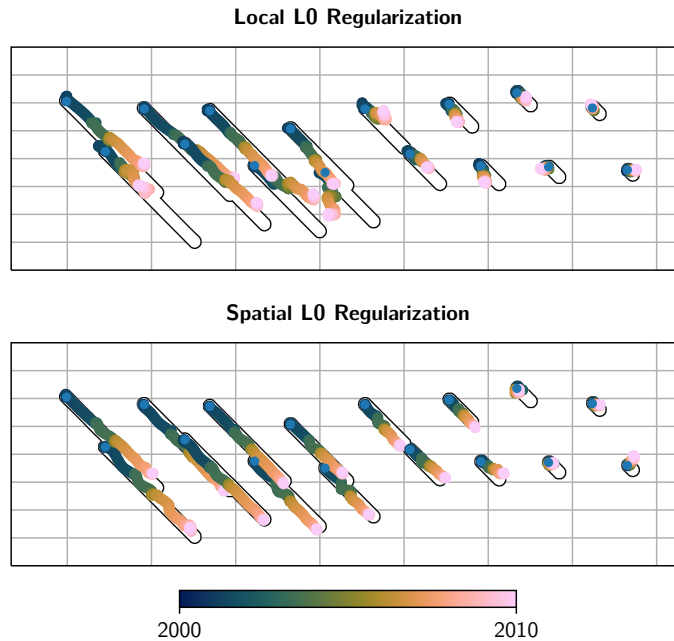


Figure 2.3: Map view of the transient motion of the synthetic network over the entire timespan. The top panel shows the result without spatial regularization, the bottom one with. Markers correspond to the position of a station relative to its initial position, with colors corresponding to time. The white background curves with black outlines are the true synthetic transient, which is clearly better matched by the solution incorporating spatial awareness.

of splines that best describes the transient signal common to all stations (compare Fig. B.4), and the better recovery of the secular velocity. Section B.2 explores the dependence of the model error on the number of stations for a different synthetic network, with Fig. B.6 further validating our claim that incorporating information from nearby stations improves the quality of the resulting model fit.

2.4.2 Long Valley Caldera

The code for this analysis, as well as additional discussion, can be found in Example 1 of the online documentation.

To demonstrate DISSTANS with real data, we consider timeseries from the Long Valley Caldera (LVC) region in California, USA. Because of the geophysical interest into the magmatic, seismic, and hydrological processes at work there, the LVC has been monitored by an ever-expanding network of GNSS stations since the late 1990s (e.g., Ji et al., 2013; Montgomery-Brown et al., 2015; Silverii et al., 2020). The displacement timeseries are complemented by detailed maintenance and seismic catalogues, which are crucial for determining the best set of steps to include in

the fitting process. In this example, the goals are threefold: (1) to illustrate the example workflow proposed in Section A.7, (2) to present the best-fit transient model to the periods of unrest in the Long Valley Caldera, and (3) to showcase the importance of allowing the seasonal signal models to vary in amplitude over time. Any in-depth physical modeling of the extracted constituents is beyond the scope of this study.

Setup

The data and corresponding maintenance and seismic events catalog are downloaded with DISSTANS-included tools from the GNSS timeseries repository maintained by the University of Nevada at Reno’s Nevada Geodetic Laboratory (Blewitt et al., 2018). Only stations with a reliability of over 50% (i.e., observations on more than half of the days the station was active) and an observation record at least one year long are considered, and outliers in each timeseries (more than 10 standard deviations away from the median) as well as the common mode error are removed. With help of the available maintenance catalog, we iteratively identify steps in the data. This process is aided by the step detector and visualization routines included in DISSTANS. For the final fit, we use following models: polynomial, sinusoidal, and steps (to recover everything except the transients; unregularized); a set of splines (to recover the transient episodes, containing timescales between months and multiple years, hundreds of individual splines; spatially regularized); and a varying-amplitude sinusoid (for deviations from the nominal, unregularized seasonal signal; locally regularized).

Transient Signals

The timespan between 2012 and 2015 (approximately) is dominated by a significant expansion of the caldera’s dome, as observed by both the GNSS network and Interferometric Synthetic Aperture Radar (InSAR) timeseries (Montgomery-Brown et al., 2015; Silverii et al., 2020). Fig. 2.4 shows the horizontal component estimated by DISSTANS in map view: the radial extension of the network from the center of the dome is clearly visible.

Fig. 2.5 shows the extracted transient constituent of selected stations in and around the resurgent dome for the entirety of the available data. Two periods of significant expansion can clearly be distinguished: around 2002–2003, and from 2011–2021. A smaller period of unrest can be seen throughout the network between approximately




Figure 2.4: Modeled horizontal transient displacements of selected stations in the Long Valley Caldera region during the period between 2012 and 2015. The traces begin at the nominal location of each station, with the color of the markers corresponding to the time. Background satellite imagery by Earthstar Geographics & Esri.

2008–2010, and the station CASA allows us to see a period of extremely rapid expansion around 1998. These results are comparable to Silverii et al. (2020, Fig. S3a), where transients were recovered using non-parametric multiyear filters, even though the directions of maximum displacements are different. Crucially, however, we did not enforce the secular long-term motion to be zero during a specific timespan. As a result, many stations appear to never reach a “steady-state” matching the general plate motion, because the transient motion, even when regularized, is dominant for large parts of each timeseries. (A priori removal of a secular trend can easily be done with DISSTANS, if desired.)

Estimation of the transient signal directly affects the secular velocity estimate, for which different published values for the stations in and around the Long Valley Caldera exist, which in turn enables a more straightforward validation than comparing extracted transients between methods or studies. In Section B.3, Figs. B.7–B.9, we compare our results with the MIDAS-derived secular velocities (Blewitt et al., 2016) and the Geodesy Advancing Geosciences and EarthScope (GAGE) facility’s secular velocities (Herring et al., 2016).




Figure 2.5: Modeled horizontal transient displacement (colored lines) of selected stations (names on the left) from Fig. 2.4, projected along the direction of maximum displacement during the period of 2012–2015. The directions (in grey to the right of the timeseries) are measured counterclockwise from east. CA99's direction is used for CASA. Black dots are the joint model's residuals, centered on the transient model.




Figure 2.6: Modeled seasonal vertical displacement timeseries for station P647 (see Fig. 2.4 for location). Blue and orange lines correspond to the annual and biannual constituents, respectively, and the black line is their sum. The deviation component only includes the annual frequency.

Seasonal Signals

Traditional least-squares model fitting for GNSS timeseries usually either approximate the seasonal signal as having a constant amplitude and phase over the entire timespan considered (or piecewise within that timeseries) (e.g., Heflin et al., 2020), or estimate a more accurate seasonal deformation signal from filtering or component-analysis methods (e.g., Silverii et al., 2020). The two approaches are usually acceptable, as either the resulting residuals are insignificant, or are not prone to producing large seasonal residuals in the first place. Given our transient modeling of even small timescales (down to the order of less than 100 days), our method does suffer from these seasonal residuals, as annual rain- and snowfall can vary widely, especially in the Sierra Nevada. In fact, because seasonal residuals are highly correlated between stations, they are not removed by our spatial L0 regularization. Modeling the seasonal signal as the sum of both an unregularized, constant, nominal constituent, and a simple, L1-regularized, station-specific deviation constituent of the same nominal frequencies that is allowed to vary in amplitude (and by construction, instantaneous phase) over time, the solver is once again able to separate seasonal (i.e., periodic) signals from (aperiodic) transient motion (see Appendix A.2). One example of the resulting seasonal fit in the vertical direction is shown in Fig. 2.6. Variations in the amplitude, and sometimes instantaneous phase, are clearly visible, demonstrating the importance of properly estimating and removing the seasonal sig-

nal at stations that are affected by major hydrological processes. Fig. B.10 shows the annual model’s vertical amplitude and phase in map view, and Fig. B.11 shows the overall vertical seasonal constituent for the stations in Fig. 2.5 for the entire timespan.

2.5 Discussion

The choice to incorporate process-agnostic, spatial awareness into the timeseries decomposition problem by means of a parametric, spline-based model that requires regularization and iteration may possibly appear odd — after all, vbICA and comparable methods already have an inherent sense of space. However, even though basis decompositions have a spatial component, the geometry of the network is neglected (e.g., relative distance between stations). Network geometry and extent become relevant when networks are large, and some signals are spatially confined: different processes at different locations may be mapped into the same component, complicating its interpretation. Furthermore, in order to obtain a clean decomposition using vbICA or similar methods, maintenance and earthquake coseismic offsets still have to be removed ahead of time, as well as the linear secular trends. Therefore, not only do these non-parametric decomposition approaches require a significant amount of preprocessing in the first place, the separation of preprocessing and actual decomposition precludes a straightforward way to quantify the covariance between the constituents. Using parametric models that are both process-aware (such as secular, seasonal, and maintenance offset models) and process-agnostic (using a dictionary of splines for transients and seasonal variations), by contrast, offers this correlation by design, while the spatial L0 regularization accomplishes the goals for sparsity and spatial awareness. We note that DISSTANS still does offer PCA/ICA, for example used for the common mode removal.

Parametric approaches allow one to include prior knowledge beyond the preprocessing steps. Incorporating such knowledge is already partly possible through the choice of the model functions (e.g., inserting a postseismic displacement model after a large earthquake), but least-squares-based methods such as the one used by DISSTANS also allow analytic inclusion of a priori model parameter knowledge, which may be added in future versions.

We omit a detailed look here at hyperparameters (e.g., regularization penalties, the number of iterations), as differing goals, as well as different characteristics of the data, will have a large impact on what the “best” choice is, and general assertions

are therefore not possible. The tools presented here therefore do not relieve the user of the task of finding the best set of hyperparameters for their data and problem formulation, although the documentation includes the specific choices for the cases presented in the previous section (based on both analytic and empirical considerations), which may provide a good starting point.

An important caveat of using a fixed dictionary of splines to model transient signals is that such fits are not phase-invariant. Processes that move both in space and time (e.g., slow slip events in Cascadia, Riel et al., 2014) are “discretized” by the onset times of individual splines, such that multiple splines (of possibly different periods) are necessary to capture a potentially simple signal that migrates in time. Failing to account for different onset times throughout the network could impact the quality of fit, as well as reduce the sparsity of the solution. However, experience shows that phase invariance is not as crucial as it may seem: First, observation noise makes exact onset times of transient signals hard to determine, and simultaneously allows the solver to fit splines that are adjacent in time when the “best” onset time would be somewhere in between the splines’ onset times. Second, if the problem persists, more splines of different periods or new onset times can be easily inserted into the models (with the main drawback being higher computational costs). In neither the synthetic nor real data examples presented here did the splines’ periods or onset times have to be adjusted from an initial, default configuration to obtain a high-quality decomposition.

2.6 Conclusion

Displacement timeseries of regional GNSS networks are commonly used to monitor surface deformation, plate motion, as well as transient signals such as hydrological loading or aseismic slip events. A crucial step in these analyses is the decomposition of the input (raw) timeseries into its constituents: secular motion, periodic seasonal variations, step offsets due to earthquakes, etc. As networks continue to grow in number and size, so does the need to efficiently analyze timeseries. We combine the many features of previously published analysis methods into a single, generic, open-source framework. The DISSTANS Python package includes: (1) incorporation of spatial information through the use of a spatial L0-regularized least-squares solver, (2) CPU-based parallelization for scaling to large networks, (3) formal uncertainty quantification with covariance matrices between components and constituents, (4) a suite of supporting tools including timeseries files data management, common mode estimation, and simple, automated step detection, as well as (5) visualization

methods to accelerate data and model inspection by the user.

Validation with synthetic GNSS network timeseries shows the beneficial effect of fitting transient signals with the spatial, L0-regularized solver: transients in the data are fit sparsely both in time and space, and are able to recover the true underlying motion better than comparable solutions without spatial awareness. An analysis of GNSS displacement timeseries from the Long Valley Caldera region in the Sierra Nevada, California, USA demonstrates the viability of our approach using real data, jointly decomposing the timeseries into step offsets, secular motion, transient signals, as well as time-varying seasonal displacements.

Chapter 3

PROBABILISTIC FRAMEWORK TO ESTIMATE
RHEOLOGICAL PROPERTIES IN SUBDUCTION ZONES
USING SEQUENCES OF EARTHQUAKES AND ASEISMIC
SLIP

Part of:

Köhne, T., Mallick, R., and Simons, M. (n.d.). “Probabilistic Estimation of Rheological Properties in Subduction Zones using Sequences of Earthquakes and Aseismic Slip.” In revision.

3.1 Introduction

Constraining the effective rheology of subduction zone megathrusts is crucial to our understanding of the physics of convergent plate boundary deformation (e.g., Bürgmann and Dresen, 2008). Key questions include: How does stress accumulate, release and redistribute during the earthquake cycle? Where and how are mountain ranges sustained? How can plate-like tectonics exist? And what does our understanding imply for seismic hazard assessments? Commonly adopted rheological descriptions of the megathrust come from laboratory experiments conducted on a limited number of rock types, at length scales varying from sub-microns to meters (e.g., Blanpied et al., 1995; Marone, 1998; Scholz, 1998; Hirth, 2002; Hirth and Kohlstedt, 2004). At low temperatures, laboratory observations appear consistent with the phenomenologically motivated rate-and-state friction (RSF) law (Dieterich, 1979; Rice and Ruina, 1983; Ruina, 1983), while various Newtonian and non-Newtonian viscous flow laws control the mode of deformation at elevated temperatures (Chopra and Paterson, 1981; Karato et al., 1986; Chopra, 1997; Hirth and Kohlstedt, 2004; Rutter and Brodie, 2004).

It is not obvious how to correctly upscale laboratory-calibrated rheologies to the geological scale, given that the role of material heterogeneities in rocks and interactions between physical (and chemical) processes at various scales remain poorly understood (Handy, 1994; Yamashita et al., 2015; Fagereng and Beall, 2021; Bercovici et al., 2023). Surface geodetic measurements are sensitive to a spatially and process averaged *effective* rheology, with the averaging scale comparable to the distance between the deforming source and the surface, and the form of averaging ambiguous.

Some authors have employed the phenomenologically motivated linear Burgers rheology in an attempt to consolidate transient behavior at multiple time scales (Yuen and Peltier, 1982; Müller, 1986; Pollitz, 2003; Masuti et al., 2016; Agata et al., 2019). In particular, linear Burgers bodies are numerically convenient to implement, however, it is unclear how to interpret the micromechanics of a Burgers body (Karato, 2021; Masuti and Barbot, 2021; Masuti et al., 2023). Regardless of the source of any proposed rheology, postseismic displacement timeseries observations of GNSS-based networks near plate interfaces can be used to estimate ranges of parameters for rheological models. However, it is unclear if geodetic observations can distinguish between different models at megathrust scales.

In the case of Northeastern Japan, several constitutive laws have been used to model surface postseismic deformation in the wake of the 2011 M_w 9.1 Tohoku-oki earthquake. For example, Sun et al. (2014) modeled the postseismic displacements using only flow in the viscoelastic mantle described by a Burgers rheology (Müller, 1986; Hetland and Hager, 2005). Subsequent improvements in the methodology of the inverse problem are highlighted in the works of Freed et al. (2017), who estimated the afterslip on the fault kinematically after modeling the viscoelastic mantle using a linear viscous rheology, and Fukuda and Johnson (2021), who jointly solve for stress-driven frictional afterslip and viscoelastic deformation as a Burgers body. Both Agata et al. (2019) and Muto et al. (2019) invoke non-steady-state or transient viscous rheologies in the mantle, in the form of power-law Burgers bodies (Chopra, 1997; Masuti et al., 2016), combined with frictional afterslip on the plate interface. Notably, none of the aforementioned studies attempted to verify whether their model setup and parameters would also sufficiently reproduce the preseismic deformation phase.

Based on the wide range of plausible models matching postseismic observations, it is therefore useful to study classes of rheological models that are consistent with the *entire* interseismic period (including the post- and preseismic time periods) using a *single* set of parameter values constant in time (Hearn and Thatcher, 2015; Mallick et al., 2022). Such models may provide additional constraints on which structural features or what ranges of rheological parameters are indeed required by the data, as opposed to being unconstrained free parameters that may lead to overfitting of the observations. We also note that constant model parameters may still yield time-dependent “effective” viscosities or strength terms, depending on the chosen rheology. To create a framework for the entire earthquake cycle, we build on the

formulation of Hetland and Simons (2010), Hetland et al. (2010), and Kanda et al. (2013) to simulate interseismic creep in an idealized subduction zone. In this model, we create sequences of earthquakes and aseismic slip (SEAS) by applying coseismic slip periodically on an otherwise locked zone (asperity) and letting the rest of the system mechanically evolve following a predefined rheological model. In contrast to the highlighted studies on Northern Japan in the previous paragraph, we assume all deformation on and around the plate interface can effectively be collapsed onto the fault plane, based on numerical simulations of the extent of shear localization with power-law rheological parameters (Moore and Parsons, 2015; Mallick et al., 2022). Conveniently, earthquake cycle models become independent of the initial state once fully spun-up (i.e., the simulation becomes cycle-invariant), avoiding the large influence these initial conditions impart (e.g., Agata et al., 2019; Muto et al., 2019). They are, of course, still sensitive to other assumptions, such as the fault geometry or the rupture history.

An additional benefit to using rheological models in earthquake cycle simulations is that the *absolute stress state* and *cumulative fault slip* are inherently tracked quantities (Kanda et al., 2013). The latter in particular enables direct predictions of the spatially-variable *slip deficit*; a key metric when assessing the seismic hazard of a region that yields the maximum amount of slip an earthquake could release. Commonly, slip deficit is approximated by assuming the fault slip rate is, to first order, constant throughout the seismic cycle, in which case it is equal to the directly-observable kinematic plate coupling. Significant differences arise, however, when there is substantial postseismic afterslip, a phenomenon common to large earthquakes (e.g., Perfettini and Avouac, 2014). While this disconnect has been analyzed in numerical, synthetic studies (e.g., Hetland and Simons, 2010), it has yet to be shown using real observations; a task our earthquake cycle model will enable us to tackle.

Fig. 3.1 summarizes the general workflow of our proposed framework, which is broadly divided into a forward and an inverse part. The remainder of this chapter describes the individual components in detail. The geometrical context which makes our framework applicable to subduction zones is presented in Section 3.2. At the core of our framework lie the rheological models that can be used to model relative slip on the interface between two plates. The mathematical definition of the two models we consider in this study, power-law viscosity and rate-dependent friction, are described in Section 3.3. The forward model, summarized in Section 3.4, builds

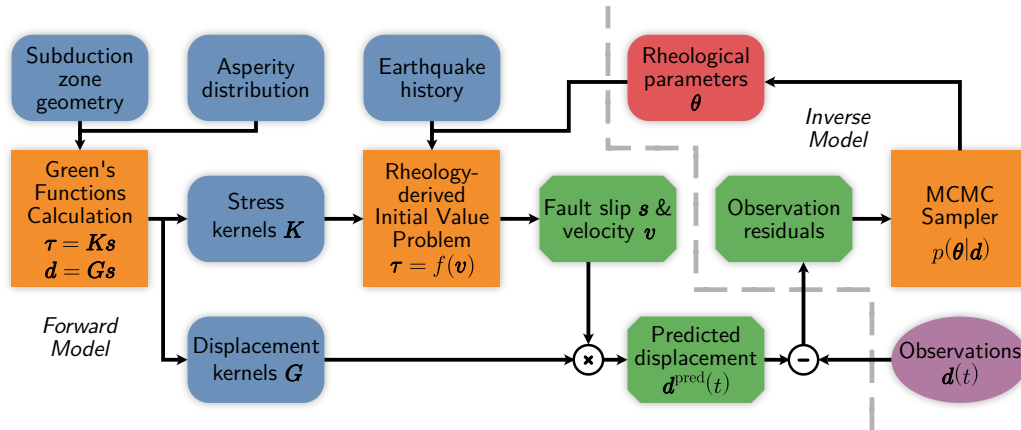


Figure 3.1: Workflow schematic, described in detail in Section 3, separated into the forward and inverse parts as shown by the dashed gray line. The forward model (Section 3.4) first uses the fault geometry (Section 3.2), asperity distribution, and earthquake history to create the stress and displacement kernels of the problem. The kernels are then used in the solution of the initial value problem given a rheological model (Section 3.3), yielding the slip history of each simulated fault patch. In the inverse problem (Section 3.5), a Markov-chain Monte Carlo sampler is used to explore the rheological parameters best matching the surface displacement timeseries.

upon the fault rheology and geometry to define a numerically-integrable initial value problem using an assumed earthquake history. Finally, the inverse model which enables the inference of model parameters given geodetic observations is presented in Section 3.5.

3.2 Fault Geometry

Geodetic observations at plate boundaries are commonly fit using models defined by discretized dislocation surfaces, i.e., finite-sized patches that express relative motion between their two sides. Savage (1983) introduced the still-popular, two-dimensional backslip model (BSM), in which a single dislocation surface represents the plate interface between the subducting and overriding plates. Over long (geologic) timescales, the incoming plate slides underneath the upper plate at a steady convergence rate, upon which the earthquake cycle is framed as a perturbation over (relatively) shorter timescales. Within the perturbation view (the “supplemental solution”, Savage, 1983), coseismic slip during earthquakes at shallow depths is balanced exactly by opposite motion during the interseismic period, resulting in a net-zero relative displacement. Since the longterm subduction is strain-free, it does not contribute to the observed deformation at the surface, making the perturbation view able to adequately model the surface deformation during both interseismic and coseismic phases (e.g., Aoki and Scholz, 2003; Suwa et al., 2006; Hetland and Si-

mons, 2010; Métois et al., 2016).

The BSM falls out of scope, however, when one approximates the steady-state convergence using dislocation surfaces (instead of “complex motions in the asthenosphere”, Savage, 1983), which would require the relative block motion between the two plates to result in a longterm uplift of the overriding plate, which is neither observed nor physical. The Elastic Subducting Plate Model (ESPM, Kanda and Simons, 2010) solves this issue by adding a second plate interface into the system, located below the original one. This additional interface can be thought of as the mechanical approximation of the bottom side of the subducting oceanic plate.

Our framework uses the ESPM directly for the 2D case, and naturally extends it for the 3D case. For the 2D case, we ignore the effect of bending and assume a planar downgoing slab. For the 3D case, we use a realistic (i.e., curved) plate interface based on geophysical observations. In both cases, we extend both the upper and lower interfaces to “infinity” to guarantee that the model domain has no impact on our results. In the 2D case, the only motion possible is along the line defining the plate interface. For the 3D model, we simulate motion in two dimensions on the fault interface, but restrict it to the quadrant $\pm 45^\circ$ around the local rake angle.

3.3 Rheological Models

In this section, we will present two rheological models that can describe the mechanical behavior of materials. For the purposes of this study, a rheological model is a constitutive law that relates stress (rate) to strain (rate). This relationship can be both linear and non-linear.

We begin our model with the partitioning of strain within any given viscoelastic element on the discretized plate interface as:

$$\epsilon = \epsilon_v + \epsilon_e \quad (3.1)$$

where the total strain, ϵ , is partitioned into an elastic ϵ_e and viscous component ϵ_v (e.g., Malvern, 1969), and all strain is parallel to the fault interface. We then assume that the relative shearing of the two sides of a plate interface can be approximated by a finite-width viscously shearing layer and v is the integrated shearing rate over the zone width, i.e., $v = 2h\dot{\epsilon}_v$, where h is the fault zone thickness and $\dot{\epsilon}_v$ is the viscous strain rate.

Slip on subduction plate boundaries is commonly modeled using a combination of a frictional and a viscoelastic constitutive law applied above the brittle-ductile tran-

sition depth and in the bulk, respectively. For the frictional behavior, rate-and-state (RSF) formulations (or their steady-state approximation, rate-dependent friction, RDF) are commonly used (e.g., Hetland et al., 2010). For viscoelastic rheologies, common formulations include power-law or linear Burgers constitutive laws (e.g., Sun et al., 2014; Johnson and Tebo, 2018; Peña et al., 2019; Peña et al., 2020). In the following sections, we will define the rheological models used for the simulations in Chapters 4–5.

3.3.1 Power-law Viscosity

The relationship between macroscopic stress and viscous strain rate in power-law models is described by

$$\alpha_n v = \tau^n \quad (3.2)$$

where α_n is a rheological strength term, v is the slip velocity on the fault patch, n is the power-law exponent, and τ is the shear traction exerted onto the fault patch (Hirth, 2002; Hirth and Kohlstedt, 2004). Typical values in high-temperature laboratory settings vary between 1–4 depending on the dominant deformation mechanism: $n = 1$ for pure linear-viscous diffusion creep, $n \approx 2$ for certain grain-boundary processes, and $3 \lesssim n \lesssim 4$ for dislocation creep.

We model each patch as a spring-dashpot (for $n \lesssim 10$) or spring-slider (for $n \rightarrow \infty$) element in series, while the elastic connection between any two pairs of patches can be modeled as a spring (Segall, 2010). By varying n , a power-law rheology includes as end member cases both linear viscosity (for $n = 1$) and rate-dependent friction (for $n \rightarrow \infty$), see, e.g., Montési and Hirth (2003), Montési (2004), and Mallick et al. (2022). Note that we assume the bulk to behave purely elastically (see Section 3.4), i.e., all viscous flow in the medium away from the plate interface is assumed to be modeled as part of the viscous shearing in the fault zone.

To enable comparisons of different models at different times in the earthquake cycle (and therefore, at different interface slip rates), we derive the *effective* viscosity as

$$\alpha_{\text{eff}} = \alpha_n^{1/n} \cdot v^{1/n-1} \quad (3.3)$$

which introduces the patch velocity v as a dependence. The effective viscosity α_{eff} is related to the fault zone viscosity η as $\alpha_{\text{eff}} = \eta/h$.

To numerically integrate the behavior of a power-law viscous element, we need the

time derivative of (3.2):

$$\frac{dv}{dt} = \frac{n}{\alpha_n} \tau^{n-1} \frac{d\tau}{dt} \quad (3.4)$$

where

$$\tau^{n-1} = (\alpha_n v)^{1-1/n} \quad (3.5)$$

simplifying to

$$\frac{dv}{dt} = \frac{n}{\alpha_n^{1/n}} v^{1-1/n} \frac{d\tau}{dt}. \quad (3.6)$$

Step changes in the stress lead to step changes in velocity as follows:

$$\tau_+ = \tau_- + \Delta\tau \quad (3.7)$$

$$\Leftrightarrow v_+ = (\alpha_n^{1/n} \sqrt[n]{v_-} + \Delta\tau)^n / \alpha_n. \quad (3.8)$$

3.3.2 Rate-dependent Friction

Rate-dependent friction introduces the dependence of the friction coefficient on the sliding velocity as follows:

$$f_{ss} = f_0 + (a - b)\zeta \quad (3.9)$$

where f_{ss} is the steady-state friction, f_0 is a reference friction, a and b are the rate-and-state frictional parameters, $\zeta = \log(v/v_0)$ is the logarithmic velocity, v is the linear velocity, and v_0 is a reference velocity. We only consider rate-strengthening rheologies, i.e., $(a - b) > 0$, since we purposefully are only interested in modeling stress-driven, stably decaying postseismic sliding (Dieterich, 1979; Rice and Ruina, 1983; Ruina, 1983). Typical laboratory-derived values for $(a - b)$ are in the range of 10^{-3} – 10^{-2} (Marone et al., 1991; Blanpied et al., 1995).

As for the power-law-viscous case where $n \rightarrow \infty$, we model each fault patch as a spring-slider element in series, with the elastic connection between any two pairs of patches modeled as a spring. The slider element translates static stress to shear resistance through the steady-state coefficient of friction f_{ss} as described by Coulomb friction:

$$\tau = f_{ss} \sigma_E$$

with the shear stress τ and the effective fault normal stress σ_E .

Taking the time derivative of the steady-state friction gives an explicit formulation for the slip acceleration $d\zeta/dt$:

$$\frac{df_{ss}}{dt} = (a - b) \frac{d\zeta}{dt}.$$

Inserting the Coulomb friction relationship and assuming constant effective normal stress leads to

$$\frac{d\tau}{dt} = \sigma_E \frac{df_{ss}}{dt}$$

which can be rearranged to give the final expression

$$\frac{d\zeta}{dt} = \frac{1}{(a - b)\sigma_E} \frac{d\tau}{dt} \quad (3.10)$$

where $\alpha_h = (a - b)\sigma_E$ is the rate-dependent frictional parameter.

Step changes in the velocity when an earthquake is occurring on the locked asperity can be calculated starting from

$$\tau_+ = \tau_- + \Delta\tau.$$

Plugging in the expression for τ yields

$$\log \frac{v_+}{v_0} = \log \frac{v_-}{v_0} + \frac{\Delta\tau}{\alpha_h}$$

and finally

$$\zeta_+ = \zeta_- + \frac{\Delta\tau}{\alpha_h}. \quad (3.11)$$

3.4 Forward Model

3.4.1 Initial Value Problem Formulation

We start from the quasistatic stress equilibrium in an elastic half-space with an embedded fault following Segall (2010). Here, the far-field loading traction τ^∞ and the elastic loading τ^{el} in the spring element of the patch are balanced by the resistive shear traction τ :

$$\tau = \tau^\infty + \tau^{\text{el}}. \quad (3.12)$$

To approximate energy radiated away during dynamic (high stressing rate) times, we add a damping term, yielding:

$$\tau = \tau^\infty + \tau^{\text{el}} - \frac{\mu}{2v_s} v \quad (3.13)$$

where μ is the shear modulus of the bulk, v_s is the shear wave velocity, and v is the sliding velocity of the patch (Rice, 1993).

We account for the finite dimensions of the mechanical problem by incorporating an elastic interaction (stress) kernel K , which we approximate using constant slip boundary elements (Crouch and Starfield, 1983). The tractions resulting from the slip interactions of all discretized fault elements are calculated as (note the vectorized, boldface notation to include all elements):

$$\boldsymbol{\tau}^{\text{el}} = \mathbf{K} \mathbf{s}. \quad (3.14)$$

The far-field plate boundary loading can be approximated analogously as:

$$\boldsymbol{\tau}^{\infty} = -\mathbf{K} \mathbf{s}^{\infty} \quad (3.15)$$

where $\mathbf{s}^{\infty} = \mathbf{v}^{\infty} t$ is the accumulated amount of slip the opposite sides on each patch should have moved relative to each other in a steadily-creeping subduction system. Using the stress kernel formulations in (3.13) and taking their time derivative yields:

$$\frac{d\boldsymbol{\tau}}{dt} = \mathbf{K} (\mathbf{v} - \mathbf{v}^{\infty}) - \frac{\mu}{2v_s} \frac{d\mathbf{v}}{dt}. \quad (3.16)$$

Here, we have two unknowns that define the mechanical system: $d\boldsymbol{\tau}/dt$ and $d\mathbf{v}/dt$. The derivatives can be taken directly from the definitions of the rheological models in Section 3.3. The final expressions can then be numerically integrated to obtain the time history of slip and velocity over the model domain (here, we use an adaptive-timestep Runge-Kutta method of order 4). To simulate the effect of earthquakes, coseismic stress changes are applied at predetermined times and with predetermined magnitudes based on the velocity step changes (also defined in the rheological models, see Sections 3.3 and 3.4.2).

For the case of **power-law viscosity**, inserting (3.2) and (3.6) into (3.16) yields

$$\frac{d\boldsymbol{\tau}}{dt} = \frac{\mathbf{K} (\mathbf{v} - \mathbf{v}^{\infty})}{1 + \frac{\mu}{2v_s} \frac{n}{\alpha_n^{1/n}} \mathbf{v}^{1-1/n}} \quad (3.17)$$

$$\frac{d\mathbf{v}}{dt} = \frac{\mathbf{K} (\mathbf{v} - \mathbf{v}^{\infty})}{\frac{\alpha_n^{1/n}}{n} \mathbf{v}^{1/n-1} + \frac{\mu}{2v_s}}. \quad (3.18)$$

The state vector we integrate is

$$\begin{bmatrix} \mathbf{s} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} \mathbf{v} \\ \frac{d\mathbf{v}}{dt} \end{bmatrix}. \quad (3.19)$$

In the case of **rate-dependent friction**, we first need to calculate $d\mathbf{v}/dt$ from $d\zeta/dt$ using the definition of the logarithmic velocity in Section 3.3.2 and eq. (3.10):

$$\frac{d\mathbf{v}}{dt} = \mathbf{v} \frac{d\zeta}{dt} = \frac{\mathbf{v}}{(a-b)\sigma_E} \frac{d\tau}{dt}.$$

Inserting this into (3.16) yields

$$\frac{d\tau}{dt} = \mathbf{K}(\mathbf{v} - \mathbf{v}^\infty) - \frac{\mu}{2v_s} \frac{\mathbf{v}}{(a-b)\sigma_E} \frac{d\tau}{dt} \quad (3.20)$$

$$\Leftrightarrow \frac{d\tau}{dt} = \frac{\mathbf{K}(\mathbf{v} - \mathbf{v}^\infty)}{1 + \frac{\mu}{2v_s} \frac{\mathbf{v}}{(a-b)\sigma_E}} \quad (3.21)$$

and

$$\frac{d\zeta}{dt} = \frac{\mathbf{K}(\mathbf{v} - \mathbf{v}^\infty)}{(a-b)\sigma_E + \frac{\mu}{2v_s} \mathbf{v}}. \quad (3.22)$$

The state vector we integrate is

$$\begin{bmatrix} s \\ \zeta \end{bmatrix} = \begin{bmatrix} \mathbf{v}_0 \exp \zeta \\ \frac{d\zeta}{dt} \end{bmatrix}. \quad (3.23)$$

Finally, to take advantage of both the per-model runtime improvement of compiled programming languages, and to enable the parallel calculation of dozens to thousands of modeled systems (which will be required by the inverse model), we implement our forward model in the CUDA computing framework. This enables the integration of the system of ordinary differential equations to be carried out on High Performance Computing clusters using many Graphics Processing Units (GPUs). Parallelization is further aided by the fact that all displacement and stress kernels are constant quantities in our model, i.e., independent of the rheological model chosen.

3.4.2 Earthquake History

Our framework of sequences of earthquakes and aseismic slip requires the definition of both the “earthquakes” and the “aseismic slip” parts. The initial value problem (IVP) formulation of Section 3.4.1 deals with the latter, leaving the characterization of the former for this section. More precisely, the earthquake history provides the initial values for the numerical IVP integrator, i.e., the slip and slip rate values s_+ and \mathbf{v}_+ (or ζ_+) after an earthquake, to be integrated forward until the next one. To fully define the rupture history, in turn, we need the timing, slip distribution, and

location of all considered earthquakes. We note that imposing earthquakes at predefined recurrence intervals using characteristic slip amounts derived from the plate convergence rate is equivalent to simulating the stress history at the asperities and letting them fail whenever a critical stress or slip deficit value is exceeded.

To use the state-change equations (3.7) and (3.11) from Section 3.3, we need to calculate $\Delta\tau$ at all the simulated (creeping) fault patches due to the earthquake slip as follows:

$$\Delta\tau = K\Delta s \quad (3.24)$$

where Δs is the coseismic slip inside the asperities, and tapered to zero outside the asperities. Since earthquakes in our model are imposed and not estimated, they must be defined *a priori* for each case study, e.g., by considering finite fault slip models created using coseismic geodetic motion or seismic data. When there are multiple asperities present, it must furthermore be specified which of these rupture at any given event time.

To define the timing of an earthquake, we start by assuming that each asperity has a characteristic recurrence time, releasing all the accumulated slip deficit at once. The recurrence times T_{rec} and slip magnitudes Δs are defined using seismic catalogs and large-scale plate convergence rates. In a next step, we may relax the assumption of uniform recurrence times and slip magnitude and posit that only after a predefined “supercycle” duration T_{cycle} , the total slip deficit has to be recovered. We can achieve this flexibility by allowing both the timing and the magnitude of each event within that time frame to vary according to a log-normal distribution. We set the length of T_{cycle} to be the lowest common multiplier of all the asperities’ recurrence times. Fig. 3.2 shows an example earthquake history realization for a case with a single asperity and uniform slip.

We acknowledge that most likely, asperities are neither constant in space, nor are their recurrence time well-modeled by a log-normal random variable (e.g., Simons et al., 2011; Kanda et al., 2013). However, we find in Section 4.3.2 that as long as the deviations from this assumption happen far enough before the period geodetically observed (and fitted), the violation of this assumption does not significantly negatively impact our recovery of rheological parameters. This situation is usually satisfied, as seismic catalogs significantly predate modern GNSS- or InSAR-based surface deformation timeseries.

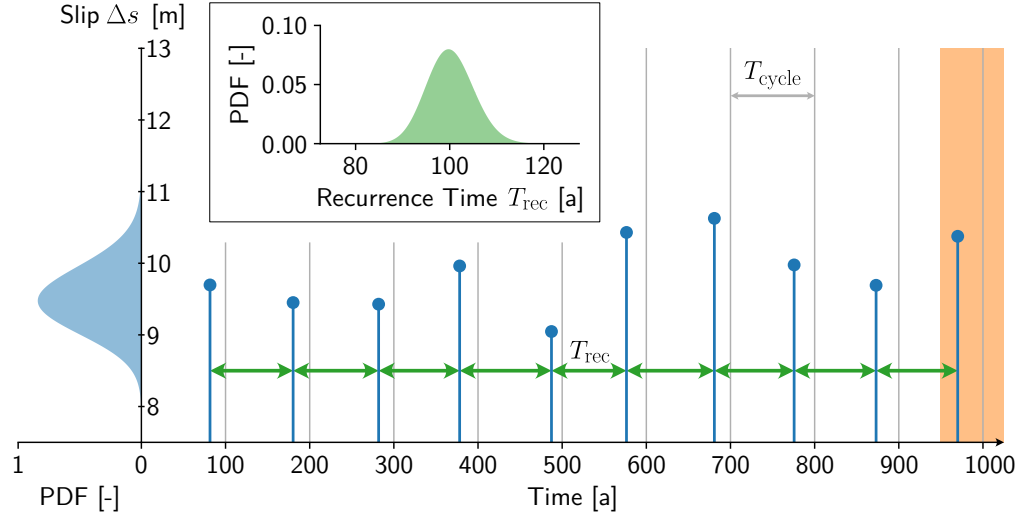


Figure 3.2: Realization of an example earthquake history where both recurrence time T_{rec} and coseismic slip amount Δs are varied over 10 cycles. The cycle length T_{cycle} is set to 100 years, and the observations begin after 950 years (orange shading). The blue dots above vertical lines mark the time (horizontal axis) and slip amount (vertical axis) of each imposed earthquake. The slip amounts for the 10 earthquakes are sampled from a log-normal distribution with scale parameter $\sigma = 0.05$ and mean $\mu = 9.5$ m; the probability density function (PDF) is shown in blue on the left. The recurrence times are sampled from a log-normal distribution with scale parameter $\sigma = 0.05$ and mean $\mu = 100$ a; the PDF is shown in green in the insert at the top.

3.5 Inverse Model

3.5.1 Bayesian Inference

For the inversion scheme, we use the Markov-chain Monte Carlo (MCMC) software AlTar which implements the CATMIP algorithm described in Minson et al. (2013), Duputel et al. (2014), Jolivet et al. (2014), Duputel et al. (2015), and Jolivet et al. (2015), which in turn uses Bayesian inference to recover parameters that best match input data. The AlTar software was modified for this study to allow for the forward model to be a general function (in this case, an initial value problem solver), rather than the default matrix multiplication for linear inverse problems.

For both 2D and 3D cases, using synthetic or real data, the inverse model aims to fit surface observations \mathbf{d} by finding the set of model parameters $\boldsymbol{\theta}$ that recreate the observations best. We denote real observations derived from geodetic networks as \mathbf{d}^{obs} , and synthetic observations derived from a predetermined set of parameters as $\mathbf{d}^{\text{target}} = g(\boldsymbol{\theta}_{\text{target}})$ using the forward model $g(\cdot)$ (described in Section 3.4). The log-likelihood of a given sample evaluated by AlTar is calculated from the residual displacements $\mathbf{d}^{\text{target/obs}} - g(\boldsymbol{\theta})$, which we assume to be normally distributed

with zero mean. We can then formulate the Bayesian problem as $p(\boldsymbol{\theta}|\mathbf{d}^{\text{targetobs}}) \propto p(\mathbf{d}^{\text{targetobs}}|\boldsymbol{\theta})p(\boldsymbol{\theta})$, where $p(\boldsymbol{\theta})$ is the prior probability density function (PDF). We furthermore assume $p(\mathbf{d}^{\text{targetobs}}|\boldsymbol{\theta}) \sim \mathcal{N}(\mathbf{d}^{\text{targetobs}}|g(\boldsymbol{\theta}), \mathbf{C}_d)$, where \mathcal{N} is the normal distribution and $\mathbf{C}_d = \sigma_d \mathbf{I}$ is the synthetic, diagonal observation covariance matrix with uniform standard deviation σ_d .

3.5.2 Exploring Uncertainty in the Assumptions

In its basic form, the forward model $g(\cdot)$ is a deterministic initial value problem fully defined through the rheological model parameter set $\boldsymbol{\theta}$ and the hyperparameters describing the fault geometry, the asperity distribution, and the earthquake history (see Fig. 3.1). Because of the deterministic aspect of $g(\cdot)$, all uncertainties in the posterior estimates are either due to observational noise, inherent non-uniqueness in the forward model, or incompatibility of the model used to generate the synthetic target displacement timeseries and the models assumed in the inversion. However, the influence of the hyperparameters on the model predictions (i.e., surface displacement timeseries as output by $g(\cdot)$) may be significant. In our framework, we are able to explore the effects of variable hyperparameters on the recovery of $\boldsymbol{\theta}$ since we can change the hyperparameters individually for each sample (i.e., each tested forward model). Varying hyperparameters makes the forward model $g(\cdot)$ non-deterministic for the purposes of the MCMC estimation, which is only aware of the parameters $\boldsymbol{\theta}$. The non-deterministic behavior will manifest itself within AITar as prediction uncertainty \mathbf{C}_p (as opposed to data uncertainty captured by \mathbf{C}_d) and change the uncertainty in the posterior probability density function $p(\boldsymbol{\theta}|\mathbf{d})$ (Minson et al., 2013; Duputel et al., 2014). Indeed, one could run a large set of forward models only varying a single hyperparameter, and calculate an empirical covariance matrix that AITar could take into account. Since we obtain satisfying results without this additional computational step, the exploration of empirically-derived \mathbf{C}_p is outside the scope of this study.

As an example of the flexibility of our framework to explore uncertainties in the hyperparameters, every time the solver calls $g(\cdot)$, we can vary the characteristic earthquake recurrence time, directly affecting the length of the interseismic period preceding the observed earthquake (see Fig. 3.2). We can go further and create sequences of earthquakes preceding the observations that have randomized lengths of the interseismic period. These variations can also be applied to the amount of slip for any given earthquake in the cycle. Together, the variability of the earthquake history hyperparameters enables the exploration of the importance of uncertainties

in these values, as one would undoubtedly have when using real-world earthquake history datasets.

3.6 Summary

In this chapter, we have presented a framework to probabilistically invert surface geodetic measurements for rheological model parameters of subduction zones, based on previous work by Kanda et al. (2013). The geometry of the plate interface is modeled following the Elastic Subducting Plate Model (Kanda and Simons, 2010). The forward model takes the geometry, an assumed earthquake history, and the type of rheological model considered as inputs, and produces surface deformation timeseries as an output. The inverse model makes use of the AlTar software (Minson et al., 2013) to probabilistically determine the best-fitting rheological parameters using a Markov-chain Monte Carlo approach. To validate this framework, we explore its performance and sensitivity using a power-law viscous rheology in Chapter 4. We then apply it to a case study of the Northern Japan megathrust in Chapter 5 assuming rate-dependent friction on the interface.

*Chapter 4***FRAMEWORK PERFORMANCE AND SENSITIVITIES FROM
SYNTHETIC 2D MODELS**

Part of:

Köhne, T., Mallick, R., and Simons, M. (n.d.). “Probabilistic Estimation of Rheological Properties in Subduction Zones using Sequences of Earthquakes and Aseismic Slip.” In revision.

4.1 Introduction

This chapter aims to investigate the performance and sensitivity of the SEAS inversion framework presented in Chapter 3. Crucially, our primary goal is to confirm that the proposed method is able to recover rheological parameters in synthetic test cases. Our secondary goal is to assess the influence of various model assumptions in an effort to guide the applicability of the framework to real data.

We start by defining a representative, simplified, two-dimensional subduction zone geometry and rupture history in Section 4.2. To explore a wide range of potential model behavior, we choose the power-law viscous constitutive equations from Section 3.3.1, whose endmembers include both linear viscous flow and frictional resistance. We then assume a target effective viscosity profile, and generate synthetic target surface displacement timeseries using the forward model presented in Section 3.4. For our reference case presented in Section 4.3, we invert this target dataset to recover best-fitting viscosity profiles. We also construct additional test cases that restrict the data available to the inversions and/or loosen several model assumptions to gauge the degree to which these changes affect our inversion results. Finally, we discuss our findings in the context of subduction zone studies in Section 4.4, before proceeding with real observations in Chapter 5.

4.2 Model Setup**4.2.1 Fault Geometry**

Our 2D fault geometry directly follows the setup described in Section 3.2 (Fig. 4.1). The shallow part of the upper plate interface is locked and only slips coseismically (the asperity). The rest of the upper plate interface creeps interseismically as a response to the stress induced by the imposed coseismic slip in the asperity and as




Figure 4.1: Fault setup for our simulations, based on the Elastic Subducting Plate Model (ESPM, Kanda and Simons, 2010). The downgoing slab is modeled by two interfaces bounding an effective elastic plate of constant thickness H . The upper and lower interfaces experience reverse and normal motion, respectively. The upper interface is furthermore assumed fully coupled (i.e., locked) from the trench down to a depth of D_{shallow} . Along the non-coupled upper interface, shear resistance is described using rheological models (see Section 3.3) with depth-varying parameters with knickpoints at D_{middle} and D_{deep} . Displacement timeseries are evaluated at discrete locations along the free surface and labeled S1–5. The length scales, observer coverage, and dipping angle are loosely based on the Northern Japan subduction zone.

a function of the chosen rheological model. The lower plate interface is assumed to be creeping at the plate convergence rate, which is a convenient simplification for approximating the stresses resulting from bending the slab as it passes through the trench (Mckenzie, 1977; Mueller and Phillips, 1991; Kanda and Simons, 2010). It also ensures the longterm motion of the subducting slab below the overriding plate (Kanda and Simons, 2010). The entire fault is assumed to be embedded in a linear-elastic half-space, and the creeping part of the upper plate interface is simulated through a collection of line dislocation elements (“patches”). For the rest of this study, we will omit the *upper* and *lower* qualifiers for the plate interface for conciseness, and *plate interface* will always refer to the upper plate interface unless otherwise noted. We compute the displacement and stress kernels necessary for the forward problem using Davis (2017).

4.2.2 Synthetic Target Observations

We start our exploration of the framework with a reference case assuming power-law viscosity ($\tau^n = \alpha_n v$) as the rheological model, which can be used to approximate a wide variety of behavior, including friction and linear viscosity, depending on the exponent n (Montési and Hirth, 2003; Montési, 2004; Mallick et al., 2022). The profile of the *effective* viscosity, derived using eq. (3.3), is shown in Fig. 4.4a, and loosely follows the assumed trend of having a frictional, relatively weak shallow

Table 4.1: Rheological parameters α_{eff}, n used to create the synthetic target dataset of surface displacement timeseries, and their respective knickpoint depths D . Note that α_{eff} is defined in (3.3) and depends on a reference slip velocity.

	Shallow	Middle	Deep
D [km]	35	50	100
α_{eff} [Pa s/m]	10^{12}	10^{15}	10^{13}
n [–]	10	6	1

interface, strengthening towards intermediate depths, and then further transitioning into a weaker, viscous interface at larger depths (see Table 4.1 for values). For the rupture history, we choose an earthquake slip of 9.5 m with a recurrence time interval of 100 years (approximately the convergence rate between the Pacific Plate and Northern Japan) up to a depth of 35 km. The generated surface displacements $\mathbf{d}^{\text{target}}$ are shown in Fig. 4.2, covering a time period of approximately 15 years before the earthquake and 10 years after the earthquake. For the inversions, we add normally-distributed noise of a standard deviation of 10 mm to create the synthetic, noisy observations.

4.2.3 Fitted Parameters

In the inverse model for the reference case, we aim to fit the power-law viscosity parameters α_{eff}, n (see Section 3.3.1 for a mathematical description of the parameters) at three different depths, as well as the depths themselves, fully defining the viscosity profile that was used to generate the target displacement timeseries. The three different depths are denoted “shallow”, “middle”, and “deep”, leading to the following parameter set:

$$\boldsymbol{\theta} = \{ \alpha_{\text{eff,shallow}}, \alpha_{\text{eff,middle}}, \alpha_{\text{eff,deep}}, n_{\text{shallow}}, n_{\text{middle}}, n_{\text{deep}}, D_{\text{middle}}, D_{\text{deep}} \} \quad (4.1)$$

We do not estimate D_{shallow} , corresponding to the locking depth, which we keep constant to not have to recreate the fault mesh and kernels at every sampling step. In later inversions, we estimate the rate-dependent frictional parameter $\alpha_{\text{h}} = (a - b)\sigma_E$ at different depths instead of α_{eff}, n (see Section 3.3.2 for a full mathematical description).




Figure 4.2: Simulated surface displacement timeseries used as input for the reference case. The solid lines are the timeseries d^{target} , with colors corresponding to observer location. The labels refer to the names of the stations, with the range of trenchward distances given as well. The black dots are the synthetic observations which include a 10 mm standard deviation Gaussian noise. At approx. 14.2 years (vertical gray line), an earthquake occurs and starts a postseismic transient process. The coseismic offset (including the effect of tapered slip) are both removed in this plot, and in the observations used in the inversion.

4.3 Inversion Results

The output of the Altar inversion routine is the posterior sample distribution of each parameter in the set θ , which approximates the posterior probability density function of each parameter. Derived from θ is the resulting depth-varying profile of the effective viscosity. We present the results of several experiments, using corner plots visualizing the posterior probability densities (revealing parameter covariances), as well as plots of the error on effective viscosity with depth. To quantify the quality of fit to both the target (synthetic) surface displacement timeseries, and the target underlying viscosity profiles, we furthermore define two error metrics, $\delta_\alpha(t_i)$ and δ_d , the average viscosity and surface displacement errors, respectively:

$$\delta_\alpha(t_i) = \frac{1}{L} \sum_{m=1}^{N_p} l_m \left(\frac{1}{N_s} \sum_{k=1}^{N_s} \left(\log_{10} \alpha_{\text{eff},m,k}^{\text{pred}}(t_i) - \log_{10} \alpha_{\text{eff},m}^{\text{target}}(t_i) \right)^2 \right)^{1/2} \quad (4.2)$$

$$\delta_d = \frac{1}{2N_t N_o} \sum_{i=1}^{N_t} \sum_{j=1}^{N_o} \sum_{c=\text{H,V}} \left(\frac{1}{N_s} \sum_{k=1}^{N_s} \left(d_{i,j,c,k}^{\text{pred}} - d_{i,j,c}^{\text{target}} \right)^2 \right)^{1/2} \quad (4.3)$$

where t is time, i is the index of the N_t observations, j is the index of the N_o observer stations at the surface, c is one of the horizontal or vertical data components H, V, respectively, k is the index of the N_s MCMC samples, m is the index of the N_p simulated patches, l_m is the patch length, $L = \sum_m l_m$ is the total depth of the simulated part of the plate interface, $\alpha_{\text{eff},m,k}^{\text{pred}}$ is the predicted effective viscosity, $\alpha_{\text{eff},m}^{\text{target}}$ is the target effective viscosity, $d_{i,j,c,k}^{\text{pred}}$ is the predicted surface displacement, and $d_{i,j,c}^{\text{target}}$ is the target surface displacement. In words, the viscosity error $\delta_\alpha(t)$ is defined as the depth-weighted average root-mean-square difference between the predicted viscosities by each MCMC sample and the target viscosity, and the surface displacement error δ_d is the spatiotemporal average of the root-mean-square difference between the observations predicted by each MCMC sample and the target surface displacements. A value of $\delta_\alpha = 1$ would therefore imply that the viscosity at any given depth is approximated to one order of magnitude on average. For the rate-dependent friction inversions, δ_α is undefined.

4.3.1 Reference Case: Power-law Viscosity

The reference case inversion for 4800 samples with a chain length of 50 completed in 7.25 hours using all 32 threads on 3 CPUs. We verify that additional samples do not change the recovered posterior covariance matrix. The synthetic target surface displacement timeseries are on average well fit, with $\delta_d = 1.4$ mm being much smaller than the observation uncertainty of 10 mm (compare Fig. 4.2). Additionally, the

transition depths are recovered well, and to a certain extent, the effective viscosity α_{eff} (assuming plate convergence rate as the conversion velocity) and exponent n at the upper transition depth. Fig. 4.3 shows the posterior covariance matrix by means of a corner plot. There are large uncertainties in the recovery of the shallow and deep effective viscosities and exponents. However, a lot of the uncertainty stems from strong correlation between n and α_{eff} , suggesting that there is some information content recovered when combining the two parameters. Using the simulated velocity profiles just before and after the applied earthquake, we can convert the covariance in the parameter samples into uncertainty in the depth-dependent effective viscosity profile as shown in Fig. 4.4b–c. Indeed, the recovery of the viscosity profile suggests that the MCMC sampler has identified the depth- and velocity-dependent effective viscosity as the recovered variable. Overall, the viscosity error is significantly less than one order of magnitude, and does not increase significantly with depth, with $\delta_\alpha = 0.22, 0.18$ for the preseismic and postseismic timesteps, respectively.

4.3.2 Adding Observational Biases and Uncertainties

We further examine the effect on the posterior distribution of estimated parameters for different test cases approaching real-world observing conditions. Such conditions can be approximated by reducing the input data (e.g., shortening the timespan or only using a single data component) or loosening the earthquake history assumptions. In total, we ran the following inversions for power-law viscous rheologies:

Case (1): The reference case presented in Section 4.3.1.

Case (2): Only the postseismic observations are used by the MCMC sampler.

Case (3): Only the *horizontal* postseismic observations after an “observational gap” of 14 days are used by the MCMC sampler.

Case (4): Like (3), although the coseismic slip applied in the locked zone in each MCMC sample is also varied according to a log-normal distribution (compare Fig. 3.2), yielding variations of $\approx 10\text{--}15\%$.

Case (5): Each MCMC sample has an irregular earthquake cycle, where the recurrence interval is varied according to a log-normal distribution (compare Fig. 3.2), yielding variations of $\approx 10\text{--}15\%$.

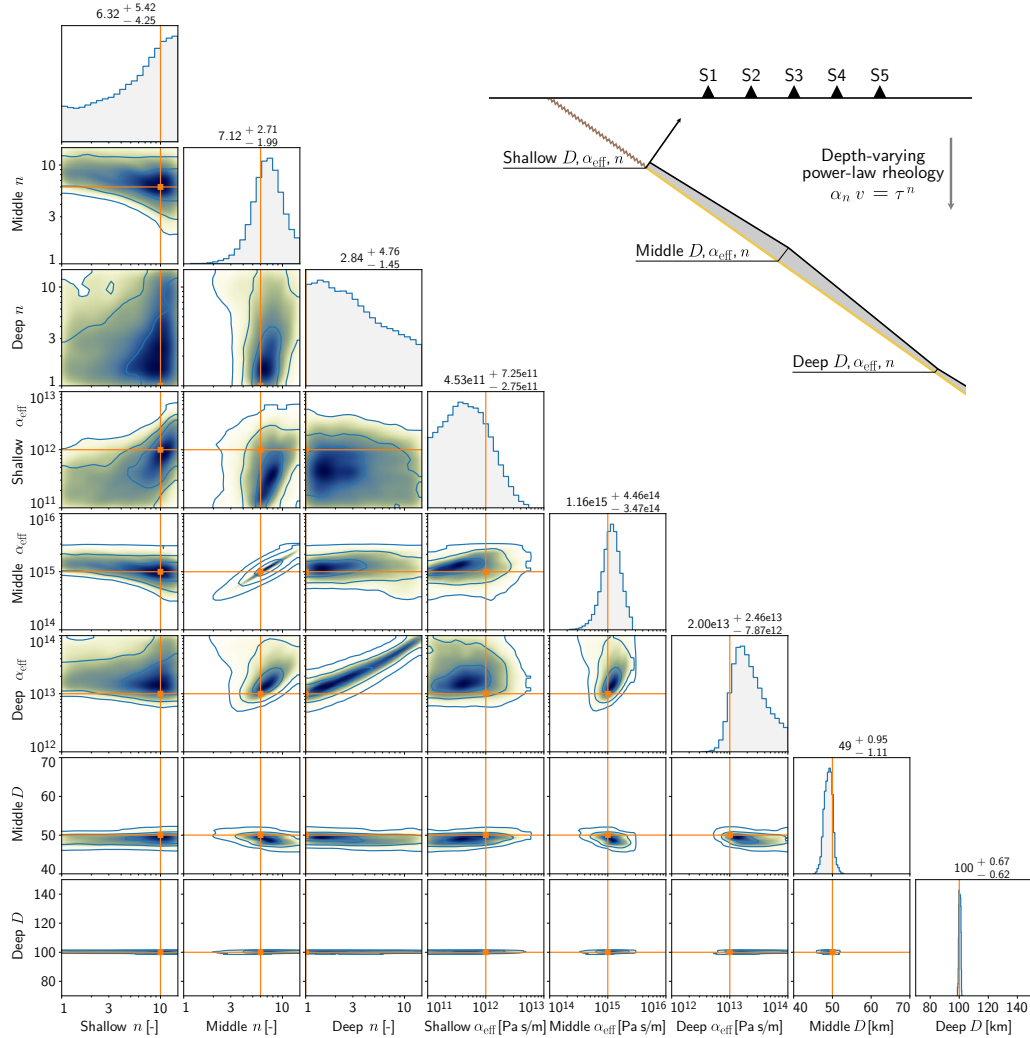


Figure 4.3: Corner plot of the posterior covariance matrix for the reference case as approximated by the MCMC inversion process for the eight estimated parameters defining the depth-dependent power-law rheology structure (see Fig. 4.4a). D denotes the knickpoints in the depth-variable viscosity structure, and α_{eff} and n denote the effective viscosities and stress exponents at different depths, respectively. The figures on the diagonal represent smoothed 1D histograms of the marginalized posterior probability density functions (PDF) for each parameter. The numerical values of the median as well as the ± 34 th percentile (equivalent to one standard deviation for a normal distribution) are shown above the diagonal. The lower, off-diagonal plots are smoothed 2D histograms of the posterior PDF, with contour lines indicating the 1, 2, and 3 standard deviation ranges from the mean. The orange lines are the values θ_{target} used in generating the target synthetic data. The inset in the top right corner (simplified version of Fig. 4.1, described in detail in Section 4.2) relates the eight estimated parameters to their definition within the fault geometry. The effective viscosity α_{eff} is related to the rheological strength term α_n as $\alpha_{\text{eff}} = \alpha_n^{1/n} \cdot v^{1/n-1}$ (see Section 3.3).

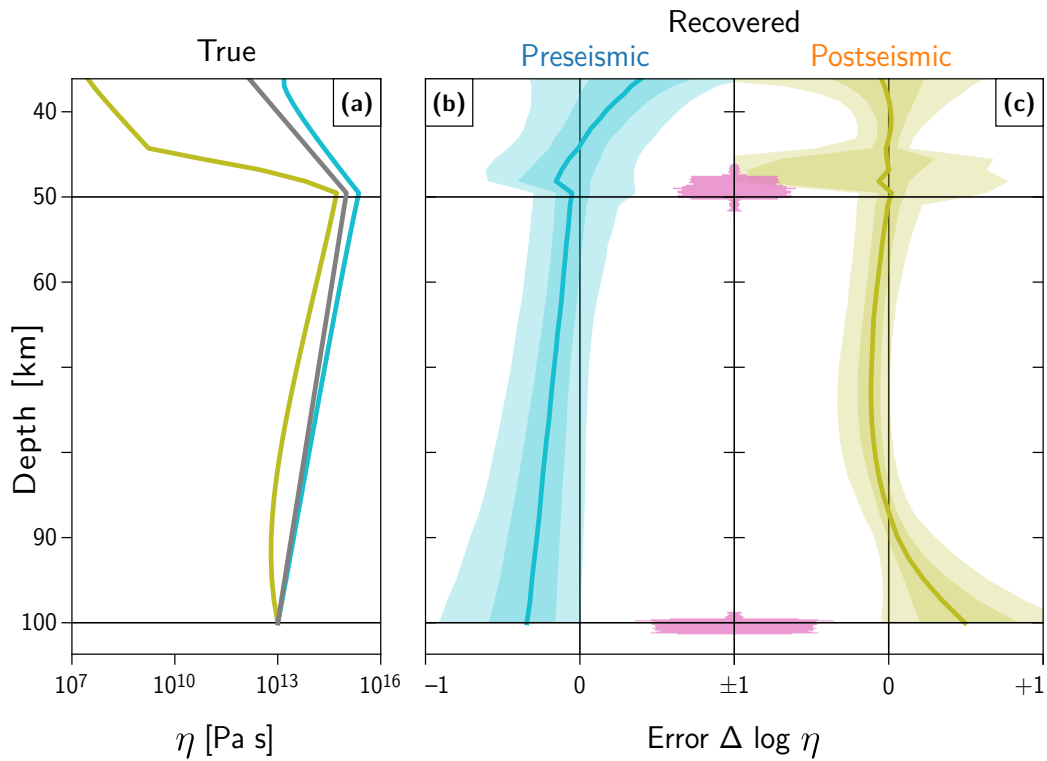


Figure 4.4: Target effective viscosity structure and posterior error for the reference case, converted to viscosity assuming $h = 1$ m. The target viscosity structure used in all cases is given in the leftmost panel (a) for the depth range of creeping fault patches. The gray curve uses the plate convergence rate to convert from the power-law rheological parameters to effective viscosity, and the blue and orange curves use the pre- and postseismic fault patch velocities, respectively. Panels (b–c) show the error in the recovered viscosity structure, as well as the recovered transition depths. Panel (b) corresponds to the preseismic apparent viscosity, and panel (c) to the postseismic one. The blue and orange solid lines show the median error between the logarithm of the recovered (predicted) viscosity profile and the target values, where ± 1 refers to an error of one order of magnitude. The light and medium shaded areas around the solid lines represent the 90th and 50th percentile ranges, respectively. The green shaded areas are horizontal histograms for the estimate of the transition depths.

Case	δ_α (preseismic) [-]	δ_α (postseismic) [-]	δ_d [m]
(1)	0.217546	0.178917	0.001402
(2)	0.181823	0.156073	0.001378
(3)	0.379139	0.415743	0.004918
(4)	0.434136	0.493933	0.010953
(5)	0.224513	0.186335	0.001546

Table 4.2: Average viscosity and surface displacement errors for different inversion experiments. Metrics defined in eqs. (4.2) and (4.3)

The recovered effective viscosity profiles for each case are presented in Fig. 4.5. Values for the error metrics $\delta_\alpha(t)$ and δ_d (defined in Section 4.3) for the different cases are given in Table 4.2 and support the visual interpretation of Fig. 4.5.

Cases (1) and (2) yield effectively the same results, showing that essentially all the information content is in the postseismic period (i.e., where the timeseries has curvature). Between Cases (2) and (3), the uncertainty in the lower transition depth D_{deep} and the uncertainty in the postseismic effective viscosity at D_{deep} is increased, although still within reasonable bounds.

Case (4) represents a more “realistic” case of using geodetic data to estimate fault rheological parameters, where observations were only collected starting some time after the earthquake (e.g., because a network had to be deployed, or the network malfunctioned due to the earthquake and had to be repaired). Vertical GNSS data frequently suffers from higher noise levels and is thus ignored in this case. Furthermore, models of coseismic slip are always uncertain, and thus we vary the slip by $\approx 10\text{--}15\%$ following the log-normal distribution in Fig. 3.2. The recurrence time is assumed to be constant and perfectly known as previously. In this case, the average surface displacement errors increase significantly compared to reference case (1), by approximately an order of magnitude, to $\delta_d = 11.0$ mm, which is around the observational noise level of 10 mm. The larger error is likely due to significant differences in the state of stress immediately after the earthquake — for a power-law viscoelastic process, the initial condition influences the relaxation time and the amplitude of the decay. Changes in the relaxation time or decay amplitude, in turn, directly affect the surface displacement timeseries, which is the observable for the inversion. Despite the increased displacement error, the average posterior viscosity errors only increase moderately by a factor of 2–3 (see Table 4.2), and remain well below an order of magnitude with values of $\delta_\alpha = 0.43, 0.49$ for the preseismic and postseismic timesteps, respectively.




Figure 4.5: Viscosity recovery in different test cases, explained in detail in Section 4.3.2. Shaded areas, lines, and histograms represent the same as in Fig. 4.4. In the leftmost panel, light blue and green still refer to the pre- and postseismic times, but all other colors now indicate the test case.

In case (5), we assume there is only historic information regarding the timing of the last significant earthquake on the megathrust. Consequently, the recurrence interval of the earthquake is uncertain and is a hyperparameter that we vary by $\approx 10\text{--}15\%$ following the log-normal distribution in Fig. 3.2. Here, the error metrics are $\delta_\alpha = 0.22, 0.19$ (for the preseismic and postseismic timesteps, respectively) and $\delta_d = 1.5$ mm, barely different compared to case (1). Compared to case (4), the variability in the recurrence time interval has a much smaller negative effect on the fitting of the surface observations as well as the recovery of the viscosity profile. We interpret this difference to be due to the recurrence time interval being much larger than the relaxation timescale. Consequently, any error in the timing of the previous earthquake does not significantly affect the slope or curvature of the displacement timeseries, thereby making the MCMC sampler effectively insensitive to the hyperparameter variability.

All the presented test cases in this section share the rupture history (9.5 m of slip every 100 years), observer geometry (five stations between 200–500 km), and target viscosity structure (see Fig. 4.4). These model characteristics directly affect the target timeseries to be fit by the inverse model, e.g., through the amount of curvature present in the postseismic relaxation phase. We do not show here additional test cases of more observing stations, different magnitude earthquakes, and other viscosity structures, as we find that they do not significantly change the ability of our framework to recover the target effective viscosity structure, which is our main goal. The recovery is not significantly hindered by these changes because there are no model prediction errors in these test cases, i.e., all physical mechanisms on the forward model are also present in the inverse model. Therefore, the leftover uncertainties and correlations are dominated by effects of non-uniqueness when estimating processes at depth from the surface. We show the impact of adding viscoelasticity into the forward model to assess its influence on the recovery of rheological parameters in Section 4.3.5, but a full exploration of model prediction errors is beyond the scope of this dissertation.

4.3.3 Logarithmic Timeseries Fits

Since modeling postseismic surface displacement timeseries using logarithmic functions is common in geodetic analyses, we perform such a fit to the modeled synthetic surface observations from the reference case (which uses power-law viscosity on the plate interface) to assess its applicability. We use the following

functional form:

$$d(t) = p_1 \log(p_2(1+t)) + p_3 + p_4 t$$

where d is the surface displacement in meters, t is time in years, and the p_i are the parameters to be estimated. In particular, p_2 is directly related to the logarithm's decay rate. The fit is performed using SciPy's `curve_fit` function (Virtanen et al., 2020), and uses the posterior ensemble of surface displacements (i.e., after the completed inversion process) as the observations (including their standard deviations). We vary the data available to the fit between 1 year, 3 years, and the entire 11 years after the earthquake.

Fig. 4.6 shows the fit to the posterior ensemble. The logarithmic functional form is able to capture the observed transient patterns extremely well. The decay rate parameters p_2 vary between 0.004–0.1 and 0.0–2.355 for the horizontal and vertical components, respectively, showing that even in our simple model, a single source process does not yield the same decay rates at different observer locations.

4.3.4 Inversions Assuming Rate-dependent Friction

We assess whether the choice of a power-law rheology formulation is too restrictive, i.e., whether it actually includes frictional behavior as one of its endmember models as described in Section 3.3 for $n \rightarrow \infty$. To that end, we take the surface observations produced by reference case (1), but assume rate-dependent friction on the creeping fault patches instead of the power-law rheology used to create the target timeseries. We can therefore estimate the depth-dependent values of α_h instead of the pairs α_{eff}, n in the parameter set θ , see eq. (4.1). Just as α_n in Section 3.3, α_h is assumed to vary linearly in logarithmic space with the knickpoints $D_{\text{middle}}, D_{\text{deep}}$. The inversion converges nicely, and Fig. 4.7 shows a corner plot of the posterior density probability function. Both the transition depths and the optimal frictional parameters are well recovered. Indeed, Fig. 4.8 shows that forward models using the estimated rate-dependent parameters reproduce the observed surface displacements very well, with only minor imperfections occurring early on in the postseismic period. The average displacement error is $\delta_d = 3.0$ mm and therefore well below the observational noise level of 10 mm. This analysis shows that our power-law model can be adequately reproduced with rate-dependent friction.




Figure 4.6: Logarithmic functional fits to the surface displacement timeseries. The different colored lines are the ensemble mean of the posterior modeled surface displacement timeseries. They almost cover the black line, which is the functional fit calculated from the entire data timespan. The gray solid and dashed lines correspond to fits using 3 and 1 years of data, respectively.

4.3.5 Effects of Neglecting a Viscoelastic Bulk

Our framework and all previous test cases assume that viscoelastic flow in the half-space bulk can be sufficiently approximated by an appropriate shearing motion on the plate interface; see Sections 3.1 and 3.3, as well as Mallick et al. (2022). To test the extent to which this assumption is valid, we build a slightly modified forward model of our two-dimensional subduction zone, which does contain viscoelastic flow in the bulk. The changes, shown in Fig. 4.9, are as follows: (1) the interface between 35 and 50 km depth exhibits rate-strengthening frictional behavior, (2) the continental crust includes a small “cold nose” above the frictional regime, where it is thicker than usual, such that it touches the frictional interface, (3) the interface below 50 km obeys a no-slip boundary condition with the viscoelastic bulk (no frictional interface or viscous channel), (4) the bulk obeys linear Maxwell viscoelasticity, and (5) the domain is purely elastic below 200 km depth. To aid our interpretation, we use depth-dependent viscosity values for the bulk below the continental plate, derived from geodetic data from the Northern Japanese subduction zone (Freed et al.,

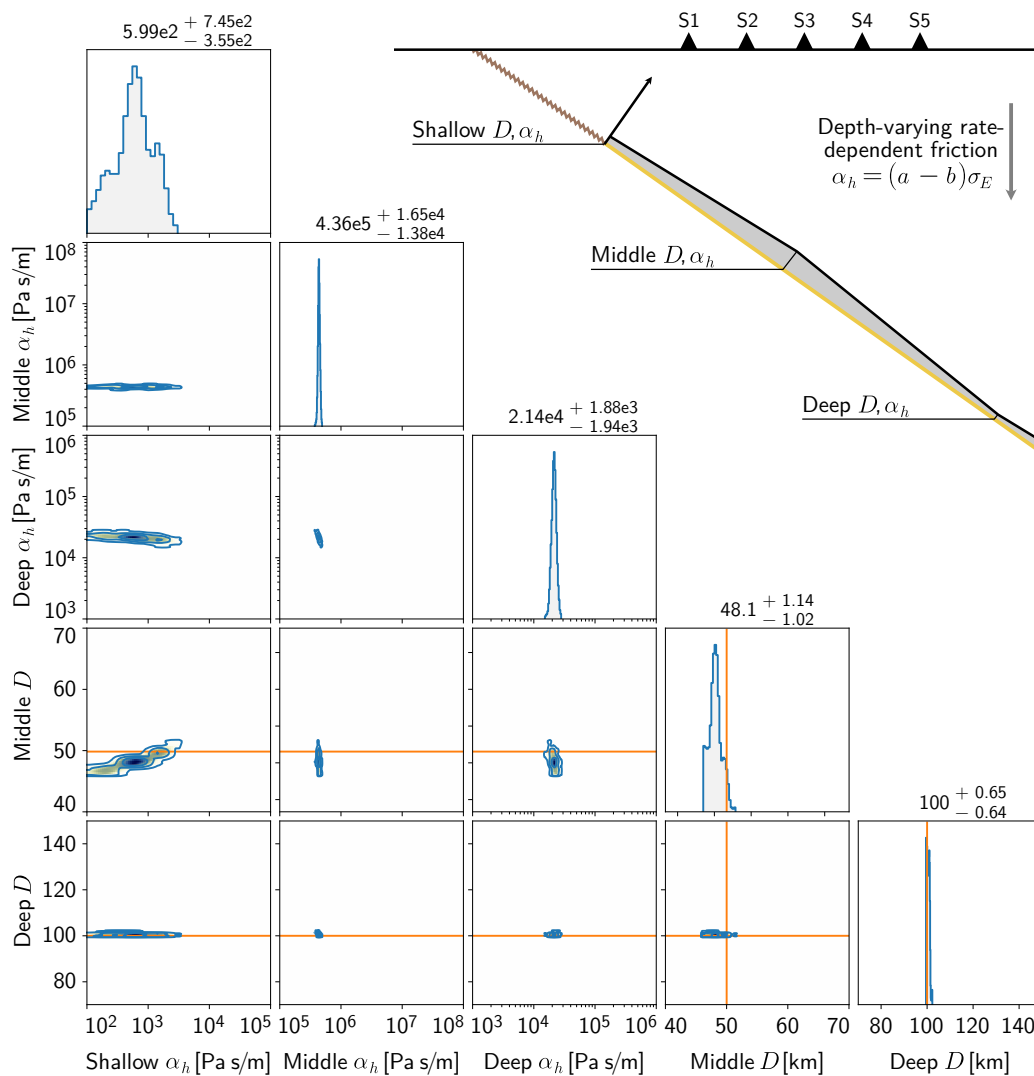


Figure 4.7: Corner plot of the posterior probability density for the rate-dependent rheology inversion. It follows the same notation as Fig. 4.3, with the rate-dependent frictional parameter α_h substituting the power-law parameters α_n and n . No orange lines are present for the frictional parameter plots since there is no straightforward conversion from the “true” effective viscosity used to generate the timeseries into the rate-dependent formulation.




Figure 4.8: Surface displacement observations (black dots) and fits (pink lines) using rate-dependent parameters. 960 samples are taken from the posterior probability density function, displacement timeseries are calculated, and then plotted over the noisy observations. Within uncertainties, the sampled timeseries all overlap.

2017), which is broadly compatible to our simplified fault geometry. The viscosity below the oceanic plate is set to 10^{20} Pa s everywhere. We then compute the synthetic, target surface observation timeseries after the spin-up period similar to the previous test cases, and repeat the generation of target timeseries for different values of the rate-dependent frictional parameter $\alpha_h = (a - b)\sigma_E$ to gain more insight. These target datasets are given to our probabilistic inverse model from Section 4.3.4, which assumes rate-strengthening friction on the interface, but crucially still does not model viscoelastic effects. The estimated parameters are the same as in Section 4.3.4, and we take $\alpha_{h,\text{shallow}}$ to be the approximation of the frictional parameter in the input dataset (between 35–50 km). The inverse model therefore aims to fit data generated by a model *containing* a viscoelastic bulk, using forward models which *do not contain* a viscoelastic bulk.

Fig. 4.10 shows the target synthetic data as well as the model predictions after estimating the best-fitting depth-dependent frictional parameter for one of the target datasets where $\alpha_h = 1$ MPa. Overall, the forward models without viscoelastic effects reproduce, to first order, the target data generated with a viscoelastic bulk.




Figure 4.9: Fault setup of the tests with forward models containing a viscoelastic bulk. The downgoing slab is modeled by two interfaces bounding an effective elastic plate of constant thickness H . Above it, the continental crust is modeled as an elastic layer and a “cold nose” touching the slab up to a depth of 50 km. The upper and lower interfaces experience reverse and normal motion, respectively. The upper interface is furthermore assumed fully coupled (i.e., locked) from the trench down to a depth of 35 km (brown line). Between 35–50 km, the interface is creeping following rate-strengthening friction with a constant value of α_h (yellow line). Every other interface obeys a no-slip boundary condition with the viscous bulk (purple lines). The viscosity below the oceanic plate η_o is constant, whereas the one below the continental plate η_c is depth-dependent. Everything else as in Fig. 4.1.

The horizontal postseismic phase is fit particularly well (and is fit even better when neglecting vertical observations during the probabilistic inversion). The average displacement error is $\delta_d = 57.8$ mm, well above the observational uncertainty of 10 mm. However, most of this misfit appears to be driven by the vertical postseismic signal.

We investigate whether the obvious misfit, especially compared to all the test cases in the previous sections, inhibits the inversion routine to recover the target rate-strengthening frictional parameter used to generate the input data. Fig. 4.11 shows the positive correlation between the target (input) and estimated (output) rate-strengthening parameter $\alpha_h = (a - b)\sigma_E$, as well as their average displacement misfits, for five test cases. Although there is a systematic offset below $\alpha_h \approx 3$ MPa, the errors amount to less than half an order of magnitude, and do not increase with the displacement error δ_d . In fact, at the upper end of our tested target α_h values, the error in frictional properties appears to decrease with increasing displacement errors, however it is beyond the scope of this section to investigate if this trend holds for higher α_h values.




Figure 4.10: Surface displacement observations (black dots) from a forward model including a shallow frictional interface up to 50 km depth and a viscoelastic bulk. In pink to green lines, the best-fitting surface displacements from forward models only containing a frictional interface up to 150 km depth and no viscoelastic bulk. Each fitted line is a sample from the posterior probability density function of a probabilistic inversion of the tested observations. The preseismic and the horizontal postseismic observations are well-matched by the simpler forward model.

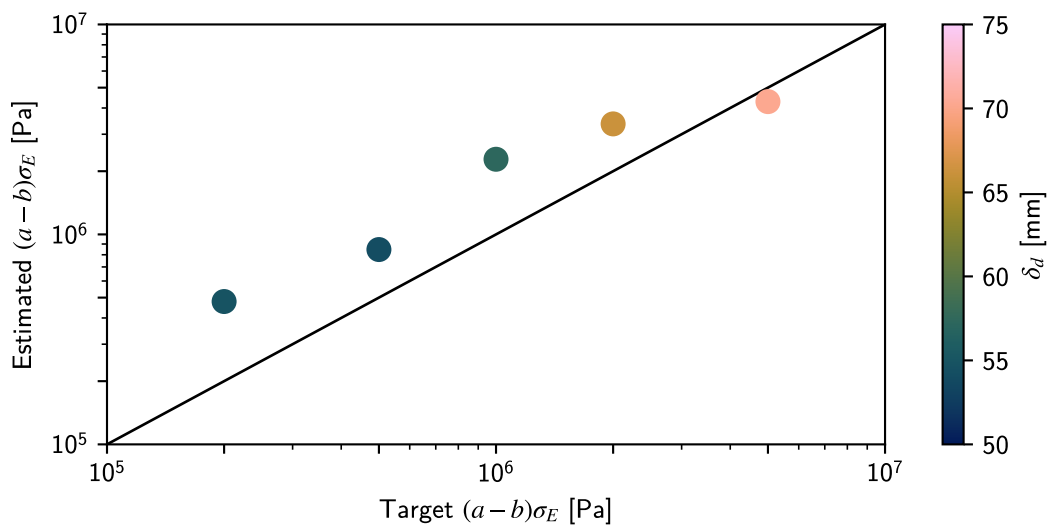


Figure 4.11: Correlation between the target, synthetic (input, horizontal axis) and the estimated (output, vertical axis) rate-strengthening parameter $\alpha_h = (a - b)\sigma_E$ for the test case where the forward model includes viscoelastic bulk flow, but the inverse model does not. The diagonal line represents perfect 1:1 correlation. The displacement error δ_d for each case is given in the color of the markers.

4.4 Implications

Our reference case presented in Section 4.3.1 concerns the inversion of surface displacements generated by sequences of earthquakes and aseismic slip occurring on a synthetic fault interface governed by a depth-variable power-law rheology. The posterior sample covariance of our inversion of the rheological parameters shows that the ability to distinguish between friction-dominated and power-law-dominated regimes (as determined by the power-law exponent n) at depth is limited (Fig. 4.3). While the estimates lie within two standard deviations of the values used to generate the target synthetic data, there remain large uncertainties and strong correlations between parameters, hindering the unique recovery of α_n and n .

These conclusions are derived from the exploration of our 2D model using synthetic data, and come with the expected limitations of model simplicity, differences between 2D and 3D displacement fields, and explored parameter space. However, these results may indicate the limited potential of using land-based geodetic observations to distinguish between different creep mechanisms in subduction zones, as suggested by various studies. For example, Qiu et al. (2018) and Tang et al. (2019) analyzed postseismic surface displacement timeseries for the Sumatran megathrust and the Taiwan orogenic belt, respectively, and found that the stress-strain rate relationship cannot be fit using a power-law rheology of a single exponent n , leading them to discard nonlinear viscosity as a candidate rheology for their study region entirely. We have shown in Section 4.3.1 that a depth-varying power-law exponent n can produce a more complex behavior at the surface (observationally similar to one produced using rate-dependent friction, see 4.3.4), and would therefore caution against discarding power-law models based on such analysis. It may be possible, however, to reconcile their preferred model of a power-law Burgers rheology with that of a depth-variable pure power-law rheology, if the Kelvin solid within the Burgers body represents the variations of the parameters with depth.

Ingleby and Wright (2017) also discuss the ability to fit postseismic surface velocity timeseries using a variety of models (including power-law creep, frictional afterslip, and a Burgers body), suggesting afterslip or high- n power-law creep are the dominant driving mechanisms in their data. However, by assuming spatially uniform model parameters, one may be significantly underestimating the effects that depth-variable parameters can have on such velocity observations, which undermines their conclusion. Periollat et al. (2022) show that uniform models (in depth and time) are inadequate for modeling the postseismic surface displacements following the 2011

M_w 9.1 Tohoku-oki earthquake, but rather than relaxing this assumption, they invoke a combination of transient brittle creep and viscoelastic relaxation to fit the data.

Other studies neglect the impact of the strong α_n - n correlation and importance of depth-dependent parameters by attempting to decouple viscoelastic relaxation from afterslip completely, fitting one of the two processes to surface observations before using the residuals to fit for the other (Freed et al., 2017; Peña et al., 2019; Peña et al., 2020). It is thus unclear whether their separation of the contributions is purely a result of their individual model assumptions, or physical. A Bayesian approach such as the one presented by Fukuda and Johnson (2021) which fits different contributions jointly may be necessary to disentangle the effects of different processes. Lastly, compared to our framework where we estimate the stress state before an earthquake by spinning up the entire earthquake cycle until stationary behavior is reached, it is unclear to what extent the initial conditions assumed by many of the aforementioned studies affect their results (e.g., Govers et al., 2018).

Even though our reference case shows that the recovery of the power-law parameters α_n and n is non-unique, converting the recovered parameters to an effective viscosity α_{eff} using the simulated interface velocities right before and after the applied earthquake shows that the inversion scheme can successfully infer the depth-varying profile of effective viscosity (Fig. 4.4b–c). We recognize that the estimation of an effective viscosity over time is indirectly a regularized inversion for a spatiotemporal characteristic displacement pattern. In other words, regardless of the constitutive law (since our power-law model can approximate a wide range of phenomenological behavior), the solver is able to match the surface observations as long as the slip timeseries of the fault patches is sufficiently well recovered using the parameters of the forward model. The equivalence between estimating effective viscosity and a characteristic spatiotemporal slip pattern is to be expected since our model involves a purely elastic bulk and, as a consequence, the fault slip velocity is uniquely defined by the stress state. Estimation of slip over time on discretized fault patches at depth is inherently non-unique, and is often regularized in space using standard smoothness operators (Harris and Segall, 1987; Fukuda and Johnson, 2008) and filtered in time using Kalman filters (Segall and Matthews, 1997; Bekaert et al., 2016) or Principal/Independent Component Analysis (Kositsky and Avouac, 2010; Gualandi et al., 2016). It is therefore natural to interpret the rheological constitutive law and assumed shape of the depth-dependence of the fault properties as a physics-based

regularization operator in our framework.

The inverse methodology in our study allows for strength profiles with no constraints on the effective viscosity depth gradients other than the goodness of fit to the surface deformation timeseries, and that the effective viscosity varies logarithmically with depth (consistent with the Arrhenius-like temperature dependence of effective viscosity, Hirth, 2002; Hirth and Kohlstedt, 2004). Relative effective viscosity as a function of depth, including any positive or negative gradients, are therefore readily extracted from our inversion results for further study. We highlight here that a key component of our framework is that we do not require two independent frameworks to model frictional afterslip and viscous creep, which are commonly required by other models (e.g., Hu et al., 2016; Freed et al., 2017; Agata et al., 2019; Muto et al., 2019; Fukuda and Johnson, 2021).

In our forward models assuming depth-dependent power-law rheology, we find that our surface displacement timeseries appear similar to logarithmic functions. In Section 4.3.3, we show that the synthetic surface observations can be fully explained by a functional model of a linear trend and logarithm, even though the fitted parameters change significantly when different input timespans are considered. The ability to fit surface observations with a logarithmic functional form has been used to argue that friction is a dominant shear resistance mechanism at fault interfaces (e.g., Ingleby and Wright, 2017; Tang et al., 2019; Periollat et al., 2022). Since the synthetic fault in our reference case is not described using a frictional model but still produces logarithm-like surface transients, we caution against inferring a rheological model from functional fitting. Similarly, we are able to fit the target observations from the synthetic reference case assuming rate-dependent friction on the interface, further suggesting that the ability to fit timeseries using a frictional model is not a sufficient condition to infer actual frictional behavior at depth (see Section 4.3.4).

In Section 4.3.5, we test our assumption that all significant bulk motion can be approximated by an equivalent motion on the fault interface, and that the recovered rheological parameters are not significantly impacted. For this experiment, we create target datasets that include effects of a depth-dependent viscoelastic bulk around a rate-strengthening frictional interface of variable $\alpha_h = (a - b)\sigma_E$. These datasets are then given to our probabilistic inversion framework, which does not contain viscoelastic effects, to estimate the best-fitting depth-dependent rate-dependent frictional parameters. We find that while the postseismic vertical motion exhibits measurable misfit between the target data and posterior sample timeseries, the postseis-

mic horizontal and preseismic motion are well recovered. The recovered frictional properties α_h of the interface are positively correlated with the input, target values, and match to within half an order of magnitude. Furthermore, the error in α_h does not appear to be correlated with the displacement timeseries misfit (at least for the range of α_h values tested here). Our test therefore suggest that the model assumption of representing all motion in the halfspace as slip on the fault interface is viable, and produces adequate estimates of rate-strengthening frictional properties even with significant viscoelastic bulk motion.

Finally, the quality of recorded earthquake history as it concerns date, magnitude, and location, varies significantly between different areas of the globe. While such a dataset is required to apply our method to current geodetic observations, we show in Section 4.3.2 that realistic uncertainties do not significantly affect the quality of the recovered viscosity profile. Indeed, the large number of samples in the MCMC process makes the inversion robust enough to still recover a good approximation to the effective viscosity profile at depth when the surface observations are significantly affected by errors in the assumed coseismic slip. Similarly, since most transient afterslip occurs in the months to decades after an earthquake, it makes sense that our inversion is insensitive to variations in the assumed recurrence time interval for earthquakes far in the past. The relative robustness of our inversion results to uncertainties in the coseismic slip and the recurrence time interval suggests that our method should perform well when applied to real data.

4.5 Conclusion

In this chapter, we have applied a probabilistic inversion framework to synthetic test cases of a simplified 2D subduction zone exhibiting cyclic earthquake behavior. Our model setup involves a shallow asperity, rupturing only coseismically, above a variable-strength interplate shear zone exhibiting either power-law viscous or rate-dependent frictional behavior. Using surface displacement observations at discrete observer locations as input, we are able to recover the transition depths in the viscosity structure as well as first order constraints of the effective viscosity with depth. Our results further suggest that uncertainties in the recurrence times and slip amounts of observed or inferred earthquakes may reduce the certainty with which we infer rheological parameters, but do not add significant bias. Lastly, we show that the ability to fit surface displacement timeseries using logarithmic functions or mechanical models of postseismic relaxation assuming rate-dependent friction on the plate interface is not a sufficient criterion to infer that frictional behavior is indeed

the main driving mechanism of shearing at depth.

*Chapter 5***ESTIMATING THE FRICTIONAL PARAMETERS OF THE
NORTHERN JAPANESE MEGATHRUST USING
EARTHQUAKE CYCLE SIMULATIONS****5.1 Introduction**

The mechanical behavior of rocks on and around plate interfaces at depth is an important factor in our understanding of how subduction zones are formed and sustained. Rheological models describing the stress-strain relationship of fault material furthermore directly inform the seismic hazard of a megathrust because of the rheologies' effect on stress accumulation and diffusion (Stern, 2002). However, laboratory-derived parameter values for such models are not easily scaled up to entire plate boundaries (e.g., Handy, 1994; Yamashita et al., 2015; Fagereng and Beall, 2021; Bercovici et al., 2023), and as a consequence, geodetically-derived parameter values may differ significantly (e.g., Marone, 1998; Thomas et al., 2017; Mallick et al., 2022). To estimate rheological parameters on regional scales, postseismic transient deformation following large megathrust earthquakes provides a natural experiment ripe for analysis, since the observed surface motion is the direct effect of the stress evolution in the subsurface, which in turn is dependent on the rheological characteristics of the fault zone. One of the geodetically best instrumented subduction zones in the world is the Northern Japanese subduction zone, with more than one thousand GNSS stations on land, dozens of observing stations on the seafloor, and an extensively documented historic seismic record. In particular, postseismic motion from the 11 March 2011 M_w 9.1 Tohoku-oki earthquake has been extensively studied, since it was not only well observed, but due to its large magnitude, the associated tectonic signal is also easily extracted from raw observations, and the postseismic period is still ongoing.

Given the non-uniqueness of geodetic inferences of complex behavior at depth given only surface observations, however, different rheological parameter values and even different types of rheological models all seem to explain the postseismic transient deformation to a reasonable degree (e.g., Sun et al., 2014; Freed et al., 2017; Perfettini et al., 2018; Agata et al., 2019; Muto et al., 2019; Fukuda and Johnson, 2021). For example, within the cited studies, the assumed subduction zone structure, the

modeled physical processes, and the relative importance of viscoelastic deformation on the timeseries, all vary significantly (described in more detail in Sections 3.1 and 5.4). What is common to studies of postseismic deformation, however, is that they do not attempt to model the complete earthquake cycle, i.e., the interseismic period, and the longterm accumulation of slip and stress. Taking into consideration all the available observations may therefore constitute a path towards differentiating between the necessity of various model elements (Hearn and Thatcher, 2015; Mallick et al., 2022).

We address the question of inferability of rheological properties and structures from geodetic data by inverting observations from before and after the 2011 Tohoku-oki earthquake. Using the framework presented in Chapter 3, we simulate the entire earthquake cycle in Northern Japan for thousands of years until reaching time-invariance, and then use a probabilistic, Bayesian Markov-chain Monte-Carlo solver to estimate best-fitting, constant frictional parameters of the plate interface. Our primary goal is to probe the degree to which interseismic data is relevant for the inference of rheological properties and to provide a perspective on the geodetic evidence for different physical processes assumed to happen at depth, thereby demonstrating that inversions using simplified earthquake cycle models are computationally feasible today. In addition, our model readily provides cumulative slip deficit as well as kinematic plate coupling values for the interseismic period, two quantities relevant in seismic hazard assessments that have been shown to differ in synthetic numerical simulations (e.g., Hetland and Simons, 2010). Our secondary goal is therefore to assess the difference between these two quantities in the Northern Japanese subduction zone.

Section 5.2 describes the inputs necessary for our framework, focusing on the 3D fault geometry as well as the assumed earthquake rupture catalog. We then present our inversion results in Section 5.3, both in terms of pre- and postseismic surface displacement fits as well as the inferred rheological parameters. Finally, we discuss the implied preseismic kinematic coupling state, the interpretability of the recovered effective viscosity of the plate interface, and the importance of viscous bulk flow in Section 5.4.

5.2 Model Setup

This section presents the inputs necessary for the probabilistic inversion framework presented in Chapter 3. We first present the discretized fault geometry (including

the size, shape, location, and rupture history of the asperities) used in the forward model. We then describe the observations used in the inverse model, as well as the rheological parameters being fitted.

5.2.1 Fault Model

Our model requires both a discretized fault interface, split into creeping and locked regions, as well as an assumed rupture history.

Asperity Catalog

We start from the collection of asperities identified by Kanda et al. (2013), which includes five minor ones (in the regions of Fukushima, Miyagi, Sanriku, Tokachi, and Nemuro) as well as two major ones (the Tohoku and Hokkaido M_w 9+ source regions, see also Nanayama et al., 2003). These asperities were modeled using elliptical shapes, span characteristic slip magnitudes from 3 m to over 100 m, and were assumed to have recurrence times derived from ruptures before the year 2000.

Kanda et al. (2013) were only attempting to model interseismic velocity fields derived from pre-2000 data, and so we cannot use the same asperity catalog and rupture history in our framework that includes data from 1996–2024, because the expected occurrence times of their future ruptures do not match up precisely enough with real earthquake occurrences. For example, it would expect the Miyagi asperity to rupture with approximately M_w 7.5 in 2018, the Tokachi asperity in 2002 (which in reality occurred in 2003), or the Nemuro asperity in 2013 (the only M_w 7+, however, being in 2004).

Furthermore, the dominant nature of the 11 March 2011 M_w 9.1 Tohoku-oki earthquake on the GNSS observations means that in order to adequately fit the postseismic data, our model cannot have a fully locked region where significant afterslip occurred. This observational constraint means that the Fukushima asperity as used by Kanda et al. (2013) has to be moved or eliminated.

We address the misalignments of the asperity catalog in the following ways. First, given the afterslip distribution of the Tohoku-oki main shock as well as the coseismic slip of the same-day M_w 7.9 aftershock (see, e.g., Simons et al., 2011), we move the Fukushima asperity south and rename it the “Ibaraki” asperity (keeping the original recurrence interval and characteristic slip magnitude). Second, we set the last rupture at the Miyagi asperity to coincide with the Tohoku-oki main shock, which indeed could have had slip inside the Miyagi asperity (see the published finite fault

models referenced in Appendix C.1). Third, since the postseismic deformation of the 2004 Nemuro asperity earthquake is much smaller than the 2003 Tokachi asperity earthquake and therefore does not have a significant influence on the land-based GNSS stations, we assume they ruptured jointly in 2003.

An additional issue is the computational cost of our framework in 3D, which is something we have to consider in all aspect of our setup. In our tests, we find that the runtime of forward-integrating the model increases linearly with the number of individual earthquakes present in a full cycle. The dates of last occurrence and the recurrence times from Kanda et al. (2013) imply an earthquake cycle length of 9000 years and a total of 1050 earthquakes (including occasional joint occurrences when downsampled to years). By shifting in time the last occurrences as described in the previous paragraph, we create more joint ruptures, and are able to slightly decrease the total number of earthquakes to 990. As we find in Section 4.4, our model is only weakly sensitive to uncertainties in recurrence times. Applying this insight to our dataset means that we can slightly modify the assumed recurrence times without expecting significant changes to the model predictions during our observational period. We therefore shift the asperity recurrence times to the closest multiple of 40 or 50 years and set the recurrence time of the Tohoku asperity to 1000 years, which has the desired effect of reducing the total number of earthquakes per cycle down to 140, a decrease of about 87%. Associated with these shifts is a reduction of the earthquake cycle length to 2000 years. In our early tests, and supported by block motion studies of Northern Japan (e.g., Loveless and Meade, 2010), we find that the relative motion between the islands of Honshu and Hokkaido can lead to significant misfits in our model results, because our ESPM formulation assumes the overriding plate to be fully rigid. We therefore restrict our study region to Northern Honshu, and therefore remove the Tokachi, Nemuro, and Hokkaido asperities altogether. This final step reduces the number of individual earthquakes to 100 (at an unchanged cycle length of 2000 years), and the number of asperities to 4.

The resulting asperity earthquake catalog is shown in Table 5.1 and visualized in Fig. 5.1. To calculate the characteristic slip on each asperity, we assume a nominal plate convergence rate of 9.5 cm/a, ensuring that the plate convergence is fully recovered after a full cycle. The magnitude of each characteristic earthquake is then simply a function of the area of the asperity. Note that during the observational period, our model uses realistic slip distributions rather than the uniform characteristic slip (see below). Furthermore, for the numerical integration, we use the convergence

	Asperity	T_{rec} [a]	s_0 [m]	Last Observed	Last Assumed
1	Ibaraki	80	7.6	1938-11-05 M_w 7.8	2011-03-11
2	Miyagi	40	3.8	1978-06-12 M_w 7.7	2011-03-11
3	Sanriku	40	3.8	1994-12-28 M_w 7.8	1994-12-28
4	Tohoku	1000	95	2011-03-11 M_w 9.1	2011-03-11

Table 5.1: Recurrence interval T_{rec} , characteristic slip magnitude s_0 , and last observed and assumed dates of rupture for the 4 asperities identified on the megathrust offshore Northern Honshu.




Figure 5.1: Visualization of Table 5.1, using the assumed last rupture dates. Colors denote the different asperities as shown in subsequent figures.

rates derived from the plate motion’s Euler pole to calculate the realistic slip in each asperity and creeping patch instead of the nominal slip from Table 5.1. Lastly, due to computational considerations, we do not randomly sample recurrence time or slip magnitude realizations as in Section 4.3.2.

Geometry

The target interface we model is the convergent plate boundary between the down-going Pacific and the overriding Eurasian plate. The shape of the interface is well constrained using seismic data merged by Hayes et al. (2018) into their Slab2 product. To convert Slab2 into the format necessary for our framework, we need to discretize the interface into triangular patches, separated into the creeping fault and the locked asperity patches.

To augment the asperities in the rupture catalog with shapes, locations, and size, we rely on previous work determining the coseismic slip extents of historic earthquakes inside the modeled asperities. We first create a uniformly discretized, triangular surface mesh derived from Slab2’s Kuril arc isodepth contours using Coreform Cubit 2023.11 (Coreform LLC, 2023). We then map published slip models onto this mesh

(see Appendix C.1 for a full list of the slip models used). If there are multiple models available, we calculate the average slip for each earthquake and on each fault patch. We then define asperities as the contiguous regions where the average coseismic slip is at least 90% of the maximum average slip amount.

With the asperities determined, we are able to refine the uniform input mesh around the asperities, and coarsen the mesh far away from the asperities, in order to decrease the computational cost when integrating our model forward. Fig. 5.2 shows the resulting mesh and asperity shapes for a mesh refined down to a minimum characteristic patch length scale of 25 km for the region offshore Honshu. For computational reasons, we create a coarser mesh with a minimum characteristic patch length scale of 45 km, shown in Fig. 5.3. We verify in Appendix C.2 using forward models that the inversions run with this coarse mesh produce similar results. Finally, we compute the displacement and stress kernels from the triangular mesh using Thompson et al. (2023).

Coseismic Slip

Directly applying the characteristic slip for the 2011 M_w 9.1 Tohoku-oki earthquake uniformly across the identified Tohoku asperity would yield a coseismic and post-seismic response unlikely to fit our GNSS observations well. While the characteristic, nominal slip of 95 m would recover the accumulated plate motion in our model, we do not expect this simple expectation to hold in nature. Earthquake “cycles” are an idealized view at best, and earthquake sequences are never expected to precisely recover the accumulated convergence at any given time. Each earthquake, even if recurring on the same asperity with some regularity, may have more or less slip than on average required by plate motion. In fact, we know from the published coseismic slip models that a slip of 95 m was only reached over a small area, if at all. We mitigate the negative effects from this by instead directly using the averaged slip model already computed to define the Tohoku asperity (see above), upscaled to preserve the event magnitude. To not affect our goal of modeling the entire earthquake cycle, where all deformation is recovered after one full cycle, we only use this modified slip distribution during the observational period.

Our forward model relies on calculating stress changes based on the applied coseismic slip. If there is a sharp boundary between where we impose coseismic slip (in the asperities) and where slip has to naturally follow by creep (outside the asperities), we generate stress singularities which are both unphysical and introduce




Figure 5.2: Map view of the 3D subduction zone interface between the incoming Pacific plate from the east and the overriding Eurasian plate in the west. The triangular discretized fault patches are shaded by depth, and the asperities are marked by different colors (introduced in Fig. 5.1). White triangles denote the locations of GNSS stations processed for the inversion. Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022).




Figure 5.3: Same as Fig. 5.2, but using a coarser approximation of the interface to reduce the computational cost of the forward model.

numerical instabilities. To prevent this issue, we taper the slip into the creeping regions (similar to, e.g., Hetland et al., 2010; Kanda et al., 2013).

5.2.2 Observations

The inverse model in our framework is constrained by timeseries of observed surface displacements. These observations have to be preprocessed in order to contain only signals of the physical processes being modeled; otherwise, there would be large, systematic misfits. We perform this separation of signals using the DISSTANS software presented in Chapter 2 (Köhne et al., 2023).

The Geospatial Information Authority of Japan (GSI) distributes raw observation files (i.e., in the RINEX format) for approximately 1500 continuously-operating

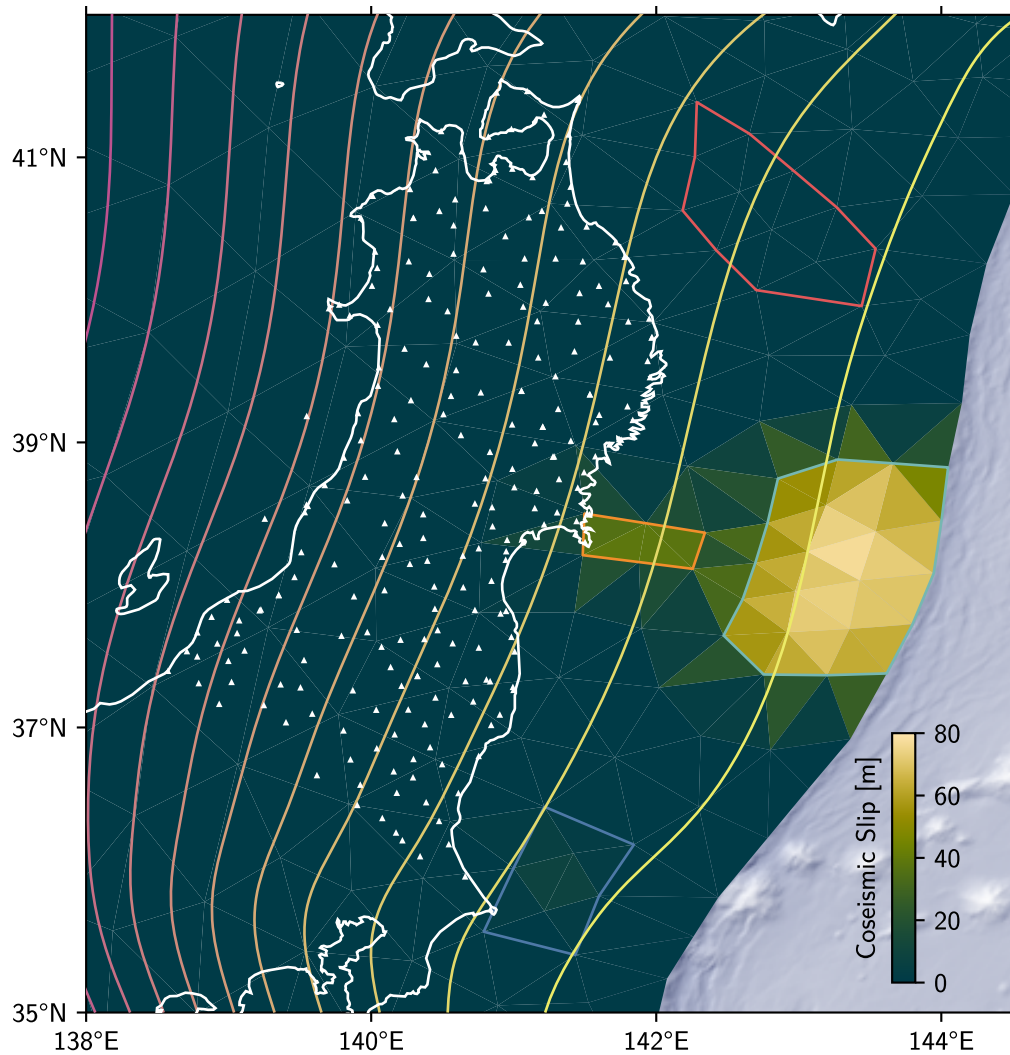


Figure 5.4: Coseismic slip distribution for the 2011 Tohoku-oki earthquake, including in adjacent asperities. The asperities included in the inversion framework are shown by the colored outlines (same colors as in Fig. 5.1). Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). Yellow to pink lines show the plate interface at 20 km depth intervals (Hayes et al., 2018), including the plate boundary at the trench (Bird, 2003). White triangles denote the locations of GNSS stations processed for the inversion.

GNSS stations on the Japanese islands. We start by obtaining the daily position timeseries of these stations (Moore, 2024) processed using JPL's GipsyX software (Bertiger et al., 2020). Using DISSTANS, we first apply quality metrics, filter out outliers, and remove the common mode error from the preprocessed timeseries. We then run multiple iterations of fitting simple secular and seasonal functional models to find step offsets in the data induced by maintenance events or earthquakes. We also include logarithmic models of varying timescales to model the large signals of the 2003 M_w 8.2 Tokachi-oki and 2011 M_w 9.1 Tohoku-oki earthquakes to improve the recovery of steps.

Once the majority of steps detected in the data are included in the functional model (based on visual inspection of the root-mean-square residuals), we furthermore add transient models to the fitting process. The first is an overcomplete dictionary of splines, spatially regularized, to capture any transient effects such as volcanic unrest, slow slip events, or unmodeled step offsets. The second is a locally regularized sinusoidal model of seasonally varying amplitude, intended to capture the yearly fluctuations in the hydrological cycle. The collection of models (secular, seasonal, step offsets, postseismic relaxation, and transients) is then fit iteratively using the parallelized mixed-regularization solver routines of DISSTANS.

The hyperparameters required by the solver that best restrict the overfitting of motion using the dictionary of splines while still being able to fit the data well are found using a grid search. Since simply minimizing the RMS of the residual timeseries at each station will lead to overfitting, we define an alternative metric to minimize. Based on the validation done in Appendix B.3, the metric we use is the RMS of the residual between the modeled horizontal secular velocity field and the velocity field defined by an Euler pole fitted to the modeled secular velocities. This ensures that the secular component in the collection of models fitted is primarily capturing the regional plate motion, and that vice-versa, the transient model only includes local or small-scale regional signals. We also calculate the RMS of the residual between the modeled vertical secular motion and a best-fit linear ramp across the observed domain, representing the first-order behavior expected in a subduction zone. We find that while the two RMS metrics are not perfectly correlated, they do share a joint minimum in hyperparameter space; we choose this hyperparameter combination moving forward. Fig. 5.5 shows a visual representation of the hyperparameter combinations tested and their respective RMS residual metrics.

In the next step, we separate the model constituents associated with the subduction




Figure 5.5: Visualization of the grid search performed to optimize the DISSTANS post-processing hyperparameters that produce the input timeseries for the inverse model. All cases assume a base scalar L0 penalty factor (penalty of `spatialfit()`) of 100. The hyperparameters explored are a factor f to increase the penalty for the vertical component, the eps (ϵ) and scale (s) values for the `ReweightFunction`, and the `spatial_reweight_percentile` (q) of `spatialfit()`. For each hyperparameter combination, the root-mean-squared residuals between the modeled secular velocity field and smooth secular velocity fields (as described in Section 5.2.2) are shown as two triangular halves of the rectangular grid element. The color of each triangular patch corresponds to the value of the RMS residual. The optimal hyperparameter combination minimizing both horizontal and vertical residuals is highlighted by the black rectangle (including the corresponding RMS residual values).

zone (i.e., the processes modeled in the forward part of our framework) from all others. The constituents considered are the secular, linear plate motion, as well as the postseismic relaxation of the 2004 and 2011 earthquakes. (The coseismic offsets induced by these earthquakes are an input to our framework, not an observable to be fit.) We then perform an additional quality check by comparing the secular velocity at each station with the motion expected if the stations were tied to the Pacific plate. We remove stations whose secular velocity deviates by more than 100% of the entire plate convergence velocity.

Our model framework is not tied to any terrestrial reference frame, and rather assumes all motion to be relative to the stable overriding plate. Therefore, before we can use our postprocessed observations, we rotate our observations into the Eurasian reference frame. To take into account the relative motion of Northern Japan with respect to stable Eurasia and the effect of the locked asperities on the fault, we furthermore estimate a reference frame correction by iteratively fitting the predicted interseismic velocity field from our model setup (assuming an Euler-pole-derived plate convergence) to the observations. The result is a reference frame relative to “stable Northern Japan” and already informed by our fault setup.

Finally, since the information content of the GNSS observations relevant to our framework is longterm secular motion and transient relaxation after large earthquakes, we downsample the timeseries. By choosing logarithmically-spaced time intervals, we capture the early postseismic phases as well as longterm trends. For each interval, we fit a polynomial of degree 5 to the observations, and then extract the value of the best-fit polynomial at the center of the considered interval. The resulting timeseries contains the same tectonic signals as the original one, while reducing the number of data points by up to two orders of magnitude (depending on the length of the input timeseries).

5.2.3 Fitted Parameters

In our framework, we integrate forward the slip history of every creeping patch, implying that every patch needs to be fully-defined in terms of its rheological properties. While we assume that all patches have to follow the same rheological model, each patch can have its own defining parameters within that constitutive law. Estimating these parameters individually for each patch would, however, be computationally prohibitive in our MCMC framework, especially because of the inherent non-uniqueness of the recovered parameters. Instead, we opt to parametrize the vari-




Figure 5.6: Example timeseries decomposition for station 0031 for the timespan 1996–2024 in the East, North, and Up components. The top panel shows the input data (daily position solutions), with outliers removed. Using DISSTANS, the timeseries is split into the constituents modeled by our inversion framework (secular and postseismic motion) and those not modeled (coseismic, seasonal, and other transient signals such as volcanic or slow slip events). To preserve potential unmodeled signal, the residuals are not removed from the postprocessed dataset.

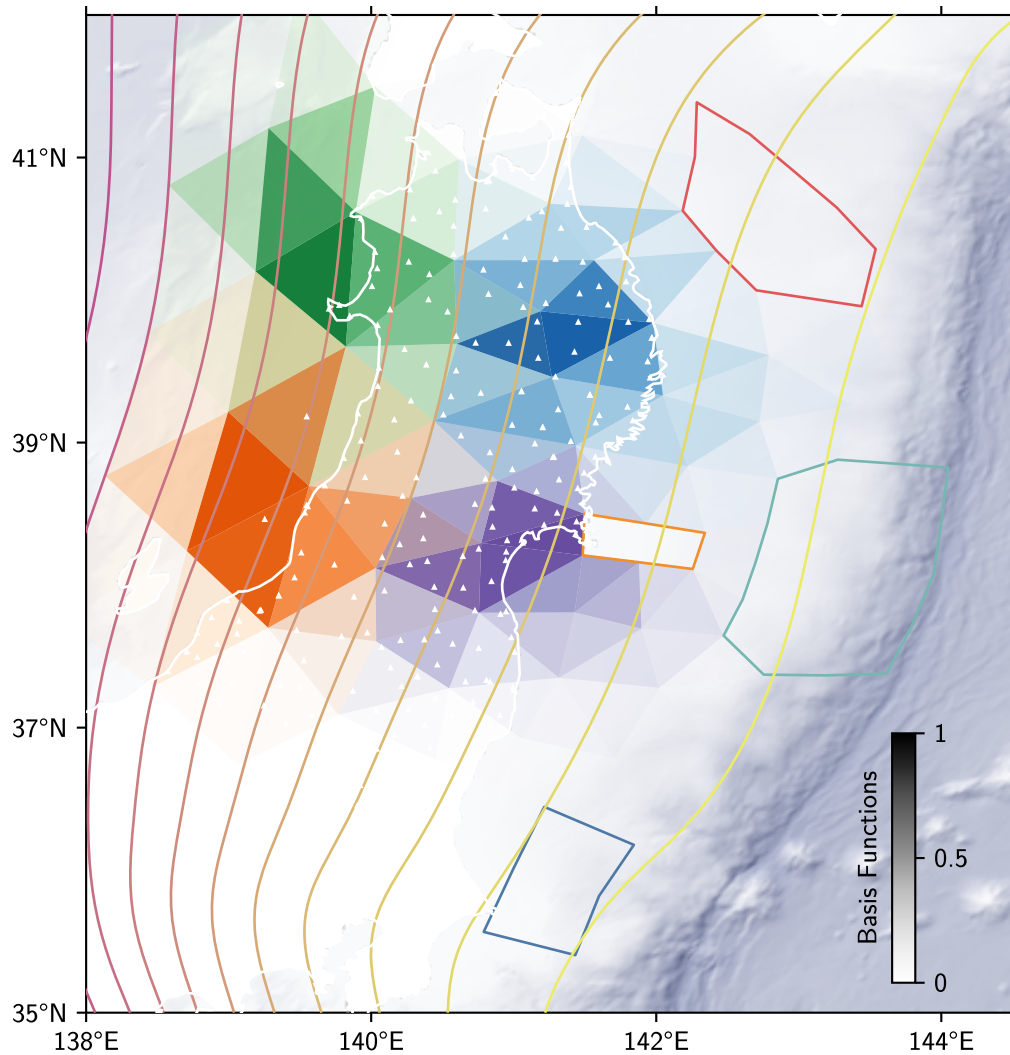


Figure 5.7: Values of four example basis functions in the down-dip and along-strike dimensions. The value of each basis is given as the darkness and transparency of the color, with the colors corresponding to the different bases, and lighter values implying smaller values. At each patch, the sum of all basis functions approximately sums to one, creating a suitable basis to parametrize a field of scalar values.

ation of the parameters throughout our model domain using piecewise linear basis functions (splines) on the 2D surface of the interface. The coordinate system for the basis functions is distance along strike (from the North) and depth, enabling our model to recover rheological variations in a spatially heterogeneous way. Fig. 5.7 shows the values of four example basis functions on the interface. Since the sole variable in a rate-dependent frictional framework is $\alpha_h = (a - b)\sigma_E$ [Pa], which can vary by multiple orders of magnitude, the specific fitted parameters are the $\log_{10} \alpha_h$ coefficients of the piecewise linear spanning functions.

5.2.4 Inversion Performance Metrics

This section presents the results of estimating spatially-variable rate-dependent frictional parameters on the Northern Japanese megathrust using our probabilistic, SEAS-simulating framework. To quantify the quality of the model fit to the data, we define the average surface displacement error δ_d similar to Section 4.3 as

$$\delta_d = \frac{1}{3N_t N_o} \sum_{i=1}^{N_t} \sum_{j=1}^{N_o} \sum_{c=E,N,U} \left(\frac{1}{N_s} \sum_{k=1}^{N_s} \left(d_{i,j,c,k}^{\text{pred}} - d_{i,j,c}^{\text{observed}} \right)^2 \right)^{1/2}. \quad (5.1)$$

When comparing different model runs, we furthermore need a metric to compare posterior distributions of the fault frictional parameter α_h on each simulated fault patch. Here, we resort to a scaled version of the Wasserstein 2-Distance (WD). Since the formal definition of the WD for empirical samples is somewhat unintuitive, we introduce the WD assuming two normal distributions $\mathcal{N}_i(\boldsymbol{\mu}_i, \mathbf{C}_i)$. In this case, it can be calculated as follows:

$$\text{WD} = \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 + \text{trace} \left(\mathbf{C}_1 + \mathbf{C}_2 - 2 \left(\mathbf{C}_2^{1/2} \mathbf{C}_1 \mathbf{C}_2^{1/2} \right)^{1/2} \right).$$

Assuming unit diagonal variances, the WD collapses to the Euclidean distance between the means. To further help the interpretation, we divide this number by the square root of the dimensionality n (i.e., the number of simulated fault patches):

$$\text{SWD} = \frac{1}{\sqrt{n}} \text{WD}. \quad (5.2)$$

Like this, the Scaled Wasserstein 2-Distance (SWD) yields an ‘‘average’’ distance between the means of each parameter individually.

5.3 Inversion Results

For the elastic halfspace, we choose a shear modulus of 45 GPa and a Poisson’s ratio of 0.25. To minimize reference frame issues, we restrict ourselves to modeling the subduction zone offshore Northern Honshu, i.e., excluding the island of Hokkaido in the north and the region close to the junction with the Izu-Ogasawara and Nankai trough to the south, see Fig. 5.3 (compare, e.g., Loveless and Meade, 2010; Perfettini and Avouac, 2014; Loveless and Meade, 2016). Correspondingly, we only use data from the area onshore of our study region, leaving 242 stations. Since we are not modeling the Hokkaido part of the subduction zone, we remove the effect of the 2003 M_w 8.2 Tokachi earthquake on the remaining stations. We downsample the timeseries to a total of 24 timesteps, 12 before and after the 2011 M_w 9.1

Tohoku-oki earthquake, respectively, using time windows that logarithmically increase in length. We calculate the best-fit approximation to the data at the center of each window using a polynomial least-squares regression. Following the kinematic analysis presented in Appendix C.3, we set the observation uncertainty to the scalar value of 5 cm. To avoid numerical issues due to high coseismic stress changes at the sharp edges just outside the asperities, we smoothly taper the applied coseismic stress changes to zero over a distance of 60 km based on the considerations of Hetland et al. (2010). For our model, we choose a $1 - \text{erf}$ functional form, which avoids areas of negative curvature in the tapering function (which would be unexpected for coseismic slip fields outside of asperities; as opposed to, e.g., a Gaussian function), and still has a compact support (as opposed to, e.g., an inverse quadratic function). Plate convergence of the interface is defined by a relative Euler pole derived from the DISSTANS-estimated secular velocity estimated at the GNSS stations. Outside the region we are estimating, we enforce an effectively constant plate convergence by setting the rate-dependent frictional parameter to a very large value. We balance the desire for spatial resolution on the interface while preventing overfitting by parametrizing the frictional parameter using four parameters down-dip and five parameters along-strike. This discretization totals 20 parameters to be estimated by the AITar MCMC solver, over an along-strike distance of about 700 km and down to 300 km of depth. We assume a uniform prior over the parameters bounds of $5 \leq \log_{10} \alpha_h \leq 10$. 32 Nvidia P100 GPUs run the AITar solver for 16384 chains (chain length of 10), finishing after 6.7 hours and 22 iterations. We verify that additional samples do not change the recovered posterior covariance matrix in Section 5.3.2

5.3.1 Results of the Reference Model

The output of our model comprises slip (rate) on the fault interface and surface displacement (rate) over the entire observational timeframe, as well as the preferred rheological properties of the interface. To gain intuition for the output, Fig. 5.8 shows the mean modeled fault slip rate on the plate interface at different times within the observational period. For selected patches (marked by black outlines), the cumulative fault slip timeseries are shown in Fig. 5.9. From the modeled fault slip, a wide range of behaviors can be observed. For the pre-2011 (preseismic) timestamps, slip rates are mostly close to the longterm plate convergence rate, except near the locked asperities. Postseismically, large slip rate increases of one order of magnitude or more can be observed close to the slipping asperities (see Fig. 5.4). Close to the

asperities, tapered slip additionally leads to coseismic motion. (Due to the tapered slip applied for numerical reasons, some patches adjacent to the asperities experience a stress drop, leading them to barely slip at all temporarily. As it concerns the surface observations, the effect of these fully-stopped patches is averaged out by neighboring patches which experience large stress and therefore slip rate increases. After most postseismic stress is relaxed, these stuck patches start to slip again as expected.)

Fig. 5.10 shows the observed and predicted displacement timeseries at the surface for select stations inland of the 2011 M_w 9.1 Tohoku-oki rupture. The timeseries are fit well qualitatively at most stations, capturing both the preseismic velocity and the magnitude and curvature of the Tohoku-oki postseismic transient relaxation. As commonly seen when studying subduction zone earthquakes, it is harder to fit the vertical postseismic transient compared to the horizontal one (further discussed in Section 5.4.3). The average displacement error is $\delta_d = 4.0$ cm, which is close to the kinematically estimated optimal misfit of $\delta_d^{\text{kin}} = 2.4$ cm (see Appendix C.3). Fig. 5.10 also shows the predicted surface displacement timeseries calculated from the rheological parameters recovered in the reference case, but using the fine mesh in the forward model. In general, the coarse and fine meshes produce reasonably similar displacement fields, a result further substantiated by the analysis in Appendix C.2, justifying our computationally-motivated choice to perform inversions only on the coarse mesh. Since the rheological parameters have not been optimized for the fine mesh, however, we continue to show reference model results with the native coarse mesh.

The quality of fit to the interseismic, secular velocity field is mainly determined by the geometry of the fault interface, the assumed plate convergence rate, and most importantly the location, shape, and size of the fully-coupled asperities. Fig. 5.11 shows the pre-2011 velocity field residuals, showing an overall good agreement and indicating that the choice of our locked asperity is sufficiently close to the apparent kinematic coupling in the later interseismic period. South of 39°N , the horizontal residuals are very small and only show secondary effects of the full locking of the asperities (which were constrained by historic ruptures, not coupling maps). North of 39°N , there is a consistent misfit in the westward direction, potentially related to intraplate deformation. In fact, the northern westward motions do not contribute to excess strain compared to the south, as the residual horizontal strain rates show (Fig. 5.12). Closer to the asperities at the coasts, however, the horizontal velocities

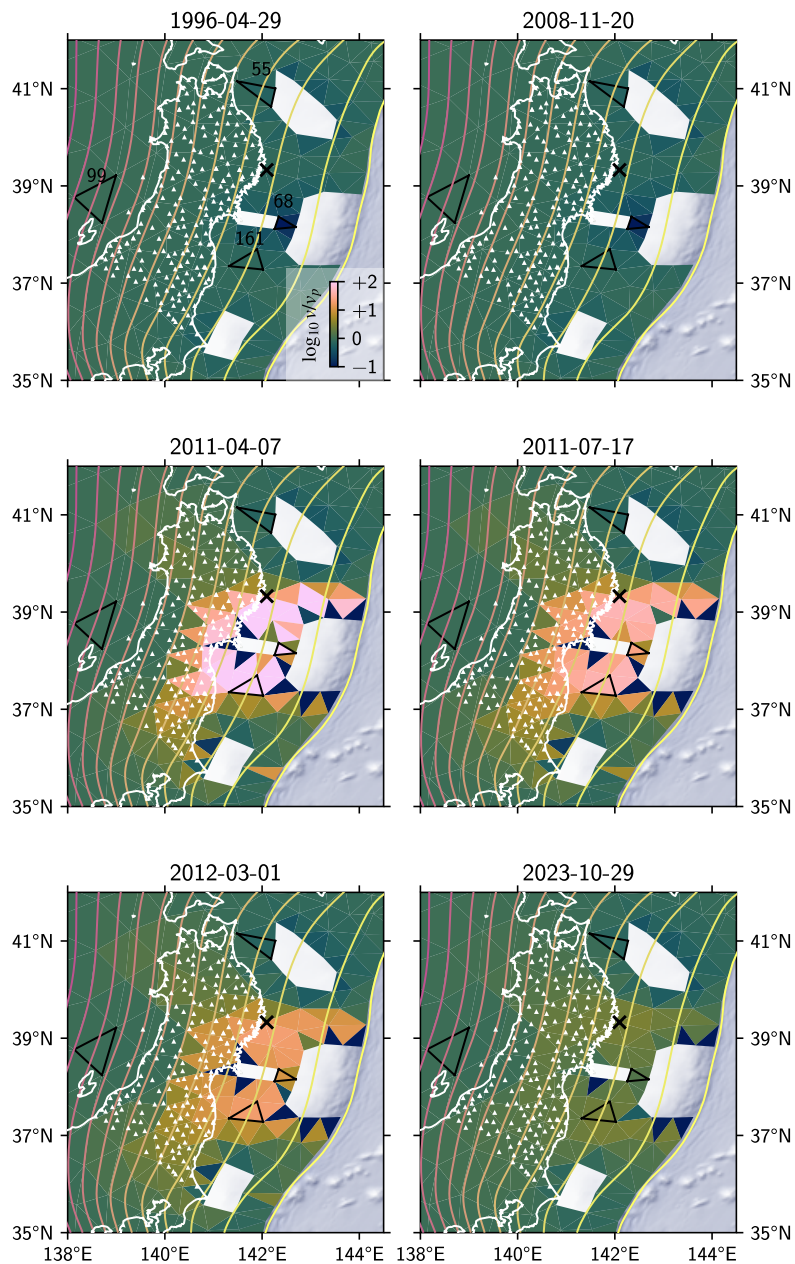


Figure 5.8: Mean fault slip rate on the Northern Japanese megathrust for the reference model inversion at six different times (two pre- and four postseismic). The discretized interface is colored according to the base-10-logarithm of the ratio between inferred slip rate and the assumed plate convergence rate. The cumulative fault slip timeseries at the patches presented with black outlines are shown in Fig. 5.9. The black cross indicates the location of the Kamaishi repeating earthquake sequence, its relation to the results is discussed in Section 5.4.6. Additional map elements as in Fig. 5.4.




Figure 5.9: Cumulative fault slip since 1996 at select patches on the interface. The vertical axis is displacement [m] and the horizontal axes is time. Lines and shaded areas show the sample mean and ± 3 standard deviations of the forward model output based on the estimated rheological parameters. The range of plate rates on the interface is shown by the shaded gray region for comparison.

imply unmodeled strain rates. Nevertheless, the absolute magnitude of the misfit is less than 10% of the plate rate, which indicates that this misfit should not negatively affect the interpretability of our model results. In the vertical direction, velocities are generally underestimated, but regional and local differences inhibit the emergence of a clear pattern. Indeed, Fig. 5.13 shows that this variability is most likely due to unmodeled local processes in data, suggesting that local effects are dominant during the interseismic phase. Qualitatively, our interseismic residuals are comparable to those reported by Kanda et al. (2013).

Figs. 5.14–5.17 show the post-Tohoku-oki cumulative displacement field. All components agree with the data qualitatively, and the model reproduces the first-order directions and magnitudes of the transient relaxation. The vertical hinge line, shown in Fig. 5.16, is especially well recovered. Fig. 5.17 again shows, however, the out-size influence of the definition of the asperities, as systemic misfits are mostly close to the fully-coupled patches. Furthermore, there is a small but systematic eastward misfit on the west coast of Honshu.

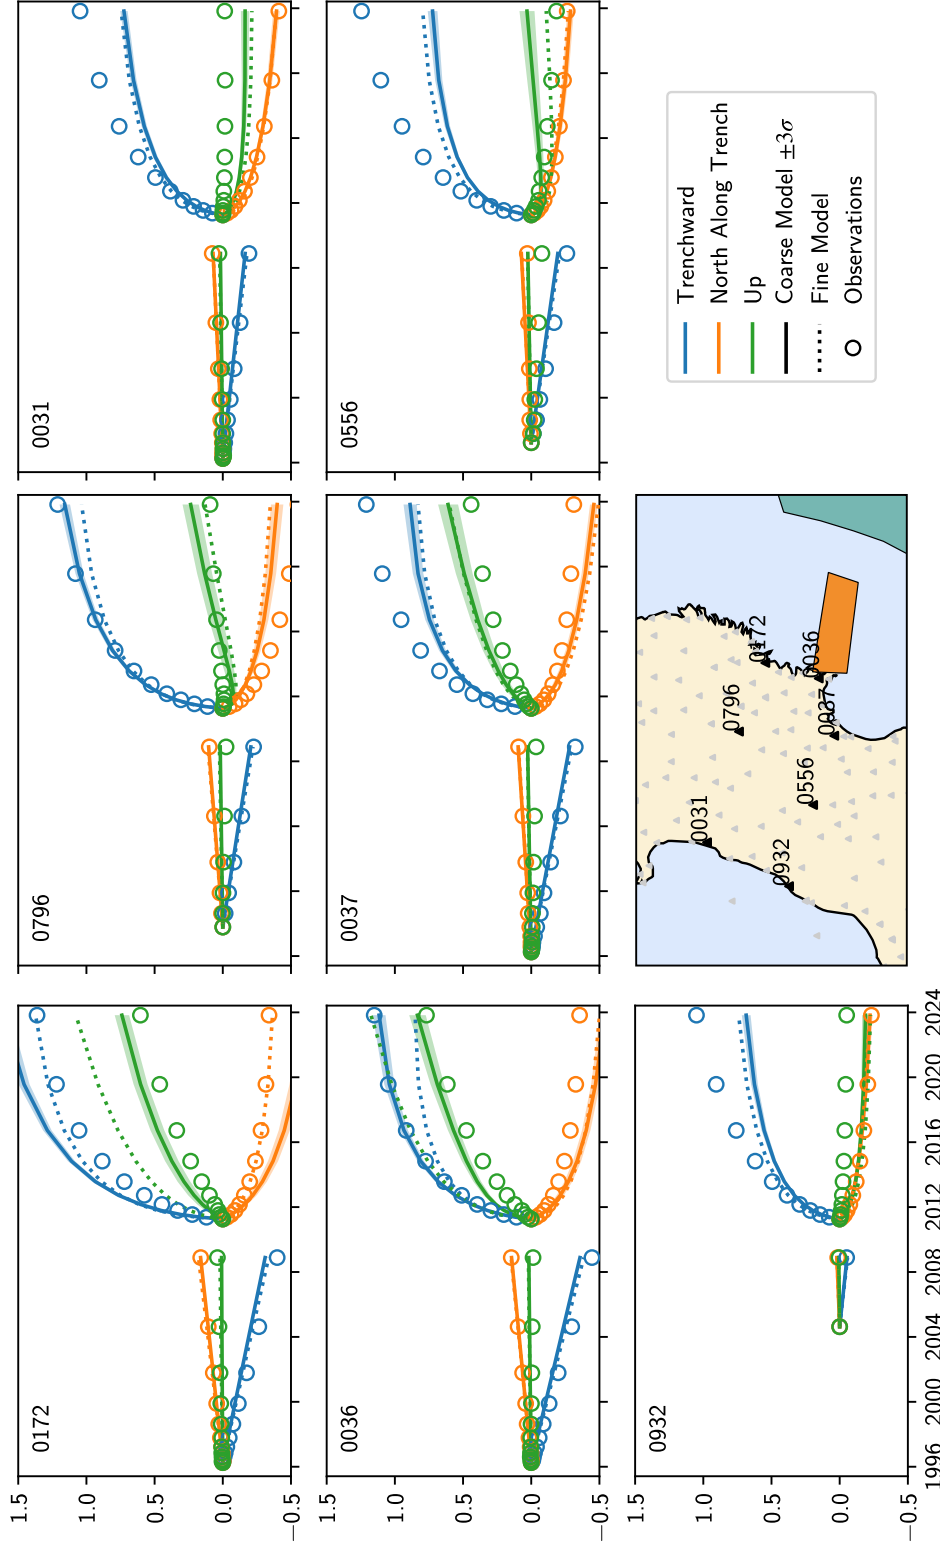


Figure 5.10: Observed and modeled timeseries for the reference model inversion at select stations. The vertical axis is displacement [m] and the horizontal axis is time. Trenchward, along-strike, and vertical data components are shown in blue, orange, and green, respectively. Circles show the observations to be fitted, reset to zero after the 2011 Tohoku-oki earthquake. Lines and shaded areas show the sample mean and ± 3 standard deviations of the forward model output based on the estimated rheological parameters in the reference case using the coarse mesh. The dotted lines show the forward model output using the fine mesh using the same estimated rheological parameters. The center bottom panel shows a zoomed-in map view of the station network, where the stations shown in this figure are labeled and marked with black polygons. The Miyagi and Tohoku asperities are shown as orange and green polygons.




Figure 5.11: Interseismic velocity residuals for Northern Honshu Island, Japan. The residuals are defined as the difference between the observed velocities and the sample mean velocities of the forward model output. The arrows show the horizontal residual, colored by the vertical residual. The asperities included in the inversion framework are shown by polygons (same colors as in Fig. 5.1). Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). Yellow to pink lines show the plate interface at 20 km depth intervals (Hayes et al., 2018), including the plate boundary at the trench (Bird, 2003).




Figure 5.12: Interseismic maximum residual shear strain rate (square root of the second invariant of the deviatoric strain rate tensor), calculated from the horizontal velocity residuals shown in Fig. 5.11 using the VISR algorithm (Shen et al., 2015) as implemented in DISSTANS (Köhne et al., 2023). Other map content as in Fig. 5.11.




Figure 5.13: Interseismic observed (filled circles) and modeled (background field) vertical velocities for Northern Honshu. Other map content as in Fig. 5.11.

Finally, Fig. 5.18 shows the mean and standard deviation of the posterior distribution of the rate-dependent frictional parameter $\alpha_h = (a - b)\sigma_E$, mapped onto each fault patch. In terms of marginalized standard deviation of the posterior, the recovered parameters below and trenchward of the used GNSS stations are very well constrained (Fig. 5.18b). We obtain values of approximately 100 kPa around the Tohoku asperity, which slightly decrease underneath Honshu (Fig. 5.18a). At larger depths around 39°N, we obtain an increase in frictional parameter, which is also correlated with higher uncertainty and will be discussed in Section 5.4.3. Towards the north and south, slightly outside the data coverage, we find values even lower at 10 kPa.

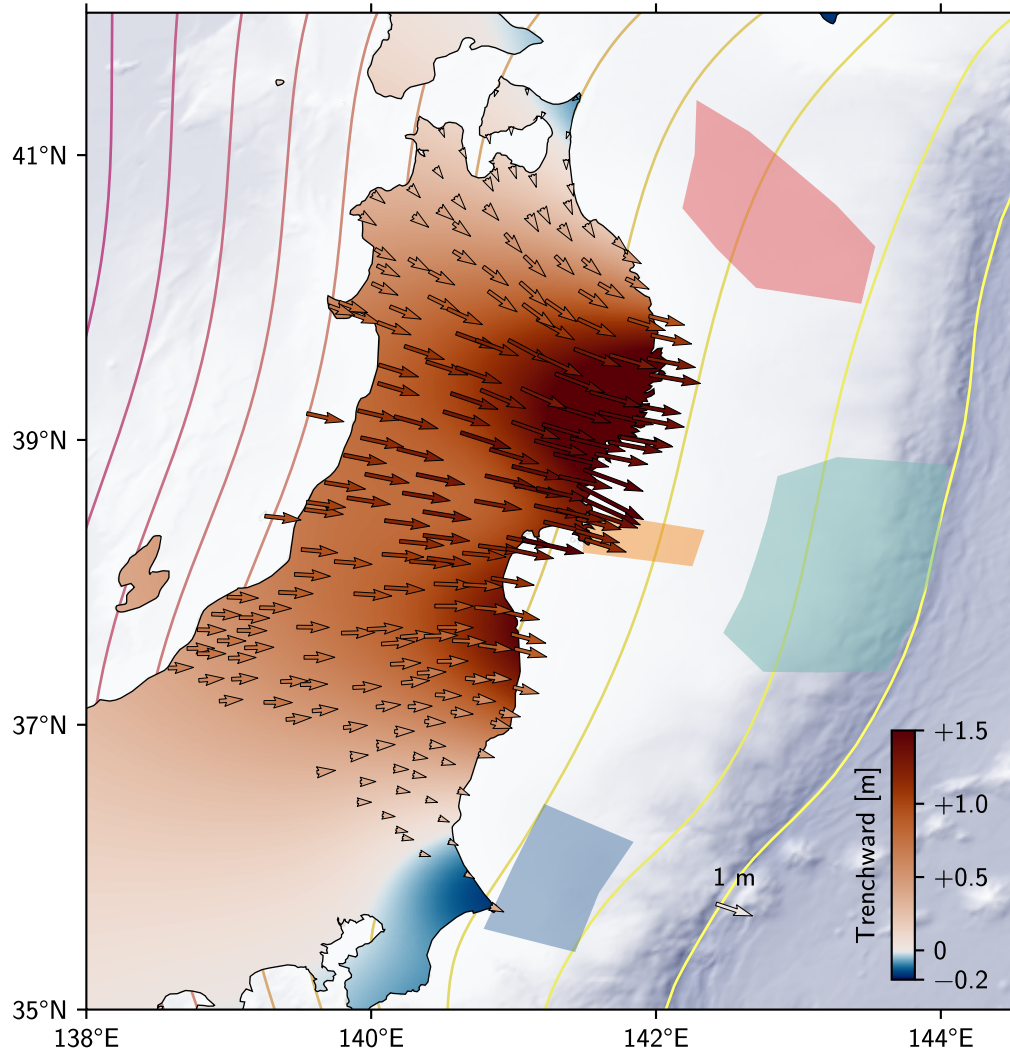


Figure 5.14: Postseismic observed and modeled trenchward cumulative displacements for the timeframe 2011–2024 in Northern Honshu. The background field shows the modeled trenchward cumulative displacement. The arrows show the observed cumulative displacement at each of the considered GNSS stations, colored by its trenchward component, using the same color scheme as for the model output. The closer the arrow colors are to the background field color where the arrow originates, the better the fit. Other map content as in Fig. 5.11.




Figure 5.15: Same as Fig. 5.14 but for the along-strike (North along trench) component.

5.3.2 Model Sensitivities

First, we assess in case (2) whether the number of chains in our MCMC inversions is sufficient. Then, to assess the robustness of the recovery of the spatially-variable frictional parameter on the plate interface, we perform additional inversions with different observations or different α_h spatial parametrizations (i.e., the number of splines spanning the simulated part of the interface). Cases (3)–(6) vary the observational periods used to constrain the inverse model, including cases that only consider pre- or post-Tohoku-oki data. Cases (7) and (8) vary the amount of spatial variability of α_h possible on the interface, and case (9) assesses the importance of the vertical data component. Specifically, we varied the following for each case:

Case (1): The reference case presented in Section 5.3.1.




Figure 5.16: Same as Fig. 5.14 but for the vertical component (with filled circles instead of arrows).

Case (2): With four times as many chains in the MCMC inversion.

Case (3): Only with observations from the time frame 1996–2012.

Case (4): Without observations from the year 2011.

Case (5): Only with observations prior to the 2011 Tohoku-oki earthquake.

Case (6): Only with observations after the 2011 Tohoku-oki earthquake.

Case (7): Without along-strike variability (i.e., 4 splines down-dip and 1 along-strike).






Figure 5.17: Postseismic cumulative displacement residuals for the timeframe 2011–2024 in Northern Honshu. The residuals are defined as the difference between the observed cumulative displacement in 2024 and the sample mean displacement of the forward model output. The arrows show the horizontal residual, colored by the vertical residual. Other map content as in Fig. 5.11.



(a) Mean



(b) Standard Deviation

Figure 5.18: See next page.

Figure 5.18: Mean and standard deviation of the base-10-logarithm of the frictional parameter $\alpha_h = (a - b)\sigma_E$ [Pa] of the discretized plate interface. Assumed parameter bounds are $5 \leq \log_{10} \alpha_h \leq 10$. North of the Sanriku, south of the Ibaraki asperities, and below 200 km, the frictional parameter is set to $\alpha_h = 10^{10}$ Pa to mimic steady plate convergence far away from the model domain, and so we do not interpret the apparent variations or inferences in these border regions (red dashed line). Background imagery shows shaded relief on land and on the ocean floor (ETOPO 2022 dataset, NOAA National Centers for Environmental Information, 2022). The asperities are marked by different colors (introduced in Fig. 5.1). Yellow to pink lines show the plate interface at 20 km depth intervals (Hayes et al., 2018), including the plate boundary at the trench (Bird, 2003). White triangles denote the locations of GNSS stations processed for the inversion.

Case (8): With only 2 splines in each direction (i.e., 2 splines down-dip and along-strike each).

Case (9): Only with horizontal observations.

Table 5.2 shows the average displacement errors δ_d for these test cases with respect to the entire timeseries, as well as their difference in terms of δ_d and the α_h Scaled Wasserstein 2-Distances (SWD) relative to the reference case. Using more chains in the inversions as in case (2) does not improve the quality of fit to the data, as measured by the virtually-identical values of δ_d . The non-zero SWD of 0.21 reflects the fact that minor differences in the recovered α_h values are still present, but we find that the small differences to the reference case are all at the borders of our simulation, where uncertainties are also higher. We therefore interpret an SWD of 0.21 to be reflective of a very similar result of α_h for the following cases.

From case (6), which contains no preseismic but all the postseismic data, it is obvious that almost all the information content to constrain the frictional parameter is contained in the postseismic period: the fit to the surface data is effectively the same, and the recovered α_h values are very similar. As the amount of postseismic data becomes less in cases (3) and (4), the fit is gradually worsened, until it can not fit the postseismic data anymore at all in case (5). Along-strike variability seems to be slightly favored by the data, as the average displacement error of case (7), which does not contain such variability, is about 15% worse. For comparison, Fig. 5.19 shows the recovered fault frictional strength for this test case, which is fairly uniform across the interface, explaining the marginally worse fit to the data. On the other hand, down-dip variability is much more important, as case (8) shows: not only is the fit to the surface data about 28% worse, the recovered values of α_h are




Figure 5.19: Same as Fig. 5.18a but for test case (7). The logarithmic uncertainty everywhere within the estimated region is 0.1 or less (not shown; compare Fig. 5.18b).

also much further away than the other test cases (except the preseismic-only case). The omission of observations in the vertical direction in case (9) changes the recovered rheological parameters to a similar degree as cases (3)–(4) and (7), the effect of which will be discussed in Section 5.4.3.

Overall, the following picture emerges from the sensitivity analysis: Preseismic data does not provide significant information for the estimation of frictional parameter. (As discussed in Section 5.3.1, preseismic data is crucial in the definition of the asperities, however.) In terms of postseismic data, both early and late periods contain complementary information that can be fit jointly. Along-strike variability is favored by the data, although it is not as important as the down-dip variability. Given these interpretations, the general correlation of δ_d and SWD, and manual, qualitative comparisons of the different model results, the rate-dependent frictional parameters shown in Fig. 5.18 can be considered robust results within their estimated uncertainties.

Case	δ_d	δ_d Increase	SWD
(1)	4.0 cm	—	—
(2)	4.0 cm	-0.7%	0.25
(3)	5.3 cm	31.7%	0.89
(4)	4.7 cm	16.8%	0.60
(5)	12.1 cm	199.8%	1.59
(6)	4.0 cm	0.2%	0.21
(7)	4.6 cm	15.2%	0.69
(8)	5.2 cm	28.3%	1.36
(9)	4.9 cm	22.7%	0.80

Table 5.2: Inversion performance of the additional test runs (2–9) relative to the reference case (1) presented in Section 5.3.1. Quality metrics defined in eqs. (5.1) and (5.2).

5.4 Discussion

The probabilistic inversion framework built on simulating sequences of earthquakes and aseismic slip (SEAS) we present in this study is able to reproduce to a reasonable degree the GNSS observations in Northern Honshu from 1996–2024, including both steady interseismic and significant postseismic deformation. Even though our model only includes interseismic creep and postseismic afterslip as physical effects, it is able to show spatially complex in deformation patterns and postseismic relaxation of multiple timescales.

5.4.1 Stress-diffusive Nature of Observations and SEAS Models

As postseismic slip on the fault interface in our model is mainly driven by coseismic stress (and therefore fault slip rate) change due earthquakes on the asperities, it follows a diffusive nature where high slip rates close to the asperity will decay in time and outwards from the asperities in space (see Fig. 5.8). As Ortega-Culaciati (2013) shows for the case of Tohoku-oki, the slip on the fault decays logarithmically in time. Equivalently, the cumulative fault slip in time windows which increase logarithmically in length, is approximately constant. Fig. 5.20 shows the excess cumulative slip in map view for six different postseismic time windows (i.e., removing the fault slip occurring due to the far-field plate loading). Indeed, the different time frames appear approximately similar in shape and magnitude, except for the first time window where most slip is still localized to the regions close to coseismic slip. Here, the patch slip rates react to details of the assumed asperity map and the tapering of coseismic slip. Fig. 5.21 show the moment release over time of the entire simulated fault interface based on the excess fault slip, exhibiting an approximately constant moment release of M_w 8.1. This is close to the M_w 8.2 value inferred kinematically

by Ortega-Culaciati (2013), even though he only considered the first 1.5 years.

The stress-diffusive nature of the postseismic relaxation also implies that at any given surface location, the amount of displacement in a time window early after the 2011 rupture should be positively correlated with displacement many years later over a time window that logarithmically increased in length. Fig. 5.22 shows the correlation of the displacement over an early and late time window for the data and reference model predictions. We find moderate correlations between the two time frames, with R^2 values of 0.66 and 0.43 for the model and the data, respectively. Fig. 5.23 shows the difference between the late displacement and the early displacement times the best-fit slope, i.e., the correlation residual. For both the data and the model, strong spatial patterns are visible, even though they are distinct. For example, the data correlation underpredicts the displacement on the coast to the north and south of the Tohoku asperity, while the model correlation underpredicts the displacement at the coast and inland west of the Tohoku asperity.

Spatially coherent correlation residuals can be explained by a spatially non-uniform slip rate diffusion, e.g., when the rheological properties vary in space. In our reference case, we estimated the fault frictional parameter and found significant spatial heterogeneity, which we therefore interpret to be the main contributing factor to the early and late time windows not being better correlated. Overall, however, we find that with $R^2 \approx 0.5$, about half of the surface displacement between early and late time windows for our observations and model predictions can be attributed to the stress-diffusive nature of postseismic relaxation. The significant correlation therefore further substantiates our initial assumption that an asperity-based earthquake cycle model can explain much of the postseismic phase.

5.4.2 Interseismic Coupling

Our model output includes slip rate at all discretized fault patches for the entire observed period. Hence, we can calculate the apparent, kinematic coupling coefficient on the interface as a function of time t , which we define as $1 - v(t)/v_p$, where v is the fault patch slip rate and v_p is the plate convergence rate at the patch. The kinematic coupling coefficient is defined independently of the stress state dictating the slip rate, and therefore its physical interpretation as hinted by Wang and Dixon (2004) is ambiguous. On one end of the spectrum, a patch may not be sliding because the frictional resistance is strong enough to prevent any slip on the fault interface. On the other end, low frictional resistance may not lead to fault slip if the local stress

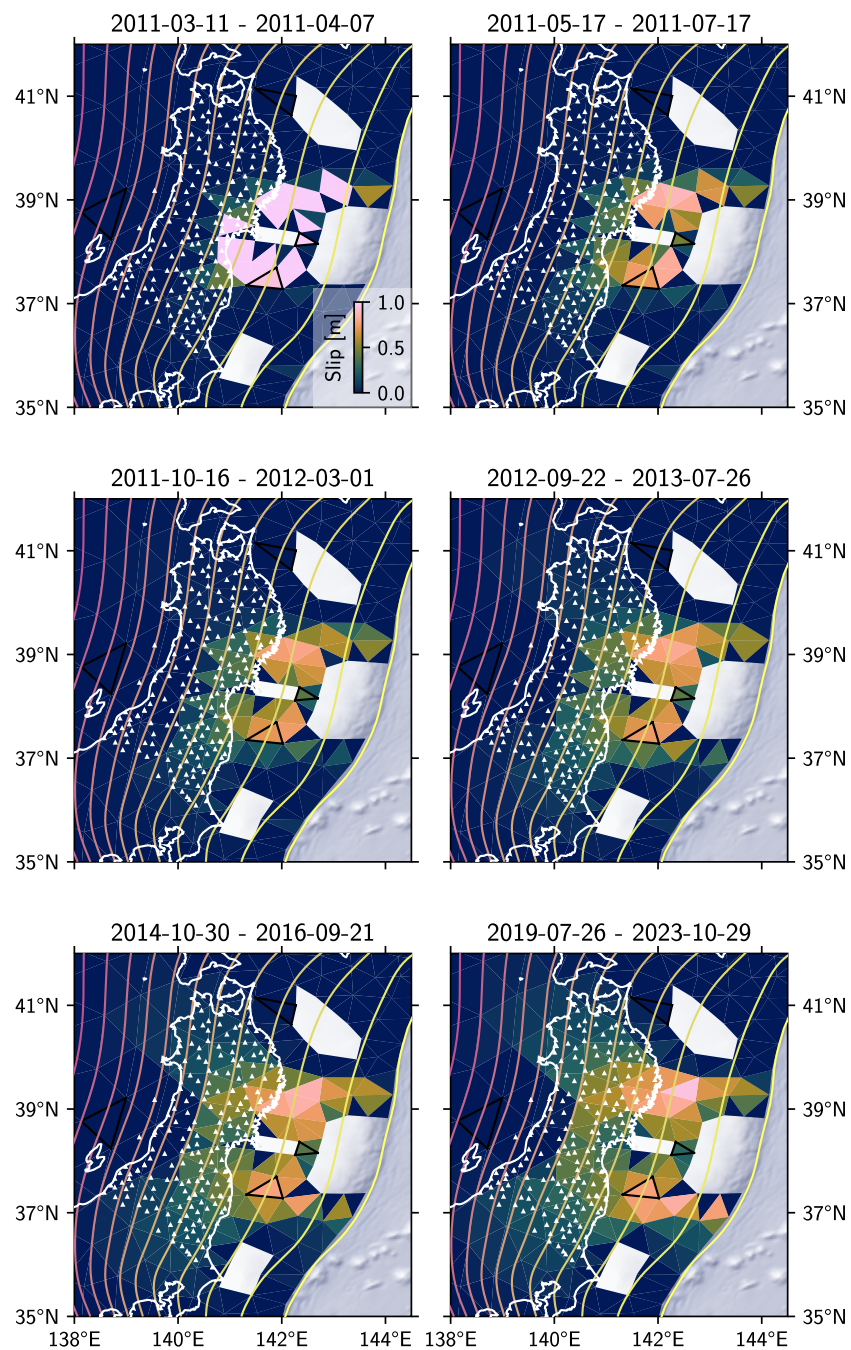


Figure 5.20: Excess fault slip on the Northern Japanese megathrust for the reference model inversion at six postseismic times. The discretized interface is colored according to the total fault slip between the logarithmically-spaced timestamps, excluding plate convergence contributions. The patches outlined in black are the same as in Fig. 5.8. Additional map elements as in Fig. 5.4.




Figure 5.21: Mean excess moment magnitude over the entire creeping fault interface for the interval between the observation timestamps. Does not include contributions of the plate convergence.




Figure 5.22: Correlation between horizontal surface displacement at the observing stations for an early (2011-05-17–2011-07-17) and late (2019-07-26–2023-10-29) time window. Blue and orange dots and dashed lines show the surface displacements at the observing stations and their best-fit linear slope for the model predictions and observations, respectively. The gray line corresponds to a 1:1 correlation, but note that linear slopes larger or smaller than one only mean the logarithmic time window growth is smaller or larger than required by the postseismic transient, respectively.




Figure 5.23: Residual between the displacement over a late time frame (2019-07-26–2023-10-29) and the displacement calculated from the best-fit correlation slope (see Fig. 5.22) using the displacement in an early time frame (2011-05-17–2011-07-17). Correlation residuals are shown for the observations (filled circles) and reference model results (background field). Other map content as in Fig. 5.11.

state does not require it to slide (e.g., when a patch is inside the stress shadow of an asperity). Kinematically, and for the purposes of this study, both of these patches are fully locked, but the interpretation as to *why* they are not sliding is still a significant physical distinction for seismic hazard estimates. For this reason, we will discuss the *slip deficit* as a more informative quantity further below.

Nonzero coupling is a direct consequence of the stress-diffusive nature of our model; i.e., nonzero coupling emerges only in areas where large earthquakes induced afterslip. Fig. 5.24a shows the coupling coefficient on the entire simulated interface, calculated using the last instantaneous slip rate before the 2011 Tohoku-oki earthquake. The prescribed locked asperities are clearly visible, and they create stress shadows in their vicinity that in turn lead to high coupling ratios (compare, e.g., Hetland and Simons, 2010; Hetland et al., 2010).

In our results, the interface remains partially coupled even to depths larger than 100 km, which qualitatively reproduces findings of deep coupling in Northern Japan (Suwa et al., 2006). Furthermore, our results match to first order the coupling coefficients from Loveless and Meade (2016) and Lindsey et al. (2021), who obtain two distinct regions of high coupling south of 40°N latitude. Differences in the location and size of areas of high coupling to our model are due to the setup of our framework: We do not invert for asperities using interseismic velocity, but rather constrain them using rupture areas of large earthquakes. These asperities are chosen such that the general afterslip patterns can be reproduced, while still providing an acceptable fit to preseismic observations.

Unlike the coupling map of Lindsey et al. (2021), which shows a very sharp transition down-dip the of the Miyagi and Tohoku asperities, however, our model predicts a smooth transition from high to low coupling. This smooth change is not entirely attributable to the stress-diffusive character of our model, since a higher contrast in the recovered variability of the frictional parameter $\alpha_h = (a - b)\sigma_E$ would directly map into a stronger contrast of interseismic coupling. A higher frictional parameter, however, would also limit the amount of postseismic relaxation possible, and so we interpret the difference between our results and those of Lindsey et al. (2021) to be again due to the inclusion of postseismic observations and our requirement that the implied frictional properties are the same at all times in the seismic cycle.

North of 40°N latitude, the differences between our model results and those discussed from the literature can likely be attributed to differences in the fitted data and model setup. For instance, neither Lindsey et al. (2021) nor our models take


into account relative block motion of Hokkaido relative to Honshu as Loveless and Meade (2016) do. In contrast to Lindsey et al. (2021), we have chosen to omit this area (both in terms of data and modeling) entirely to avoid reference frame issues.

For the purposes of seismic hazard assessment, however, kinematic coupling maps are only useful if they are a good proxy for the slip deficit (or similarly, the region expected to slip coseismically). Slip deficit, defined as $1 - \int v(t)/v_p dt$, can significantly differ from the kinematic coupling coefficient if the observed fault slip rate is not constant throughout the interseismic period. This can be the case in areas of significant afterslip (slipping faster early in the interseismic phase, and slower later) or regions of slow aseismic slip (e.g., Ozawa et al., 2013; Riel et al., 2014; Saux et al., 2022). While slow slip events are not included in our inversion (and most likely do not play a role in Northern Honshu), we do include afterslip and, as such, can recover the cumulative fault slip at each patch from our model output.


Fig. 5.24b shows the cumulative fault slip converted to slip deficit. As in the kinematic coupling map, the asperities are clearly visible as first-order features. However, two major differences to Fig. 5.24a are visible. First, the Sanriku asperity is not colored in dark black, as it ruptured in 1994 and has not yet accumulated enough slip to host another earthquake. From the kinematic coupling perspective, there would be no difference between any of the asperities. Second, and perhaps even more prominent, is the fact that almost all coupling deeper than 40 km is not present as slip deficit. Only the region around the large Tohoku asperity shows a significant slip deficit, and could therefore participate in an upcoming earthquake. The contrast of significant kinematic coupling, but negligible slip deficit at depths larger than 40 km, naturally reconciles the observations of deep coupling in Northern Japan *without* the need for the presence of earthquakes (Suwa et al., 2006). (We do not interpret the slip deficit north of the Sanriku asperity, as this is the transition region from estimated to imposed frictional parameter.)

5.4.3 Tradeoffs Between Horizontal and Vertical Observation Components

The displacement and stress kernels we derive for our forward model, by nature of the homogenous linear elastic halfspace, do not take into account the known spatial variability of the elastic modulus (including depth dependence), the presence of the rigid slab, or the effect of topography. These effects have been shown to play important roles when modeling the coseismic displacements generated by large subduction zones earthquakes (e.g., Hsu et al., 2011; Hashima et al., 2016; Langer and



(a) Kinematic Coupling



(b) Slip Deficit

Figure 5.24: Comparison of the preseismic coupling ($1 - v/v_p$) and accumulated slip deficit ($1 - \int v(t)/v_p dt$) as output from the reference inversion. Trench, slab depths, and GNSS station locations as in Fig. 5.18.


Ragon, 2023; Ragon and Simons, 2023). In particular, displacement kernels that include these heterogeneities can significantly reduce systematic misfits between the horizontal and vertical components of the displacement field, which usually trade off with each other in kinematic models assuming a homogenous half space.

We investigate the influence of the horizontal-vertical tradeoff on our framework by analyzing in more detail case (9) from Section 5.3.2, where the fault frictional parameter α_h is estimated using only horizontal observations. Fig. 5.25a shows the mean and standard deviation of the estimated frictional parameter. Qualitatively, the lack of the deep region of very high α_h values below the west coast around 39°N compared to the frictional parameter recovered from the reference case (Fig. 5.18a) is apparent. Smaller changes are higher and lower values of frictional parameter off the west and east coast of the northern tip of Honshu, respectively, although this area is also characterized by a higher uncertainty (see Fig. 5.25b). Fig. 5.26 shows the difference of the fault frictional parameter between cases (1) and (9), also highlighting the disappearance of the high α_h values at 39°N. Otherwise, the differences are mainly correlated with regions of higher posterior uncertainty. The changes in the estimated frictional parameter across the plate interface when considering vertical data or not in our inversions therefore suggest that heterogeneous fault structure in terms of elastic modulus gradients may play a significant role in the recovery of rheological properties.


5.4.4 Frictional Parameter Along the Northern Japan Plate Interface

The ALTar inversion step in our framework estimates spatially-variable fault frictional parameter $\alpha_h = (a - b)\sigma_E$ [Pa] that best fits the observed surface displacement timeseries through a piecewise linear spline representation imposed on the fault interface. Fig. 5.18 shows the resulting frictional parameter for each patch, and Fig. 5.28 shows a corner plot of the underlying spline coefficients. In general, in areas of low uncertainty, we find a general decrease of the frictional parameter with increasing depth. This finding is consistent with the frictional property inversions using slip rate timeseries by Thomas et al. (2017) at the Longitudinal Valley Fault in Taiwan for depths less than 50 km.

At all depths, we find that the frictional parameter varies along strike. One cause for this variability, especially towards the lower end of the recovered values, may be a true feature of the data: Rate-neutral regions (i.e., where $a - b \rightarrow 0$) lead to estimates of $\alpha_h \rightarrow 0$ regardless of the effective normal pressure σ_E , and could therefore be a



(a) Mean



(b) Standard Deviation

Figure 5.25: See next page.

Figure 5.25: Mean and standard deviation of the base-10-logarithm of the frictional parameter $\alpha_h = (a - b)\sigma_E$ [Pa] of the discretized plate interface for inversion case (9), ignoring all vertical data in the fitting process. Map description and additional elements as in Fig. 5.18.




Figure 5.26: Difference of the base-10-logarithm of the frictional parameter α_h between reference case (1) and the inversion test case (9), which ignores all vertical data in the fitting process. Additional map elements as in Fig. 5.18

valid interpretation. However, one of the assumptions of our model is that the fault slip behavior everywhere could be approximated by rate-strengthening friction, and the presence of rate-neutral patches may therefore violate our model assumptions.

On the higher side of the spectrum of values we recover, even at low depths (less than 40 km) where friction may be the dominant slip mechanism, our estimates reach up to $\alpha_h = 10^7$ Pa. Assuming $\sigma_E = 100$ MPa, this would imply ranges of $a - b$ between 10^{-3} – 10^{-1} , which are close to the laboratory-derived values of 10^{-3} – 10^{-2} (Marone et al., 1991; Blanpied et al., 1995). Fig. 5.27 shows the values of $a - b$ on the interface for the reference case assuming hydrostatic pressure on the fault, rather than a constant value. Here, we take $\sigma_E = \kappa (\rho_{\text{rock}} - \rho_{\text{water}}) g z$, where κ is the pres-




Figure 5.27: Rate-dependent $a - b$ parameter derived from α_h assuming hydrostatic pressure. Additional map elements as in Fig. 5.18

sure dependence, ρ are material densities, g is the gravitational acceleration, and z is the depths of the patches (see, e.g., Suppe, 2014; Thomas et al., 2017). Using depth-dependent effective normal stress leads to a larger variation of implied $a - b$ values, however, the recovered values do overlap with laboratory-derived values for significant parts of the interface. Together, even with the range of α_h values we estimate on the entire interface and the strong dependence on the assumed σ_E profile, we find that our regional-scale, earthquake-cycle-derived $a - b$ values are overall compatible to those from the rate-and-state laboratory and numerical literature (Dieterich, 1979; Rice and Ruina, 1983; Ruina, 1983; Hetland et al., 2010; Fukuda and Johnson, 2021).

The differences between our and laboratory-derived frictional properties could be explained by recognizing that, postseismically, our inversion scheme is mostly sensitive to the cumulative effect of slip rates of the interface. In turn, the evolution of the slip rate on the interface is a direct consequence of the imposed coseismic stress change and the frictional parameter α_h . Since the coseismic stress change is

not an inverted parameter, but coseismic velocity change is proportional to the stress change and α_h , the MCMC solver faces the potentially unresolvable issue of fitting both the early and late postseismic relaxation using a single parameter. Therefore, without any prior, external constraints on the bounds of α_h , we caution against directly interpreting the recovered values.

This assessment is supported by exploratory test runs we performed with different boundary conditions. We attempt to further decrease the lower bound of α_h to avoid the ill-constrained posterior distributions we obtain where the recovered α_h collapses onto the lower bound of $\alpha_{h,\min} = 10^5$ Pa (as in our reference model). We find that to the degree numerically possible, the MCMC solver always collapses the posterior distribution of some spline coefficients to the lower bound (see Fig. 5.28). We are not able to test values smaller than $\alpha_{h,\min} \approx 7 \cdot 10^4$ Pa, because the initial postseismic velocity driven by an average stress change of $\Delta\tau \approx 10^7$ Pa would tend to infinity, see (3.11), and numerical artifacts start to appear close to this value. On the other end of the spectrum, when setting a lower bound of $\alpha_h > 10^6$ Pa, the data is fit significantly worse ($\approx 50\%$ higher average surface errors at $\delta_d = 5.8$ cm). Taken together, our additional tests suggest that at least some α_h values our model infers could be biased by the prior parameter bounds we impose. Additionally, we have assumed that the effective normal stress is constant in time, which may not be an appropriate simplification if there are significant physical changes on the interface (e.g., fluid migration due to the large earthquakes). The recovered values of $\alpha_h = (a - b)\sigma_E$ may therefore be affected by both the numerical and physical limits of our framework, inhibiting the recovery of the “true” frictional properties of the Northern Japanese megathrust.

5.4.5 Importance of Bulk Flow

Our framework is simpler in terms of model complexity than other recent studies estimating rheological parameters from postseismic transients (e.g., Sun et al., 2014; Hu et al., 2016; Freed et al., 2017; Agata et al., 2019; Muto et al., 2019; Fukuda and Johnson, 2021). The main difference is that the above-cited studies all include some way of describing bulk flow, which is usually described as being necessary to fit timeseries with multiple apparent relaxation timescales, as well as the early landward motion of seafloor geodetic stations. Given the well-established benefits of including bulk flow in studies of postseismic displacement, the question is warranted why we choose to ignore methods of including viscous flow behavior in our framework (or rather, why we assume such bulk flow can be approximated by a fric-




Figure 5.28: Condensed corner plot showing the posterior sample distribution of the 20 estimated α_{η} spline coefficients spanning the interface and yielding the frictional parameter values at each patch as shown in Fig. 5.18. The figures on the diagonal represent smoothed 1D histograms of the marginalized posterior probability density functions (PDF) for each parameter. A muted orange background indicates spline coefficients at the boundary of the model domain (i.e., the northernmost, southernmost, and deepest coefficients). The numerical values of the median as well as the ± 34 th percentile (equivalent to one standard deviation for a normal distribution) are shown above the diagonal. The lower, off-diagonal plots are smoothed 2D histograms of the posterior PDF, with contour lines indicating the 1, 2, and 3 standard deviation ranges from the mean. Most non-boundary coefficients are well recovered and are similar to Normal distributions. Boundary coefficients show one-sided PDFs skewed towards the imposed limits of the parameter range.

tionally sliding interface). The answer is threefold. We want to stress, however, that none of our reasons should be construed to argue against the presence of viscoelastic effects in the bulk, as is evidenced by the aforementioned studies — only as practical reasons validating our approach and results.

First, we show in Section 4.3.5 that for a synthetic, two-dimensional subduction zone, viscoelastic bulk motion does not significantly hinder the estimation of rate-strengthening frictional properties on the fault. The recovered α_h values are within half an order of magnitude of the target values, despite displacement timeseries misfits an order of magnitude larger than in cases where no viscoelastic bulk is present.

Second, from a modeling perspective, it is still unclear to which extent including bulk viscoelastic effects (the prevalent choice for modeling flow) are necessary to fit on-land GNSS observations following the 2011 Tohoku-oki earthquake. For example, Perfettini and Avouac (2014) are able to model early postseismic deformation without any viscous component, whereas the model of Freed et al. (2017) prefers a roughly even split in the importance between afterslip and viscoelastic relaxation for on-land stations. Fukuda and Johnson (2021) and Muto et al. (2019) go even further, crediting viscoelastic effects with more than twice as much impact on the observations as afterslip. While a quantitative comparison of each of the discussed studies in terms of RMS fit to the displacement timeseries (or similar) is yet outstanding (as the misfit metrics are not standardized, if reported at all), it is safe to state that at least some of them include some amount of systematic misfit to the data; e.g., an apparent north-south contraction away from the coast in Perfettini and Avouac (2014, Fig. 10), a north-south extension at the eastern coast in Fukuda and Johnson, 2021, Fig. 4, or an underprediction of trenchward motion on land in Muto et al. (2019, Fig. 2). Therefore, we deem it not unlikely that while including viscoelastic effects to model the motion of seafloor stations appears to be necessary, the inclusion may also obstruct the study of other physical effects or processes, some of which may be more important for land-based stations. Indeed, the ability of our reference case to model the surface observations (including the pre-Tohoku-oki period) to a comparable degree as more complex models suggests current models may not be sensitive to bulk viscous behavior beyond that which can be approximated by a dislocation interface (such as in our model).

The second reason we choose to omit viscoelastic relaxation in our forward model is simply because of its associated computational cost. To invert for model parameters using MCMC methods, tens of thousands to millions of forward models

(chains) must be evaluated in a reasonable timeframe. For example, Fukuda and Johnson (2021) report that simulating “7.3 years of afterslip [...] takes ~ 7 min” on a CPU. If we assume that this runtime would translate from their 2.5D model to our 3D one (which would more likely represent a runtime increase of 1–2 orders of magnitude), and this runtime would translate to one period between two earthquakes in our simulation, simulating *over one hundred* of these periods for *each chain* (as our framework is designed to do, see Section 5.2.1) would be prohibitive in the near future. Fully-3D codes based on Finite Element Methods (e.g., Hu et al., 2016; Freed et al., 2017; Agata et al., 2019), or fully-dynamic models simulating earthquake cycles including the coseismic phase (e.g., Dal Zilio et al., 2022), would be even more prohibitive. Of course, improvements in mathematical or numerical formulations may change this perception significantly, e.g., the potential addition of viscoelastic effects to Boundary Element Methods.

5.4.6 Relation with the Kamaishi Repeating Earthquake Sequence

Repeating earthquakes offshore Kamaishi in Northeastern Japan were remarkably regular before the 2011 Tohoku-oki megathrust rupture, occurring approximately every 5.5 ± 0.7 years with magnitudes of $M_w 4.9 \pm 0.2$ (Uchida et al., 2012). After the 2011 earthquake, however, their recurrence time interval drastically shortened while rupture magnitudes increased significantly to a maximum of $M_w 5.9$ (Uchida et al., 2015). In this subsection, we aim to compare the effective slip rate of the Kamaishi repeating earthquakes (i.e., the coseismic slip divided by the length of the preceding interseismic period) with the modeled fault slip rate of the patch in our reference model which contains the locations of the repeaters. Coincidentally, the location of the Kamaishi repeaters in our reference model is close to the intersection of four patches (see Fig. 5.8), and lies within a region of strong slip rate gradients. Rather than average the values of the patches neighboring the Kamaishi epicenters, we present here the results of the forward model run on the fine mesh, using the α_h values estimated from the reference model (introduced in Sections 5.2.1 and 5.3.1, described in detail in Appendix C.2).

Fig. 5.29 shows the slip rate in the Kamaishi repeater region based on seismic analysis and our forward model. Before the Tohoku-oki rupture, the slip rate from our inversion matched closely or was slightly larger than the slip rates from Uchida et al. (2012) and Uchida et al. (2015). This is consistent with the interpretation of Uchida et al. (2012) that the source region of the repeating earthquakes is a small asperity surrounded by stable sliding on the interface, and we infer that based on our mod-




Figure 5.29: Slip rate in the Kamaishi region from the seismic analyses of Uchida et al. (2012, orange markers) and Uchida et al. (2015, green markers), and the SEAS forward model results on a fine mesh (blue line) using frictional parameters of the reference model. The date of the 2011 Tohoku-oki earthquake is shown in by the vertical black line.

eled slip rates, approximately all slip deficit is released inside the asperity with each earthquake.

After the rupture, however, the forward-modeled slip rates underestimate initially, and then overestimate the seismically-derived estimates. Qualitatively, the SEAS model includes a finite rise time to reach maximum postseismic slip rate before slowly decaying, while the estimates from Uchida et al. (2012) and Uchida et al. (2015) start at their maximum value and then rapidly decay. This discrepancy may be explained by the difference in considered interface area: The source area of the repeating earthquakes is less than 1 km^2 , while the patch of the fine mesh containing the Kamaishi repeaters is about 420 km^2 . A full investigation into the compatibility of our modeled results with the seismically-derived values is beyond the scope of this dissertation, however.

5.5 Conclusion

In this chapter, we showed results from probabilistic inversions using full-length surface observations timeseries from Northern Japan to infer spatially-variable rheological parameters of the subduction interface. In doing so, we have shown that probabilistic inversions of geodetic data that take into account not just a certain phase of the seismic cycle (e.g., postseismic or interseismic), but all the available

data, are now computationally feasible. Even though some simplifications were necessary (e.g., omission of bulk viscous flow), our model fits the observed data almost as well as published studies with much more complexity (e.g., Fukuda and Johnson, 2021). Given the inherent non-uniqueness when inverting surface observations for fault interface or bulk properties at depth, our results provide an additional confounding factor for the quest of discerning the occurrence and/or significance of various physical processes proposed as drivers of postseismic relaxation. Our model provides estimates of the rate-dependent frictional property $\alpha_h = (a - b)\sigma_E$, but our analysis indicates that the recovered values may be sensitive choices and limitations of our model setup rather than being directly comparable to laboratory-derived values.

Chapter 6

CONCLUSION

6.1 Summary

This dissertation spanned a wide variety of methods, observations, and scales. In Chapter 2, we presented the DISSTANS Python package, a timeseries decomposition software developed to address the need of the geoscientific community to process large datasets efficiently. Included in the package is a wide variety of standard and state-of-the-art methods, as well as a suite of visualization routines to simplify the exploration of GNSS networks and timeseries. DISSTANS was validated in the small region of Long Valley Caldera in California, identifying transient expansion and contraction episodes over a variety of timescales. This new software may provide useful in future studies around the globe, either for standard postprocessing steps, or to recover signals from various physical processes that were previously only attainable using methods tuned to a specific study area, or with great upfront time investment.

For the work in the following Chapters 3–5, we use DISSTANS to analyze over 1400 GNSS stations in Japan to extract a few hundred, high-quality, and long-duration timeseries from Northern Honshu. These timeseries were postprocessed such that they only contained subduction-zone-related motion, i.e., the effect of plate convergence as well as the postseismic transient relaxation occurring since the 2011 M_w 9.1 Tohoku-oki earthquake. Using the probabilistic inverse framework developed for this dissertation, we were able to invert the observed surface deformation in Northern Japan from 1996 to 2024 to recover rheological properties of the megathrust, as well as to produce coupling maps. The insights gained from our results may potentially change the relative importance of different physical mechanisms assumed to occur inside subduction zones. We believe we are the first to simulate entire earthquake cycles *and* fit them to both *pre- and postseismic* data, in a Bayesian probabilistic way (although simplifying the coseismic phase, and excluding it from the inversion). While heavily simplified in some aspects, our method still recovers most of the observed surface deformation signal, and shows that models of this scale are now computationally feasible (ignoring fully-dynamic rupture simulations). The progress presented may therefore represent the next step in plate

boundary deformation modeling.

6.2 Potential for Future Work

In terms of GNSS processing, we focused on Northern Honshu to avoid reference frame and other complications for our cycle inversion framework. However, a post-processed dataset of surface deformation across the entirety of Japan, split into its different constituents, would extract the effects of postseismic relaxation, slow slip events, volcanic events, and the hydrological cycle, across the Japanese islands. Such a trove of information (especially when processed in a way that can continuously update as new data comes in) may prove very valuable for a range of other studies.

We did not include seafloor geodetic stations (e.g., Iinuma et al., 2016) in either our timeseries postprocessing or our cycle inversion results. This decision was partially due because of the additional work necessary to understand how to balance different data sources, but also because the importance of viscoelastic effects to fully account for the significant landward motion of the seafloor observed after the 2011 Tohoku-oki earthquake. In the future, however, it may prove insightful to study the effect of adding seafloor observations on the recovered rheological parameters and their uncertainties. Using seafloor data may be performed even without the addition of viscoelastic relaxation into our framework, although of course, extending our framework to account for bulk flow would constitute another obvious way forward. Viscous flow is currently reserved for methods involving spectral or finite element methods, which are too computationally intensive in 3D for our framework, but progress is being made to incorporate flow into boundary element methods (R. Mallick, personal communication), or it may be approximated by ascribing a rheology to the lower plate interface, and simulating its slip evolution (in contrast to the current, constant slip rate assumption).

Observations of enhanced landward motion (ELM) after large earthquakes (such as the 2011 M_w 9.1 Tohoku-oki earthquake considered in this study) in regions farther away from the main rupture areas have recently added a phenomenon that studies of postseismic motion must take into account (Yuzariyadi and Heki, 2021; Corbi et al., 2022; D'Acquisto et al., 2023; Sun et al., 2024). Using a finite-element model, Sun et al. (2024) were able to show that a weak lithosphere-asthenosphere boundary (LAB) layer would be able to produce both the character and magnitude of the ELM, and point out that the presence of such a layer would require current modeling

frameworks of postseismic transient motion in subduction zones to be revised. In particular, since this layer would be ubiquitous, it would affect the results of every study inverting for rheological properties of megathrusts. The framework presented in this work is inherently equipped to include the potentially-weak LAB because of the mechanical equivalence of a weak LAB and the ESPM's assumption that relative motion between the downgoing slab and the oceanic mantle can be approximated by a dislocation interface. As such, by assigning the lower plate interface a viscous rheology and simulating it in the forward model, the viscosity of the LAB could be estimated. In fact, this is the same approach as if our framework were extended to include viscoelastic effects of the mantle, which could prove to be either a benefit or drawback: On the one hand, the two effects could trade off with each other and make the estimation of the lower interface viscosity ambiguous, if not impossible. On the other hand, it may turn out that a weak LAB may reconcile the landward motion of both the area above coseismic ruptures *and* farther away long the trench, making an extension of our framework into this direction a very valuable scientific target.

For our cycle inversions, we have accounted for the observation uncertainty using a constant, symmetric, diagonal covariance matrix. The omission of more complex covariance information was motivated because of the significant local-scale differences in the behavior of the stations in the GNSS network which would not be reproducible by a simple model setup such as ours. However, another reason to assume observation standard deviations much larger than their actual uncertainties was the unmodeled influence of the prediction error (Minson et al., 2013; Duputel et al., 2015; Langer and Ragon, 2023). This systematic modeling error, and foremost its covariance, is the direct cause of rheological parameters trading off with each other, as well as the large influence some of our assumptions can have (e.g., the assumed plate convergence or the elastic plate thickness). A continuation of this work could take the prediction error into account by numerically estimating the covariance that (un)modeled parameters have on the spatiotemporal pattern of the surface observations. These effects have only been marginally explored so far in Chapter 4 when varying earthquake recurrence times or magnitudes.

We have shown in Section 5.4.3 that our proposed framework is affected by the inherent tradeoffs of simple, homogenous half-space models between horizontal and vertical predicted surface motion (e.g., Hsu et al., 2011; Hashima et al., 2016; Langer and Ragon, 2023; Ragon and Simons, 2023). Depending on whether vertical data is

used in our inversions, we either recover an area of very high frictional parameter at about 150 km depth down-dip of the Tohoku asperity (when using vertical data) or not. The effects of heterogeneous fault structure at depth (e.g., a depth-dependent elastic modulus or the strength contrast between the downgoing slab and the surrounding mantle) are naturally incorporated in postseismic relaxation studies that are based on computationally intensive Finite Element Methods (FEM) (e.g., Sun et al., 2014; Hu et al., 2016; Freed et al., 2017; Muto et al., 2019). Even though our framework is not amenable to using FEM for our forward model, the linear elastic nature of the bulk material in our fault structure enables us to incorporate more realistic fault zone models through the use of displacement and stress kernels that are derived by FEM. Without viscous bulk flow, such kernels are constant in time and can therefore be precomputed; after which, the computational cost of our forward models is comparable to using kernels derived from a homogenous half-space. Using kernels computed from Finite Element Methods that take into account the slab rigidity, a depth-dependent elastic modulus, as well as topography, could therefore be a natural avenue for improvement of our model.

The source of the very low frictional values of the subduction interface recovered in Chapter 5 and are shown in Figs. 5.18 and 5.28 should also be the target of future investigation. The kinematic inversion results, shown in Appendix C.3 which estimate above-plate-rate slip rates throughout the entire postseismic period, hint at a process that is happening on a large scale. It may be the case that we have not sufficiently eliminated the effects of intraplate block motion in Honshu, or that convergence rates between the Pacific and Eurasian plates are not well approximated by an Euler pole rotation. Deciphering the cause (data postprocessing, imperfect kernels, or other assumptions inherent to our model design) could benefit the physical interpretation of our results for the creeping regions of the Northern Japanese megathrust.

To reduce the outsize influence that the assumptions of the location, shape, and size of the chosen, fully-locked asperities have on the model results, another avenue to extend and improve our cycle inversion framework would be to estimate the location of asperities and to relax the assumption that they cannot host any slip between ruptures. Herman and Govers (2020) demonstrated that the probabilistic estimation of binary fault locking probabilities with interseismic velocities is feasible, and that it produces results generally consistent with coupling map interpretations. Estimating the shape, size, and location of the asperities jointly with rheology in our model

could significantly decrease the post-2011 misfit to the observations, and by doing so, allow the inversion to focus more on physically plausible rheological parameters than on finding the least-suboptimal kinematic fit to the data. At the same time, the estimation of asperities using earthquake cycles that inherently take into account stress shadows would be the first of its kind.

Earthquake cycle simulations are naturally geared towards modeling longterm motion and deformation. For the case of Northern Japan, it is still unclear how the observed shortterm, geodetically-derived vertical surface velocities fit into the larger picture of geologically-derived rates of the last thousands of years. Specifically, the subsidence observed in GNSS data in the Northern Honshu region of Sendai before and during the 2011 Tohoku-oki earthquake (see Figs. 1.2 and 1.3) is opposite in sign to the geologic uplift rate derived from marine terraces (e.g., Matsu'ura et al., 2009; Hashima and Sato, 2017). Only the postseismic phase after the Tohoku-oki earthquake (see Fig. 1.4) produces notable uplift in the Sendai region. Assuming the 2011 Tohoku-oki earthquake is characteristic of the recurring, shallow, large earthquakes offshore Northern Honshu, the postseismic phase is therefore the only candidate time period where the longterm uplift can originate. Compared to the geodetically-derived subsidence of about 2 mm/a preseismically (extrapolated onto a 1000-year recurrence time interval) and about 0.4 m coseismically, postseismic uplift would need to amount to 2.6 m to finally yield 0.2 m of anelastic accumulated uplift per cycle (approximately 7%).

In our framework, we rely on the Elastic Subducting Plate Model, in which the downgoing slab subducts steadily in the longterm. Due to the fully-elastic nature of our model, however, it does not reproduce anelastic deformation in the overriding plate that leads to uplift or subsidence on geologic timescales. Any remaining surface deformation after an entire supercycle is purely due to numerical limitations in the integration and fault discretization. Nevertheless, it would be straightforward to introduce a mechanism that “converts” a certain amount of elastic into permanent, anelastic deformation. This framework could then be constrained both by modern (e.g., GNSS- or InSAR-based) and geological surface displacement observations, and could complement existing studies invoking specific physical mechanisms to accumulate longterm deformation (e.g., Hashima and Sato, 2017).

Further technical improvements could include the use of displacement and stress kernels which were not derived from assuming uniform values across each patch, and fitting surface strain (rates) instead of horizontal displacement (rates). The for-

mer could provide a means to avoid stress singularities close to the modeled asperities as discussed in Section 5.3.1 for Fig. 5.8. This may furthermore improve the physical interpretability of the model results, as well as aid in the MCMC sampling process, as numerical artifacts are decreased. The latter would eliminate the dependence of the model results on the assumed overriding reference frame, which is an issue particularly in Northern Japan, where plate boundaries are unclear, and a “stable plate interior” may not exist.

The applicability and potential benefit of our framework extends well beyond Japan. Based on geomorphology, geobiology, and paleoseismology, we today know that most well-studied subduction zones (including the Sumatran Sunda, South American, Japanese, and Cascadia ones) around the world exhibit quasi-periodic recurrence of large megathrust earthquakes embedded in super-imposed cycles spanning many segments (Philibosian and Meltzner, 2020), making them potential targets for our proposed method. For example, the Chilean subduction zone has hosted three M_w 8+ earthquakes in the past 20 years, each rupturing highly-coupled segments (Métois et al., 2016). Applying our method, which is already capable of handling multiple asperities as well as ingesting observations of multiple earthquakes during the observational period, would enable the study of stress interactions between ruptures on differing asperities, and of the potential evolution of plate coupling and slip deficit in time. The latter in particular could provide valuable insights into the behavior of subduction zones that have hosted multiple large earthquakes since the start of modern geodetic measurements.

BIBLIOGRAPHY

- Agata, R., Barbot, S. D., Fujita, K., Hyodo, M., Inuma, T., Nakata, R., Ichimura, T., and Hori, T. (2019). “Rapid Mantle Flow with Power-Law Creep Explains Deformation after the 2011 Tohoku Mega-Quake.” In: *Nature Communications* 10.1, p. 1385. DOI: 10.1038/s41467-019-08984-7.
- Agrawal, A., Verschueren, R., Diamond, S., and Boyd, S. (2018). “A Rewriting System for Convex Optimization Problems.” In: *Journal of Control and Decision* 5.1, pp. 42–60. DOI: 10.1080/23307706.2017.1397554.
- Altamimi, Z., Métivier, L., Reischung, P., Rouby, H., and Collilieux, X. (2017). “ITRF2014 Plate Motion Model.” In: *Geophysical Journal International* 209.3, pp. 1906–1912. DOI: 10.1093/gji/ggx136.
- Ammon, C. J., Lay, T., Kanamori, H., and Cleveland, M. (2011). “A Rupture Model of the 2011 off the Pacific Coast of Tohoku Earthquake.” In: *Earth, Planets and Space* 63.7, p. 33. DOI: 10.5047/eps.2011.05.015.
- Aoki, Y. and Scholz, C. H. (2003). “Vertical Deformation of the Japanese Islands, 1996–1999.” In: *Journal of Geophysical Research: Solid Earth* 108.B5. DOI: 10.1029/2002JB002129.
- Bailey, R. A. (1989). *Geologic Map of the Long Valley Caldera, Mono-Inyo Craters Volcanic Chain, and Vicinity, Eastern California*. USGS Numbered Series 1933.
- Bedford, J. and Bevis, M. (2018). “Greedy Automatic Signal Decomposition and Its Application to Daily GPS Time Series.” In: *Journal of Geophysical Research: Solid Earth* 123.8, pp. 6992–7003. DOI: 10.1029/2017JB014765.
- Bekaert, D. P. S., Segall, P., Wright, T. J., and Hooper, A. J. (2016). “A Network Inversion Filter Combining GNSS and InSAR for Tectonic Slip Modeling.” In: *Journal of Geophysical Research: Solid Earth* 121.3, pp. 2069–2086. DOI: 10.1002/2015JB012638.
- Bercovici, D., Mulyukova, E., Girard, J., and Skemer, P. (2023). “A Coupled Model for Phase Mixing, Grain Damage and Shear Localization in the Lithosphere: Comparison to Lab Experiments.” In: *Geophysical Journal International* 232.3, pp. 2205–2230. DOI: 10.1093/gji/ggac428.
- Bertiger, W., Bar-Sever, Y., Dorsey, A., Haines, B., Harvey, N., Hemberger, D., Heflin, M., Lu, W., Miller, M., Moore, A. W., Murphy, D., Ries, P., Romans, L., Sibois, A., Sibthorpe, A., Szilagyi, B., Vallisneri, M., and Willis, P. (2020). “GipsyX/RTGx, a New Tool Set for Space Geodetic Operations and Research.” In: *Advances in Space Research* 66.3, pp. 469–489. DOI: 10.1016/j.asr.2020.04.015.
- Bird, P. (2003). “An Updated Digital Model of Plate Boundaries.” In: *Geochemistry, Geophysics, Geosystems* 4.3. DOI: 10.1029/2001GC000252.

- Blanpied, M. L., Lockner, D. A., and Byerlee, J. D. (1995). “Frictional Slip of Granite at Hydrothermal Conditions.” In: *Journal of Geophysical Research: Solid Earth* 100.B7, pp. 13045–13064. DOI: 10.1029/95JB00862.
- Blewitt, G. (2015). “GPS and Space-Based Geodetic Methods.” In: *Treatise on Geophysics (Second Edition)*. Ed. by G. Schubert. Oxford: Elsevier, pp. 307–338. DOI: 10.1016/B978-0-444-53802-4.00060-9.
- Blewitt, G., Hammond, W., and Kreemer, C. (2018). “Harnessing the GPS Data Explosion for Interdisciplinary Science.” In: *Eos* 99. DOI: 10.1029/2018E0104623.
- Blewitt, G., Kreemer, C., Hammond, W. C., and Gazeaux, J. (2016). “MIDAS Robust Trend Estimator for Accurate GPS Station Velocities without Step Detection.” In: *Journal of Geophysical Research: Solid Earth* 121.3, pp. 2054–2068. DOI: 10.1002/2015JB012552.
- Bürgmann, R. and Dresen, G. (2008). “Rheology of the Lower Crust and Upper Mantle: Evidence from Rock Mechanics, Geodesy, and Field Observations.” In: *Annual Review of Earth and Planetary Sciences* 36.1, pp. 531–567. DOI: 10.1146/annurev.earth.36.031207.124326.
- Burnham, K. P. and Anderson, D. R. (2002). “Information and Likelihood Theory: A Basis for Model Selection and Inference.” In: *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. New York, NY: Springer, pp. 49–97. DOI: 10.1007/978-0-387-22456-5_2.
- Butzer, P. L., Schmidt, M., and Stark, E. L. (1988). “Observations on the History of Central B-Splines.” In: *Archive for History of Exact Sciences* 39.2, pp. 137–156. JSTOR: 41133848.
- Candès, E. J., Wakin, M. B., and Boyd, S. P. (2008). “Enhancing Sparsity by Reweighted ℓ_1 Minimization.” In: *Journal of Fourier Analysis and Applications* 14.5, pp. 877–905. DOI: 10.1007/s00041-008-9045-x.
- Chopra, P. N. and Paterson, M. S. (1981). “The Experimental Deformation of Dunite.” In: *Tectonophysics. The Effect of Deformation on Rocks* 78.1, pp. 453–473. DOI: 10.1016/0040-1951(81)90024-X.
- Chopra, P. N. (1997). “High-Temperature Transient Creep in Olivine Rocks.” In: *Tectonophysics. The Adolphe Nicolas Volume* 279.1, pp. 93–111. DOI: 10.1016/S0040-1951(97)00134-0.
- Corbi, F., Bedford, J., Poli, P., Funiciello, F., and Deng, Z. (2022). “Probing the Seismic Cycle Timing with Coseismic Twisting of Subduction Margins.” In: *Nature Communications* 13.1, p. 1911. DOI: 10.1038/s41467-022-29564-2.
- Coreform LLC (2023). *Coreform Cubit*. Orem, UT.
- Crouch, S. L. and Starfield, A. M. (1983). *Boundary Element Methods in Solid Mechanics: With Applications in Rock Mechanics and Geological Engineering*. 1st edition. London ; Boston: George Allen & Unwin.

- D'Acquisto, M., Herman, M. W., Riva, R. E. M., and Govers, R. (2023). "Can Plate Bending Explain the Observed Faster Landward Motion of Lateral Regions of the Subduction Zone After Major Megathrust Earthquakes?" In: *Journal of Geophysical Research: Solid Earth* 128.3, e2022JB025431. DOI: 10.1029/2022JB025431.
- Dal Zilio, L., Lapusta, N., Avouac, J.-P., and Gerya, T. (2022). "Subduction Earthquake Sequences in a Non-Linear Visco-Elasto-Plastic Megathrust." In: *Geophysical Journal International* 229.2, pp. 1098–1121. DOI: 10.1093/gji/ggab521.
- Davis, T. (2017). *A New Open Source Boundary Element Code and Its Application to Geological Deformation: Exploring Stress Concentrations around Voids and the Effects of 3D Frictional Distributions on Fault Surfaces*. Aberdeen University.
- Diamond, S. and Boyd, S. (2016). "CVXPY: A Python-Embedded Modeling Language for Convex Optimization." In: *Journal of Machine Learning Research* 17, p. 83.
- Dieterich, J. H. (1979). "Modeling of Rock Friction: 1. Experimental Results and Constitutive Equations." In: *Journal of Geophysical Research: Solid Earth* 84.B5, pp. 2161–2168. DOI: 10.1029/JB084iB05p02161.
- Dong, D., Fang, P., Bock, Y., Webb, F., Prawirodirdjo, L., Kedar, S., and Jamason, P. (2006). "Spatiotemporal Filtering Using Principal Component Analysis and Karhunen-Loeve Expansion Approaches for Regional GPS Network Analysis." In: *Journal of Geophysical Research: Solid Earth* 111.B3. DOI: 10.1029/2005JB003806.
- Dong, D., Herring, T. A., and King, R. W. (1998). "Estimating Regional Deformation from a Combination of Space and Terrestrial Geodetic Data." In: *Journal of Geodesy* 72.4, pp. 200–214. DOI: 10.1007/s001900050161.
- Duputel, Z., Jiang, J., Jolivet, R., Simons, M., Rivera, L., Ampuero, J.-P., Riel, B., Owen, S. E., Moore, A. W., Samsonov, S. V., Ortega Culaciati, F., and Minson, S. E. (2015). "The Iquique Earthquake Sequence of April 2014: Bayesian Modeling Accounting for Prediction Uncertainty." In: *Geophysical Research Letters* 42.19, pp. 7949–7957. DOI: 10.1002/2015GL065402.
- Duputel, Z., Agram, P. S., Simons, M., Minson, S. E., and Beck, J. L. (2014). "Accounting for Prediction Uncertainty When Inferring Subsurface Fault Slip." In: *Geophysical Journal International* 197.1, pp. 464–482. DOI: 10.1093/gji/ggt517.
- Elson, P., Andrade, E. S. de, Lucas, G., May, R., Hattersley, R., Campbell, E., Dawson, A., Raynaud, S., scmc72, Little, B., Snow, A. D., Donkers, K., Blay, B., Killick, P., Wilson, N., Peglar, P., Ibdreyer, Andrew, Szymaniak, J., Berchet, A., Bosley, C., Davis, L., Filipe, Krasting, J., Bradbury, M., Kirkham, D., stephenworsley, Clément, Caria, G., and Hedley, M. (2022). *SciTools/Cartopy: V0.20.2*. Zenodo. DOI: 10.5281/zenodo.5842769.

- Fagereng, Å. and Beall, A. (2021). “Is Complex Fault Zone Behaviour a Reflection of Rheological Heterogeneity?” In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2193, p. 20190421. DOI: 10.1098/rsta.2019.0421.
- Freed, A. M., Hashima, A., Becker, T. W., Okaya, D. A., Sato, H., and Hatanaka, Y. (2017). “Resolving Depth-Dependent Subduction Zone Viscosity and Afterslip from Postseismic Displacements Following the 2011 Tohoku-oki, Japan Earthquake.” In: *Earth and Planetary Science Letters* 459, pp. 279–290. DOI: 10.1016/j.epsl.2016.11.040.
- Fujii, Y., Satake, K., Sakai, S., Shinohara, M., and Kanazawa, T. (2011). “Tsunami Source of the 2011 off the Pacific Coast of Tohoku Earthquake.” In: *Earth, Planets and Space* 63.7, pp. 815–820. DOI: 10.5047/eps.2011.06.010.
- Fukuda, J. and Johnson, K. M. (2008). “A Fully Bayesian Inversion for Spatial Distribution of Fault Slip with Objective Smoothing.” In: *Bulletin of the Seismological Society of America* 98.3, pp. 1128–1146. DOI: 10.1785/0120070194.
- Fukuda, J. and Johnson, K. M. (2021). “Bayesian Inversion for a Stress-Driven Model of Afterslip and Viscoelastic Relaxation: Method and Application to Postseismic Deformation Following the 2011 MW 9.0 Tohoku-Oki Earthquake.” In: *Journal of Geophysical Research: Solid Earth* 126.5, e2020JB021620. DOI: 10.1029/2020JB021620.
- Gazeaux, J., Williams, S., King, M., Bos, M., Dach, R., Deo, M., Moore, A. W., Ostini, L., Petrie, E., Roggero, M., Teferle, F. N., Olivares, G., and Webb, F. H. (2013). “Detecting Offsets in GPS Time Series: First Results from the Detection of Offsets in GPS Experiment.” In: *Journal of Geophysical Research: Solid Earth* 118.5, pp. 2397–2407. DOI: 10.1002/jgrb.50152.
- Goudarzi, M. A., Cocard, M., and Santerre, R. (2014). “EPC: Matlab Software to Estimate Euler Pole Parameters.” In: *GPS Solutions* 18.1, pp. 153–162. DOI: 10.1007/s10291-013-0354-4.
- Govers, R., Furlong, K. P., van de Wiel, L., Herman, M. W., and Broerse, T. (2018). “The Geodetic Signature of the Earthquake Cycle at Subduction Zones: Model Constraints on the Deep Processes.” In: *Reviews of Geophysics* 56.1, pp. 6–49. DOI: 10.1002/2017RG000586.
- Griffin, J. D., Stirling, M. W., and Wang, T. (2020). “Periodicity and Clustering in the Long-Term Earthquake Record.” In: *Geophysical Research Letters* 47.22. DOI: 10.1029/2020GL089272.
- Gualandi, A., Serpelloni, E., and Belardinelli, M. E. (2016). “Blind Source Separation Problem in GPS Time Series.” In: *Journal of Geodesy* 90.4, pp. 323–341. DOI: 10.1007/s00190-015-0875-4.

- Gusman, A. R., Tanioka, Y., Sakai, S., and Tsushima, H. (2012). “Source Model of the Great 2011 Tohoku Earthquake Estimated from Tsunami Waveforms and Crustal Deformation Data.” In: *Earth and Planetary Science Letters* 341–344, pp. 234–242. DOI: 10.1016/j.epsl.2012.06.006.
- Handy, M. R. (1994). “Flow Laws for Rocks Containing Two Non-Linear Viscous Phases: A Phenomenological Approach.” In: *Journal of Structural Geology* 16.3, pp. 287–301. DOI: 10.1016/0191-8141(94)90035-3.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. (2020). “Array Programming with NumPy.” In: *Nature* 585.7825, pp. 357–362. DOI: 10.1038/s41586-020-2649-2.
- Harris, R. A. and Segall, P. (1987). “Detection of a Locked Zone at Depth on the Parkfield, California, Segment of the San Andreas Fault.” In: *Journal of Geophysical Research: Solid Earth* 92.B8, pp. 7945–7962. DOI: 10.1029/JB092iB08p07945.
- Hashima, A., Becker, T. W., Freed, A. M., Sato, H., and Okaya, D. A. (2016). “Co-seismic Deformation Due to the 2011 Tohoku-oki Earthquake: Influence of 3-D Elastic Structure around Japan.” In: *Earth, Planets and Space* 68.1, p. 159. DOI: 10.1186/s40623-016-0535-9.
- Hashima, A. and Sato, T. (2017). “A Megathrust Earthquake Cycle Model for Northeast Japan: Bridging the Mismatch between Geological Uplift and Geodetic Subsidence.” In: *Earth, Planets and Space* 69.1, p. 23. DOI: 10.1186/s40623-017-0606-6.
- Hayes, G. P. (2011). “Rapid Source Characterization of the 2011 Mw 9.0 off the Pacific Coast of Tohoku Earthquake.” In: *Earth, Planets and Space* 63.7, pp. 529–534. DOI: 10.5047/eps.2011.05.012.
- (2017). “The Finite, Kinematic Rupture Properties of Great-Sized Earthquakes since 1990.” In: *Earth and Planetary Science Letters* 468, pp. 94–100. DOI: 10.1016/j.epsl.2017.04.003.
- Hayes, G. P., Moore, G. L., Portner, D. E., Hearne, M., Flamme, H., Furtney, M., and Smoczyk, G. M. (2018). “Slab2, a Comprehensive Subduction Zone Geometry Model.” In: *Science* 362.6410, pp. 58–61. DOI: 10.1126/science.aat4723.
- Hearn, E. H. and Thatcher, W. R. (2015). “Reconciling Viscoelastic Models of Post-seismic and Interseismic Deformation: Effects of Viscous Shear Zones and Finite Length Ruptures.” In: *Journal of Geophysical Research: Solid Earth* 120.4, pp. 2794–2819. DOI: 10.1002/2014JB011361.
- Heflin, M., Donnellan, A., Parker, J., Lyzenga, G., Moore, A., Ludwig, L. G., Rundle, J., Wang, J., and Pierce, M. (2020). “Automated Estimation and Tools to Extract Positions, Velocities, Breaks, and Seasonal Terms From Daily GNSS Mea-

- surements: Illuminating Nonlinear Salton Trough Deformation.” In: *Earth and Space Science* 7.7, e2019EA000644. DOI: 10.1029/2019EA000644.
- Herman, M. W. and Govers, R. (2020). “Locating Fully Locked Asperities Along the South America Subduction Megathrust: A New Physical Interseismic Inversion Approach in a Bayesian Framework.” In: *Geochemistry, Geophysics, Geosystems* 21.8, e2020GC009063. DOI: 10.1029/2020GC009063.
- Herring, T. A., King, R. W., Floyd, M. A., and McClusky, S. C. (2018). *Introduction to GAMIT/GLOBK*. Massachusetts Institute of Technology. Cambridge, MA, USA.
- Herring, T. A., Melbourne, T. I., Murray, M. H., Floyd, M. A., Szeliga, W. M., King, R. W., Phillips, D. A., Puskas, C. M., Santillan, M., and Wang, L. (2016). “Plate Boundary Observatory and Related Networks: GPS Data Analysis Methods and Geodetic Products.” In: *Reviews of Geophysics* 54.4, pp. 759–808. DOI: 10.1002/2016RG000529.
- Hetland, E. A. and Hager, B. H. (2005). “Postseismic and Interseismic Displacements near a Strike-Slip Fault: A Two-Dimensional Theory for General Linear Viscoelastic Rheologies.” In: *Journal of Geophysical Research: Solid Earth* 110.B10. DOI: 10.1029/2005JB003689.
- Hetland, E. A., Musé, P., Simons, M., Lin, Y. N., Agram, P. S., and DiCaprio, C. J. (2012). “Multiscale InSAR Time Series (MInTS) Analysis of Surface Deformation.” In: *Journal of Geophysical Research: Solid Earth* 117.B2. DOI: 10.1029/2011JB008731.
- Hetland, E. A. and Simons, M. (2010). “Post-Seismic and Interseismic Fault Creep II: Transient Creep and Interseismic Stress Shadows on Megathrusts.” In: *Geophysical Journal International* 181.1, pp. 99–112. DOI: 10.1111/j.1365-246X.2009.04482.x.
- Hetland, E. A., Simons, M., and Dunham, E. M. (2010). “Post-Seismic and Interseismic Fault Creep I: Model Description.” In: *Geophysical Journal International* 181.1, pp. 81–98. DOI: 10.1111/j.1365-246X.2010.04522.x.
- Hirth, G. (2002). “Laboratory Constraints on the Rheology of the Upper Mantle.” In: *Reviews in Mineralogy and Geochemistry* 51.1, pp. 97–120. DOI: 10.2138/gsrmg.51.1.97.
- Hirth, G. and Kohlstedt, D. (2004). “Rheology of the Upper Mantle and the Mantle Wedge: A View from the Experimentalists.” In: *Inside the Subduction Factory*. American Geophysical Union (AGU), pp. 83–105. DOI: 10.1029/138GM06.
- Houston, H., Delbridge, B. G., Wech, A. G., and Creager, K. C. (2011). “Rapid Tremor Reversals in Cascadia Generated by a Weakened Plate Interface.” In: *Nature Geoscience* 4.6, pp. 404–409. DOI: 10.1038/ngeo1157.

- Hsu, Y.-J., Simons, M., Williams, C., and Casarotti, E. (2011). “Three-Dimensional FEM Derived Elastic Green’s Functions for the Coseismic Deformation of the 2005 Mw 8.7 Nias-Simeulue, Sumatra Earthquake.” In: *Geochemistry, Geophysics, Geosystems* 12.7. DOI: 10.1029/2011GC003553.
- Hsu, Y.-J., Yu, S.-B., and Chen, H.-Y. (2009). “Coseismic and Postseismic Deformation Associated with the 2003 Chengkung, Taiwan, Earthquake.” In: *Geophysical Journal International* 176.2, pp. 420–430. DOI: 10.1111/j.1365-246X.2008.04009.x.
- Hu, Y., Bürgmann, R., Uchida, N., Banerjee, P., and Freymueller, J. T. (2016). “Stress-Driven Relaxation of Heterogeneous Upper Mantle and Time-Dependent Afterslip Following the 2011 Tohoku Earthquake.” In: *Journal of Geophysical Research: Solid Earth* 121.1, pp. 385–411. DOI: 10.1002/2015JB012508.
- Huang, D. W., Dai, W. J., and Luo, F. X. (2012). “ICA Spatiotemporal Filtering Method and Its Application in GPS Deformation Monitoring.” In: *Applied Mechanics and Materials* 204–208, pp. 2806–2812. DOI: 10.4028/www.scientific.net/AMM.204-208.2806.
- Hunter, J. D. (2007). “Matplotlib: A 2D Graphics Environment.” In: *Computing in Science Engineering* 9.3, pp. 90–95. DOI: 10.1109/MCSE.2007.55.
- Ide, S., Baltay, A., and Beroza, G. C. (2011). “Shallow Dynamic Overshoot and Energetic Deep Rupture in the 2011 M_w 9.0 Tohoku-Oki Earthquake.” In: *Science* 332.6036, pp. 1426–1429. DOI: 10.1126/science.1207020.
- Iinuma, T., Hino, R., Uchida, N., Nakamura, W., Kido, M., Osada, Y., and Miura, S. (2016). “Seafloor Observations Indicate Spatial Separation of Coseismic and Postseismic Slips in the 2011 Tohoku Earthquake.” In: *Nature Communications* 7.1, p. 13506. DOI: 10.1038/ncomms13506.
- Ingleby, T. and Wright, T. J. (2017). “Omori-like Decay of Postseismic Velocities Following Continental Earthquakes.” In: *Geophysical Research Letters* 44.7, pp. 3119–3130. DOI: 10.1002/2017GL072865.
- Ji, K. H., Herring, T. A., and Llenos, A. L. (2013). “Near Real-Time Monitoring of Volcanic Surface Deformation from GPS Measurements at Long Valley Caldera, California.” In: *Geophysical Research Letters* 40.6, pp. 1054–1058. DOI: 10.1002/grl.50258.
- Johnson, K. M. and Tebo, D. (2018). “Capturing 50 Years of Postseismic Mantle Flow at Nankai Subduction Zone.” In: *Journal of Geophysical Research: Solid Earth* 123.11, pp. 10,091–10,106. DOI: 10.1029/2018JB016345.
- Jolivet, R., Duputel, Z., Riel, B., Simons, M., Rivera, L., Minson, S. E., Zhang, H., Aivazis, M. A. G., Ayoub, F., Leprince, S., Samsonov, S., Motagh, M., and Fielding, E. J. (2014). “The 2013 Mw 7.7 Balochistan Earthquake: Seismic Potential of an Accretionary Wedge.” In: *Bulletin of the Seismological Society of America* 104.2, pp. 1020–1030. DOI: 10.1785/0120130313.

- Jolivet, R., Simons, M., Agram, P. S., Duputel, Z., and Shen, Z.-K. (2015). “Aseismic Slip and Seismogenic Coupling along the Central San Andreas Fault.” In: *Geophysical Research Letters* 42.2, pp. 297–306. DOI: 10.1002/2014GL062222.
- Kanda, R. V. S., Hetland, E. A., and Simons, M. (2013). “An Asperity Model for Fault Creep and Interseismic Deformation in Northeastern Japan.” In: *Geophysical Journal International* 192.1, pp. 38–57. DOI: 10.1093/gji/ggs028.
- Kanda, R. V. S. and Simons, M. (2010). “An Elastic Plate Model for Interseismic Deformation in Subduction Zones.” In: *Journal of Geophysical Research: Solid Earth* 115.B3. DOI: 10.1029/2009JB006611.
- Kano, M. and Kato, A. (2020). “Detailed Spatial Slip Distribution for Short-Term Slow Slip Events Along the Nankai Subduction Zone, Southwest Japan.” In: *Journal of Geophysical Research: Solid Earth* 125.7, e2020JB019613. DOI: 10.1029/2020JB019613.
- Karato, S.-i. (2021). “A Theory of Inter-Granular Transient Dislocation Creep: Implications for the Geophysical Studies on Mantle Rheology.” In: *Journal of Geophysical Research: Solid Earth* 126.10, e2021JB022763. DOI: 10.1029/2021JB022763.
- Karato, S.-I., Paterson, M. S., and FitzGerald, J. D. (1986). “Rheology of Synthetic Olivine Aggregates: Influence of Grain Size and Water.” In: *Journal of Geophysical Research: Solid Earth* 91.B8, pp. 8151–8176. DOI: 10.1029/JB091iB08p08151.
- Kobayashi, H., Koketsu, K., Miyake, H., and Kanamori, H. (2021). “Similarities and Differences in the Rupture Processes of the 1952 and 2003 Tokachi-Oki Earthquakes.” In: *Journal of Geophysical Research: Solid Earth* 126.1, e2020JB020585. DOI: 10.1029/2020JB020585.
- Köhne, T., Riel, B., and Simons, M. (2023). “Decomposition and Inference of Sources through Spatiotemporal Analysis of Network Signals: The DISSTANS Python Package.” In: *Computers & Geosciences* 170, p. 105247. DOI: 10.1016/j.cageo.2022.105247.
- Koketsu, K., Hikima, K., Miyazaki, S., and Ide, S. (2004). “Joint Inversion of Strong Motion and Geodetic Data for the Source Process of the 2003 Tokachi-oki, Hokkaido, Earthquake.” In: *Earth, Planets and Space* 56.3, pp. 329–334. DOI: 10.1186/BF03353060.
- Kositsky, A. P. and Avouac, J.-P. (2010). “Inverting Geodetic Time Series with a Principal Component Analysis-Based Inversion Method.” In: *Journal of Geophysical Research: Solid Earth* 115.B3. DOI: 10.1029/2009JB006535.
- Langbein, J. (2020). “Methods for Rapidly Estimating Velocity Precision from GNSS Time Series in the Presence of Temporal Correlation: A New Method and Comparison of Existing Methods.” In: *Journal of Geophysical Research: Solid Earth* 125.7, e2019JB019132. DOI: 10.1029/2019JB019132.

- Langer, L. and Ragon, T. (2023). “Accuracy of Finite Fault Slip Estimates in Subduction Zone Regions with Topographic Green’s Functions and Seafloor Geodesy.” In: *Journal of Geophysical Research: Solid Earth*, e2023JB026559. DOI: 10.1029/2023JB026559.
- Lay, T., Ammon, C. J., Kanamori, H., Xue, L., and Kim, M. J. (2011). “Possible Large Near-Trench Slip during the 2011 Mw 9.0 off the Pacific Coast of Tohoku Earthquake.” In: *Earth, Planets and Space* 63.7, p. 32. DOI: 10.5047/eps.2011.05.033.
- Lin, W., Conin, M., Moore, J. C., Chester, F. M., Nakamura, Y., Mori, J. J., Anderson, L., Brodsky, E. E., Eguchi, N., Expedition 343 Scientists, Cook, B., Jeppson, T., Wolfson-Schwehr, M., Sanada, Y., Saito, S., Kido, Y., Hirose, T., Behrmann, J. H., Ikari, M., Ujiie, K., Rowe, C., Kirkpatrick, J., Bose, S., Regalla, C., Remitti, F., Toy, V., Fulton, P., Mishima, T., Yang, T., Sun, T., Ishikawa, T., Sample, J., Takai, K., Kameda, J., Toczko, S., Maeda, L., Kodaira, S., Hino, R., and Saffer, D. (2013). “Stress State in the Largest Displacement Area of the 2011 Tohoku-Oki Earthquake.” In: *Science* 339.6120, pp. 687–690. DOI: 10.1126/science.1229379.
- Lindsey, E. O., Mallick, R., Hubbard, J. A., Bradley, K. E., Almeida, R. V., Moore, J. D. P., Bürgmann, R., and Hill, E. M. (2021). “Slip Rate Deficit and Earthquake Potential on Shallow Megathrusts.” In: *Nature Geoscience*, pp. 1–6. DOI: 10.1038/s41561-021-00736-x.
- Loveless, J. P. and Meade, B. J. (2010). “Geodetic Imaging of Plate Motions, Slip Rates, and Partitioning of Deformation in Japan.” In: *Journal of Geophysical Research: Solid Earth* 115.B2. DOI: 10.1029/2008JB006248.
- (2016). “Two Decades of Spatiotemporal Variations in Subduction Zone Coupling Offshore Japan.” In: *Earth and Planetary Science Letters* 436, pp. 19–30. DOI: 10.1016/j.epsl.2015.12.033.
- Mai, P. M. and Thingbaijam, K. K. S. (2014). “SRCMOD: An Online Database of Finite-Fault Rupture Models.” In: *Seismological Research Letters* 85.6, pp. 1348–1357. DOI: 10.1785/0220140077.
- Mallick, R., Lambert, V., and Meade, B. (2022). “On the Choice and Implications of Rheologies That Maintain Kinematic and Dynamic Consistency Over the Entire Earthquake Cycle.” In: *Journal of Geophysical Research: Solid Earth* 127.9, e2022JB024683. DOI: 10.1029/2022JB024683.
- Malvern, L. E. (1969). *Introduction to the Mechanics of a Continuous Medium*. Prentice-Hall Series in Engineering of the Physical Sciences. Englewood Cliffs, N.J.: Prentice-Hall.
- Marone, C. (1998). “Laboratory-Derived Friction Laws and Their Application to Seismic Faulting.” In: *Annual Review of Earth and Planetary Sciences* 26.1, pp. 643–696. DOI: 10.1146/annurev.earth.26.1.643.

- Marone, C. J., Scholtz, C. H., and Bilham, R. (1991). “On the Mechanics of Earthquake Afterslip.” In: *Journal of Geophysical Research: Solid Earth* 96.B5, pp. 8441–8452. DOI: 10.1029/91JB00275.
- Masuti, S. and Barbot, S. (2021). “MCMC Inversion of the Transient and Steady-State Creep Flow Law Parameters of Dunite under Dry and Wet Conditions.” In: *Earth, Planets and Space* 73.1, p. 208. DOI: 10.1186/s40623-021-01543-9.
- Masuti, S., Barbot, S. D., Karato, S.-i., Feng, L., and Banerjee, P. (2016). “Upper-Mantle Water Stratification Inferred from Observations of the 2012 Indian Ocean Earthquake.” In: *Nature* 538.7625, pp. 373–377. DOI: 10.1038/nature19783.
- Masuti, S., Muto, J., and Rybacki, E. (2023). “Transient Creep of Quartz and Granulite at High Temperature Under Wet Conditions.” In: *Journal of Geophysical Research: Solid Earth* 128.10, e2023JB027762. DOI: 10.1029/2023JB027762.
- Matsu’ura, T., Furusawa, A., and Saomoto, H. (2009). “Long-Term and Short-Term Vertical Velocity Profiles across the Forearc in the NE Japan Subduction Zone.” In: *Quaternary Research* 71.2, pp. 227–238. DOI: 10.1016/j.yqres.2008.12.005.
- McGuire, J. J. and Segall, P. (2003). “Imaging of Aseismic Fault Slip Transients Recorded by Dense Geodetic Networks.” In: *Geophysical Journal International* 155.3, pp. 778–788. DOI: 10.1111/j.1365-246X.2003.02022.x.
- Mckenzie, D. (1977). “The Initiation of Trenches.” In: *Island Arcs, Deep Sea Trenches and Back-Arc Basins*. American Geophysical Union (AGU), pp. 57–61. DOI: 10.1029/ME001p0057.
- McKinney, W. (2010). “Data Structures for Statistical Computing in Python.” In: *Proceedings of the 9th Python in Science Conference*, pp. 56–61. DOI: 10.25080/Majora-92bf1922-00a.
- Meade, B. J. and Hager, B. H. (2005). “Block Models of Crustal Motion in Southern California Constrained by GPS Measurements.” In: *Journal of Geophysical Research: Solid Earth* 110.B3. DOI: 10.1029/2004JB003209.
- Métois, M., Vigny, C., and Socquet, A. (2016). “Interseismic Coupling, Megathrust Earthquakes and Seismic Swarms Along the Chilean Subduction Zone (38°–18°S).” In: *Pure and Applied Geophysics* 173.5, pp. 1431–1449. DOI: 10.1007/s00024-016-1280-5.
- Minson, S. E., Simons, M., and Beck, J. L. (2013). “Bayesian Inversion for Finite Fault Earthquake Source Models I—Theory and Algorithm.” In: *Geophysical Journal International* 194.3, pp. 1701–1726. DOI: 10.1093/gji/ggt180.
- Misra, P. and Enge, P. (2010). *Global Positioning System: Signals, Measurements, and Performance*. Lincoln, Mass: Ganga-Jamuna Press.
- Miura, S., Iinuma, T., Yui, S., Uchida, N., Sato, T., Tachibana, K., and Hasegawa, A. (2006). “Co- and Post-Seismic Slip Associated with the 2005 Miyagi-oki Earthquake (M7.2) as Inferred from GPS Data.” In: *Earth, Planets and Space* 58.12, pp. 1567–1572. DOI: 10.1186/BF03352662.

- Montési, L. G. J. (2004). “Controls of Shear Zone Rheology and Tectonic Loading on Postseismic Creep.” In: *Journal of Geophysical Research: Solid Earth* 109.B10. DOI: 10.1029/2003JB002925.
- Montési, L. G. J. and Hirth, G. (2003). “Grain Size Evolution and the Rheology of Ductile Shear Zones: From Laboratory Experiments to Postseismic Creep.” In: *Earth and Planetary Science Letters* 211.1, pp. 97–110. DOI: 10.1016/S0012-821X(03)00196-1.
- Montgomery-Brown, E. K., Wicks, C. W., Cervelli, P. F., Langbein, J. O., Svarc, J. L., Shelly, D. R., Hill, D. P., and Lisowski, M. (2015). “Renewed Inflation of Long Valley Caldera, California (2011 to 2014).” In: *Geophysical Research Letters* 42.13, pp. 5250–5257. DOI: 10.1002/2015GL064338.
- Moore, A. (2024). *JPL GEONET Japan GPS Position Timeseries (2024)*. DOI: 10.48577/jpl.NAHL2B.
- Moore, J. D. and Parsons, B. (2015). “Scaling of Viscous Shear Zones with Depth-Dependent Viscosity and Power-Law Stress–Strain-Rate Dependence.” In: *Geophysical Journal International* 202.1, pp. 242–260. DOI: 10.1093/gji/ggv143.
- Mueller, S. and Phillips, R. J. (1991). “On The Initiation of Subduction.” In: *Journal of Geophysical Research: Solid Earth* 96.B1, pp. 651–665. DOI: 10.1029/90JB02237.
- Müller, G. (1986). “Generalized Maxwell Bodies and Estimates of Mantle Viscosity.” In: *Geophysical Journal International* 87.3, pp. 1113–1141. DOI: 10.1111/j.1365-246X.1986.tb01986.x.
- Muto, J., Moore, J. D. P., Barbot, S., Iinuma, T., Ohta, Y., and Iwamori, H. (2019). “Coupled Afterslip and Transient Mantle Flow after the 2011 Tohoku Earthquake.” In: *Science Advances* 5.9, eaaw1164. DOI: 10.1126/sciadv.aaw1164.
- Nagai, R., Kikuchi, M., and Yamanaka, Y. (2001). “Comparative Study on the Source Processes of Recurrent Large Earthquakes in Sanriku-oki Region : The 1968 Tokachi-oki Earthquake and the 1994 Sanriku-oki Earthquake.” In: *Comparative study on the source processes of recurrent large earthquakes in Sanriku-oki region : The 1968 Tokachi-oki earthquake and the 1994 Sanriku-oki earthquake* 54.2, pp. 267–280.
- Nakayama, W. and Takeo, M. (1997). “Slip History of the 1994 Sanriku-Haruka-Oki, Japan, Earthquake Deduced from Strong-Motion Data.” In: *Bulletin of the Seismological Society of America* 87.4, pp. 918–931. DOI: 10.1785/BSSA0870040918.
- Nanayama, F., Satake, K., Furukawa, R., Shimokawa, K., Atwater, B. F., Shigeno, K., and Yamaki, S. (2003). “Unusually Large Earthquakes Inferred from Tsunami Deposits along the Kuril Trench.” In: *Nature* 424.6949, pp. 660–663. DOI: 10.1038/nature01864.

- Nishimura, T. (2009). “Slip Distribution of the 1973 Nemuro-oki Earthquake Estimated from the Re-Examined Geodetic Data.” In: *Earth, Planets and Space* 61.11, pp. 1203–1214. DOI: 10.1186/BF03352973.
- NOAA National Centers for Environmental Information (2022). *ETOPO 2022 15 Arc-Second Global Relief Model*. DOI: 10.25921/FD45-GT74.
- Ortega-Culaciati, F. H. (2013). “Aseismic Deformation in Subduction Megathrusts : Central Andes and North-East Japan.” PhD thesis. Pasadena, CA, USA: California Institute of Technology.
- Ozawa, S., Yarai, H., Imakiire, T., and Tobita, M. (2013). “Spatial and Temporal Evolution of the Long-Term Slow Slip in the Bungo Channel, Japan.” In: *Earth, Planets and Space* 65.2, p. 67. DOI: 10.5047/eps.2012.06.009.
- Peña, C., Heidbach, O., Moreno, M., Bedford, J., Ziegler, M., Tassara, A., and Oncken, O. (2019). “Role of Lower Crust in the Postseismic Deformation of the 2010 Maule Earthquake: Insights from a Model with Power-Law Rheology.” In: *Pure and Applied Geophysics* 176.9, pp. 3913–3928. DOI: 10.1007/s00024-018-02090-3.
- (2020). “Impact of Power-Law Rheology on the Viscoelastic Relaxation Pattern and Afterslip Distribution Following the 2010 Mw 8.8 Maule Earthquake.” In: *Earth and Planetary Science Letters* 542, p. 116292. DOI: 10.1016/j.epsl.2020.116292.
- Perfettini, H. and Avouac, J. P. (2014). “The Seismic Cycle in the Area of the 2011 Mw9.0 Tohoku-Oki Earthquake.” In: *Journal of Geophysical Research: Solid Earth* 119.5, pp. 4469–4515. DOI: 10.1002/2013JB010697.
- Perfettini, H., Frank, W. B., Marsan, D., and Bouchon, M. (2018). “A Model of Aftershock Migration Driven by Afterslip.” In: *Geophysical Research Letters* 45.5, pp. 2283–2293. DOI: 10.1002/2017GL076287.
- Periollat, A., Radiguet, M., Weiss, J., Twardzik, C., Amitrano, D., Cotte, N., Marill, L., and Socquet, A. (2022). “Transient Brittle Creep Mechanism Explains Early Postseismic Phase of the 2011 Tohoku-Oki Megathrust Earthquake: Observations by High-Rate GPS Solutions.” In: *Journal of Geophysical Research: Solid Earth* 127.8, e2022JB024005. DOI: 10.1029/2022JB024005.
- Philibosian, B. and Meltzer, A. J. (2020). “Segmentation and Supercycles: A Catalog of Earthquake Rupture Patterns from the Sumatran Sunda Megathrust and Other Well-Studied Faults Worldwide.” In: *Quaternary Science Reviews* 241, p. 106390. DOI: 10.1016/j.quascirev.2020.106390.
- Pollitz, F. F. (2003). “Transient Rheology of the Uppermost Mantle beneath the Mojave Desert, California.” In: *Earth and Planetary Science Letters* 215.1, pp. 89–104. DOI: 10.1016/S0012-821X(03)00432-1.

- Qiu, Q., Moore, J. D. P., Barbot, S., Feng, L., and Hill, E. M. (2018). “Transient Rheology of the Sumatran Mantle Wedge Revealed by a Decade of Great Earthquakes.” In: *Nature Communications* 9.1, p. 995. DOI: 10.1038/s41467-018-03298-6.
- Ragon, T. and Simons, M. (2023). “A Secondary Zone of Uplift Measured After Megathrust Earthquakes: Caused by Early Downdip Afterslip?” In: *Geophysical Research Letters* 50.7, e2022GL101510. DOI: 10.1029/2022GL101510.
- Rice, J. R. and Ruina, A. L. (1983). “Stability of Steady Frictional Slipping.” In: *Journal of Applied Mechanics* 50.2, pp. 343–349. DOI: 10.1115/1.3167042.
- Rice, J. R. (1993). “Spatio-Temporal Complexity of Slip on a Fault.” In: *Journal of Geophysical Research: Solid Earth* 98.B6, pp. 9885–9907. DOI: 10.1029/93JB00191.
- Riel, B., Simons, M., Agram, P., and Zhan, Z. (2014). “Detecting Transient Signals in Geodetic Time Series Using Sparse Estimation Techniques.” In: *Journal of Geophysical Research: Solid Earth* 119.6, pp. 5140–5160. DOI: 10.1002/2014JB011077.
- Ruina, A. (1983). “Slip Instability and State Variable Friction Laws.” In: *Journal of Geophysical Research: Solid Earth* 88.B12, pp. 10359–10370. DOI: 10.1029/JB088iB12p10359.
- Rutter, E. H. and Brodie, K. H. (2004). “Experimental Grain Size-Sensitive Flow of Hot-Pressed Brazilian Quartz Aggregates.” In: *Journal of Structural Geology* 26.11, pp. 2011–2023. DOI: 10.1016/j.jsg.2004.04.006.
- Santamaría-Gómez, A. (2019). “SARI: Interactive GNSS Position Time Series Analysis Software.” In: *GPS Solutions* 23.2, p. 52. DOI: 10.1007/s10291-019-0846-y.
- Santamaría-Gómez, A. and Ray, J. (2021). “Chameleonic Noise in GPS Position Time Series.” In: *Journal of Geophysical Research: Solid Earth* 126.3, e2020JB019541. DOI: 10.1029/2020JB019541.
- Satake, K., Fujii, Y., Harada, T., and Namegaya, Y. (2013). “Time and Space Distribution of Coseismic Slip of the 2011 Tohoku Earthquake as Inferred from Tsunami Waveform Data.” In: *Bulletin of the Seismological Society of America* 103.2B, pp. 1473–1492. DOI: 10.1785/0120120122.
- Saux, J. P., Molitors Bergman, E. G., Evans, E. L., and Loveless, J. P. (2022). “The Role of Slow Slip Events in the Cascadia Subduction Zone Earthquake Cycle.” In: *Journal of Geophysical Research: Solid Earth*. DOI: 10.1029/2021JB022425.
- Savage, J. C. (1983). “A Dislocation Model of Strain Accumulation and Release at a Subduction Zone.” In: *Journal of Geophysical Research: Solid Earth* 88.B6, pp. 4984–4996. DOI: 10.1029/JB088iB06p04984.

- Schoenberg, I. J. (1973). "The Basis Property of B-Splines." In: *Cardinal Spline Interpolation*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, pp. 11–19. DOI: 10.1137/1.9781611970555.ch2.
- Scholz, C. H. (2002). *The Mechanics of Earthquakes and Faulting*. 2nd ed. Cambridge University Press.
- Scholz, C. H. (1998). "Earthquakes and Friction Laws." In: *Nature* 391.6662, pp. 37–42. DOI: 10.1038/34097.
- Segall, P. (2010). *Earthquake and Volcano Deformation*. Princeton University Press.
- Segall, P. and Matthews, M. (1997). "Time Dependent Inversion of Geodetic Data." In: *Journal of Geophysical Research: Solid Earth* 102.B10, pp. 22391–22409. DOI: 10.1029/97JB01795.
- Shao, G., Li, X., Ji, C., and Maeda, T. (2011). "Focal Mechanism and Slip History of the 2011 Mw 9.1 off the Pacific Coast of Tohoku Earthquake, Constrained with Teleseismic Body and Surface Waves." In: *Earth, Planets and Space* 63.7, pp. 559–564. DOI: 10.5047/eps.2011.06.028.
- Shen, Z.-K., Wang, M., Zeng, Y., and Wang, F. (2015). "Optimal Interpolation of Spatially Discretized Geodetic Data." In: *Bulletin of the Seismological Society of America* 105.4, pp. 2117–2127. DOI: 10.1785/0120140247.
- Silverii, F., Montgomery-Brown, E. K., Borsa, A. A., and Barbour, A. J. (2020). "Hydrologically Induced Deformation in Long Valley Caldera and Adjacent Sierra Nevada." In: *Journal of Geophysical Research: Solid Earth* 125.5, e2020JB019495. DOI: 10.1029/2020JB019495.
- Simons, M., Minson, S. E., Sladen, A., Ortega, F., Jiang, J., Owen, S. E., Meng, L., Ampuero, J.-P., Wei, S., Chu, R., Helmberger, D. V., Kanamori, H., Hetland, E., Moore, A. W., and Webb, F. H. (2011). "The 2011 Magnitude 9.0 Tohoku-Oki Earthquake: Mosaicking the Megathrust from Seconds to Centuries." In: *Science* 332.6036, pp. 1421–1425. DOI: 10.1126/science.1206731.
- Stern, R. J. (2002). "Subduction Zones." In: *Reviews of Geophysics* 40.4, pp. 3-1-3–38. DOI: 10.1029/2001RG000108.
- Sugimura, A. and Uyeda, S. (1973). *Island Arcs: Japan and Its Environs*. Developments in Geotectonics 3. Amsterdam, New York: Elsevier Scientific Pub. Co.
- Sun, T., Wang, K., and He, J. (2024). "Geodetic Signature of a Weak Lithosphere-Asthenosphere Boundary in Postseismic Deformation of Large Subduction Earthquakes." In: *Earth and Planetary Science Letters* 630, p. 118619. DOI: 10.1016/j.epsl.2024.118619.
- Sun, T., Wang, K., Iinuma, T., Hino, R., He, J., Fujimoto, H., Kido, M., Osada, Y., Miura, S., Ohta, Y., and Hu, Y. (2014). "Prevalence of Viscoelastic Relaxation after the 2011 Tohoku-oki Earthquake." In: *Nature* 514.7520, pp. 84–87. DOI: 10.1038/nature13778.

- Suppe, J. (2014). “Fluid Overpressures and Strength of the Sedimentary Upper Crust.” In: *Journal of Structural Geology*. Fluids and Structures in Fold and Thrust Belts with Recognition of the Work of David V. Wiltschko 69, pp. 481–492. DOI: 10.1016/j.jsg.2014.07.009.
- Suwa, Y., Miura, S., Hasegawa, A., Sato, T., and Tachibana, K. (2006). “Interplate Coupling beneath NE Japan Inferred from Three-Dimensional Displacement Field.” In: *Journal of Geophysical Research: Solid Earth* 111.B4. DOI: 10.1029/2004JB003203.
- Tang, C.-H., Hsu, Y.-J., Barbot, S., Moore, J. D. P., and Chang, W.-L. (2019). “Lower-Crustal Rheology and Thermal Gradient in the Taiwan Orogenic Belt Illuminated by the 1999 Chi-Chi Earthquake.” In: *Science Advances* 5.2, eaav3287. DOI: 10.1126/sciadv.aav3287.
- Tanioka, Y., Hirata, K., Hino, R., and Kanazawa, T. (2004). “Slip Distribution of the 2003 Tokachi-oki Earthquake Estimated from Tsunami Waveform Inversion.” In: *Earth, Planets and Space* 56.3, pp. 373–376. DOI: 10.1186/BF03353067.
- The pandas development team (2021). *Pandas-Dev/Pandas: Pandas 1.3.0*. Zenodo. DOI: 10.5281/zenodo.5060318.
- Thomas, M. Y., Avouac, J.-P., and Lapusta, N. (2017). “Rate-and-State Friction Properties of the Longitudinal Valley Fault from Kinematic and Dynamic Modeling of Seismic and Aseismic Slip.” In: *Journal of Geophysical Research: Solid Earth* 122.4, pp. 3115–3137. DOI: 10.1002/2016JB013615.
- Thompson, B., Vasyura-Bathke, H., Howell, A., and Meade, B. (2023). *Tbenthompson/Cutde: V23.6.25*. Zenodo. DOI: 10.5281/zenodo.8080078.
- U. S. Geological Survey (2017). *Advanced National Seismic System (ANSS) Comprehensive Catalog*. DOI: 10.5066/F7MS3QZH.
- Uchida, N. and Bürgmann, R. (2021). “A Decade of Lessons Learned from the 2011 Tohoku-Oki Earthquake.” In: *Reviews of Geophysics* 59.2, e2020RG000713. DOI: 10.1029/2020RG000713.
- Uchida, N., Matsuzawa, T., Ellsworth, W. L., Imanishi, K., Shimamura, K., and Hasegawa, A. (2012). “Source Parameters of Microearthquakes on an Interplate Asperity off Kamaishi, NE Japan over Two Earthquake Cycles.” In: *Geophysical Journal International* 189.2, pp. 999–1014. DOI: 10.1111/j.1365-246X.2012.05377.x.
- Uchida, N., Shimamura, K., Matsuzawa, T., and Okada, T. (2015). “Postseismic Response of Repeating Earthquakes around the 2011 Tohoku-oki Earthquake: Moment Increases Due to the Fast Loading Rate.” In: *Journal of Geophysical Research: Solid Earth* 120.1, pp. 259–274. DOI: 10.1002/2013JB010933.

- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., and van Mulbregt, P. (2020). “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python.” In: *Nature Methods* 17.3, pp. 261–272. DOI: 10.1038/s41592-019-0686-2.
- Walwer, D., Calais, E., and Ghil, M. (2016). “Data-Adaptive Detection of Transient Deformation in Geodetic Networks.” In: *Journal of Geophysical Research: Solid Earth* 121.3, pp. 2129–2152. DOI: 10.1002/2015JB012424.
- Wang, K. and Dixon, T. (2004). ““Coupling” Semantics and Science in Earthquake Research.” In: *Eos, Transactions American Geophysical Union* 85.18, pp. 180–180. DOI: 10.1029/2004E0180005.
- Wegener, A. (1920). *Die Entstehung Der Kontinente Und Ozeane. 2.* gänzl. umgearb. Aufl. Die Wissenschaft Bd. 66. Braunschweig: Vieweg.
- Yagi, Y. (2004). “Source Rupture Process of the 2003 Tokachi-oki Earthquake Determined by Joint Inversion of Teleseismic Body Wave and Strong Ground Motion Data.” In: *Earth, Planets and Space* 56.3, pp. 311–316. DOI: 10.1186/BF03353057.
- Yagi, Y. and Fukahata, Y. (2011). “Rupture Process of the 2011 Tohoku-oki Earthquake and Absolute Elastic Strain Release: Rupture Process of Tohoku-Oki Earthquake.” In: *Geophysical Research Letters* 38.19, n/a–n/a. DOI: 10.1029/2011GL048701.
- Yamanaka, Y. and Kikuchi, M. (2003). “Source Process of the Recurrent Tokachi-oki Earthquake on September 26, 2003, Inferred from Teleseismic Body Waves.” In: *Earth, Planets and Space* 55.12, e21–e24. DOI: 10.1186/BF03352479.
- (2004). “Asperity Map along the Subduction Zone in Northeastern Japan Inferred from Regional Seismic Data.” In: *Journal of Geophysical Research: Solid Earth* 109.B7. DOI: 10.1029/2003JB002683.
- Yamashita, F., Fukuyama, E., Mizoguchi, K., Takizawa, S., Xu, S., and Kawakata, H. (2015). “Scale Dependence of Rock Friction at High Work Rate.” In: *Nature* 528.7581, pp. 254–257. DOI: 10.1038/nature16138.
- Yamazaki, Y., Lay, T., Cheung, K. F., Yue, H., and Kanamori, H. (2011). “Modeling Near-Field Tsunami Observations to Improve Finite-Fault Slip Models for the 11 March 2011 Tohoku Earthquake: The 2011 Tohoku Near-Field Tsunami Modeling.” In: *Geophysical Research Letters* 38.7, n/a–n/a. DOI: 10.1029/2011GL049130.
- Yokota, Y., Ishikawa, T., and Watanabe, S.-i. (2018). “Seafloor Crustal Deformation Data along the Subduction Zones around Japan Obtained by GNSS-A Observations.” In: *Scientific Data* 5.1, p. 180182. DOI: 10.1038/sdata.2018.182.

- Yue, H. and Lay, T. (2013). “Source Rupture Models for the Mw 9.0 2011 Tohoku Earthquake from Joint Inversions of High-Rate Geodetic and Seismic Data.” In: *Bulletin of the Seismological Society of America* 103.2B, pp. 1242–1255. DOI: 10.1785/0120120119.
- Yuen, D. A. and Peltier, W. R. (1982). “Normal Modes of the Viscoelastic Earth.” In: *Geophysical Journal International* 69.2, pp. 495–526. DOI: 10.1111/j.1365-246X.1982.tb04962.x.
- Yuzariyadi, M. and Heki, K. (2021). “Enhancement of Interplate Coupling in Adjacent Segments after Recent Megathrust Earthquakes.” In: *Tectonophysics* 801, p. 228719. DOI: 10.1016/j.tecto.2021.228719.

DISSTANS IMPLEMENTATION DETAILS

DISSTANS is written in Python. While the main text reports results obtained with DISSTANS, we focus here on presenting the structure and methodology of the package, with little to no actual sample code. For sample code, please refer to the package documentation. The following typesetting will be used for clarity: classes are capitalized and typeset in bold monospace font (e.g., **Station**) and attributes, properties, variables, methods, functions as well as general code are typeset in regular monospace font (e.g., parameters or `import disstans`) with callables (e.g., functions and methods) additionally being trailed with parentheses (e.g., `get_mapping()`).

A.1 Structure

Fig. A.1 presents the modular structure of DISSTANS. The highest level of abstraction is the **Network** class, which serves three main purposes. First, for each station in the network, it contains a **Station** object in its `stations` dictionary attribute, which enables straightforward access. Its second use is to provide a suite of convenience methods that perform a certain task for each station. Without parallelization enabled, their only advantage is that a user does not have to write explicit for-loops. However, **Network** methods also implement an automatic switch to parallelized execution using Python’s `multiprocessing.pool` module if the configuration is set accordingly. Finally, the **Network** class contains methods that interface with all stations simultaneously; for example, the graphical user interface `gui()` and other plotting functions (more details about visualization methods in Appendix A.8). Plotting functions are based on the standard Matplotlib (Hunter, 2007) and Cartopy (Elson et al., 2022) packages.

One level down in the hierarchy is the **Station** class. Apart from storing the metadata information `name` and `location`, it is the container object for all datasets being assigned to the station; for example, raw or post-processed GNSS displacement timeseries (e.g., *Dataset 1* and *Dataset 2* in Fig. A.1). A network can contain multiple stations, and each station can contain multiple datasets, but not all datasets have to be present at all stations. The **Station** class also provides functions that directly work on contained timeseries, such as `analyze_residuals()`.

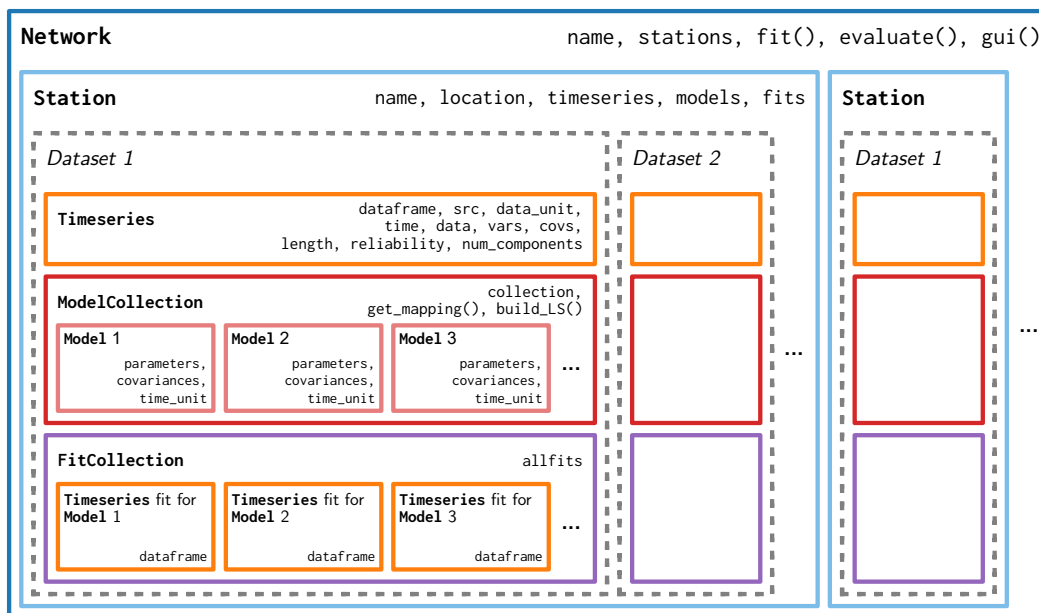


Figure A.1: Code structure of DISSTANS, explained in detail in Appendix A.1.

On the third level, for each dataset, a station contains three key elements: the actual data (in the **Timeseries** object, stored in the `Station.timeseries` dictionary), the associated models (as a **ModelCollection** object containing the individual **Model** objects, stored in the `Station.models` dictionary), and any fits to the data based on model evaluations (as a dictionary of **Timeseries** objects, one for each model, plus one for all models jointly, all stored in the `Station.fits` **FitCollection** object). Using the methods provided by the **Station** class ensures that whenever a new dataset is added (or removed), all three elements are initialized (or deleted) appropriately. While this separation might appear somewhat confusing, it is necessary to enable easy access to individual objects while preserving flexibility. For example, a **Timeseries** object is physically independent of whatever model one wants to apply to it, and therefore the code should reflect this (i.e., the **Timeseries** object should not change when a model is added or removed, or when an individual model is evaluated to yield a prediction). The separation into data, models, and fits also allows for the same dataset to easily have different models at different stations, or multiple models of the same class (e.g., two sets of step functions, one for maintenance and one for earthquake-induced steps). Using the **Timeseries** class also for fits (i.e., model-predicted timeseries) allows for the efficient re-use of practical **Timeseries** methods such as file storage or mathematical operations.

On the lowest level, the **Model** and **Timeseries** objects store their data using standard

NumPy arrays (Harris et al., 2020) and pandas DataFrames (McKinney, 2010; The pandas development team, 2021), respectively, enabling seamless integration with existing Python-based workflows.

The open-source nature of the code, along with a defined hierarchical, object-oriented structure, allows for easy modification and extension by the user through subclassing. For example, storing additional station metadata such as antenna information can easily be implemented by creating a Python class inheriting from the **Station** class and extending the initialization function to accept additional instance variables. Another example is the implementation of new user-defined models by subclassing **Model** which then seamlessly integrate into the rest of DISSTANS's workflow. Finally, loading timeseries data from a custom data format can be integrated into DISSTANS by subclassing the **Timeseries** class. In fact, all of the included models (see below) and timeseries file formats are subclasses of **Model** and **Timeseries**, respectively, and can be used as examples by users wishing to extend the code functionality.

A.2 Models

DISSTANS uses a linear combination of parametric models. Parametric models linear in their coefficients (i.e., not necessarily composed of linear functions) allow both simple unregularized as well as more complex L2, L1 or L0 regularized least-squares fitting (more detail about regularization schemes in Appendix A.3). Furthermore, estimating multiple models jointly is straightforward as they are linearly added together, and the mapping (or design) matrix is simply a horizontal stack of all the models' individual mapping matrices (everything automatically done by the **ModelCollection** class). Lastly, the formal estimated model parameter covariance matrix can usually be estimated in closed-form.

The individual **Model** classes included in DISSTANS can be separated into basic and spline models. All models can be used with one or multiple data components. The basic models currently included are: **Polynomial**, **Step**, **Sinusoid**, **Logarithmic**, **Exponential**, **HyperbolicTangent** and **Arctangent**. The basic models are either single functions (e.g., logarithm), or their functions form orthogonal bases within their class (e.g., polynomials). The spline modeling in **BSpline** or **ISpline** model is based on Hetland et al. (2012) and Riel et al. (2014), containing multiple cardinal B- or integrated-B-splines (respectively) of the same timescale and order but with different center times. The **SplineSet** combines several **BSpline** or **ISpline** models of

different timescales into one large collection, forming a linearly-dependent (overcomplete) spanning set able to approximate arbitrary functions. The `AmpPhModulatedSinusoid` estimates a sinusoid of a given nominal frequency, but allows the instantaneous amplitude and phase to vary. Time-varying properties are enabled by modeling the linear sine and cosine coefficients of the sinusoid as being defined by a linearly-independent set of B-Spline basis functions over the given time interval. Some form of regularization is necessary to gain a meaningful result when using an overcomplete set of splines.

A.2.1 Joint Mathematical Formulation

In DISSTANS, the joint mathematical formulation $g(t)$ is the sum of all the individual models contained in a `ModelCollection`. Each individual constituent g_{Model} (described by `Model` objects) can again be a linear superposition of functions g_j and corresponding coefficients m_j :

$$g(t) = \sum_{\text{num_parameters}} g_{\text{Model}}(t) \quad (\text{A.1a})$$

$$= \sum_{j=1} m_j g_j(t). \quad (\text{A.1b})$$

Here, `num_parameters` is the total number of all individual functions, and therefore also the number of all coefficients to be estimated.

While the models are continuous in time, timeseries decomposition inherently works on discrete observations d_i at times t_i . Using matrix notation, the least-squares problem can be formulated as follows:

$$\mathbf{d} = \mathbf{G}\mathbf{m} + \epsilon \quad (\text{A.2})$$

where

$$\mathbf{d} = (d_i) \in \mathbb{R}^{\text{num_observations} \times 1} \quad (\text{A.3a})$$

$$\mathbf{G} = (G_{i,j}) \in \mathbb{R}^{\text{num_observations} \times \text{num_parameters}} \quad (\text{A.3b})$$

$$= (g_j(t_i)) \in \mathbb{R}^{\text{num_observations} \times \text{num_parameters}} \quad (\text{A.3c})$$

$$\mathbf{m} = (m_j) \in \mathbb{R}^{\text{num_parameters} \times 1} \quad (\text{A.3d})$$

and $\epsilon \in \mathbb{R}^{\text{num_observations} \times 1}$ is the column vector of residuals. All solvers start from this formulation to find the best set of m_j that minimizes a given cost function dependent on ϵ (potentially including regularization criteria, see Appendix A.3). The choice of the data misfit loss function implicitly defines the assumed distribution from which ϵ

is drawn (e.g., a Normal distribution in the case of unregularized least-squares). The observations can include measurements in all three dimensions, allowing the use of cross-component covariances in the fitting process. In DISSTANS, the mapping (i.e., design) matrices \mathbf{G} are assembled by the `get_mapping()` methods, \mathbf{d} is represented by `Timeseries` objects, and \mathbf{m} is returned by the solver in `Solution` objects and added to each `Model` object.

In the following three subsections, we detail both the basic and the spline-based models.

A.2.2 Basic Models

The basic models in DISSTANS include function commonly used to model geodetic timeseries:

$$g_{\text{Arctangent}}(t) = m' \left(\frac{1}{\pi} \arctan \left(\frac{t}{\tau} \right) + 0.5 \right) \quad (\text{A.4a})$$

$$g_{\text{Exponential}}(t) = m' \left(1 - \exp \left(-\frac{t}{\tau} \right) \right) \quad (\text{A.4b})$$

$$g_{\text{HyperbolicTangent}}(t) = m' \left(\frac{1}{2} \tanh \left(\frac{t}{\tau} \right) + 0.5 \right) \quad (\text{A.4c})$$

$$g_{\text{Logarithmic}}(t) = m' \log \left(1 + \frac{t}{\tau} \right) \quad (\text{A.4d})$$

$$g_{\text{Polynomial}}(t) = \sum_l m'_l t^l \quad (\text{A.4e})$$

$$g_{\text{Step}}(t) = \sum_l m'_l H(t - t_l^{\text{step}}) \quad (\text{A.4f})$$

$$g_{\text{Sinusoid}}(t) = m'_0 \cos(\omega t) + m'_1 \sin(\omega t) \quad (\text{A.4g})$$

where all m' are just stand-ins for the overall set of m_j , τ can vary between models, the t_l^{step} are step times, and $H(t)$ is the Heaviside function. While all of these models are available out-of-the-box, the user still has to actively specify which models to use, how many of each, and with which reference times, timescales, or periods.

A.2.3 Linearly Dependent, Overcomplete Dictionary of Splines: `splineSet`

For study areas where significant transients of arbitrary shape can be found, DISSTANS offers spline-based transient modeling.

We start with the formulation of a single cardinal B-spline basis function (*spline function*) of reference time t_{ref} . “Normalized” timestamps t' can be calculated as follows:

$$t' = \frac{t - t_{\text{ref}}}{\rho}. \quad (\text{A.5})$$




Figure A.2: Example of **BSpline** (left) and **ISpline** (right) spline functions for a single timescale and order at seven different center times. Each curve on the right is the integral of the curve on the left of the same color.

By default, this single spline function is then shifted to multiple center times by using its timescale ρ , leading to different normalized timevectors for each spline function:

$$t'_j = \frac{(t - t_{\text{ref}}) - j \cdot \rho}{\rho}. \quad (\text{A.6})$$

(Here, $j = 0 \dots \text{num_splines}$ only considers the spline functions.)

To create the spline functions of a certain degree p (with order $n = p + 1$), we can then use the following relation (Schoenberg, 1973; Butzer et al., 1988):

$$g_j(t'_j) = \sum_{k=0}^n \frac{(-1)^k}{p!} \cdot \binom{n}{k} \cdot \left(t'_j + \frac{n}{2} - k\right)^p. \quad (\text{A.7})$$

This is the model represented by **BSpline**. Based on Hetland et al. (2012), this study uses the integrated form of this spline function to represent transients. Its mathematical representation is:

$$g_j(t'_j) = \sum_{k=0}^n \frac{(-1)^k}{(p+1)!} \cdot \binom{n}{k} \cdot \left(t'_j + \frac{n}{2} - k\right)^{p+1}. \quad (\text{A.8})$$

The final spline model (a single **BSpline** or **ISpline** object) over all the available center times is therefore

$$g_{\{\text{B,I}\}\text{Spline}}(t) = \sum_{j=0}^{\text{num_splines}} m_j g_j(t). \quad (\text{A.9})$$

Within the **SplineSet** class, this model can then be repeated again for different timescales ρ . Fig. A.2 shows example spline functions.




Figure A.3: Visualization of the intermediate functions used by `AmpPhModulatedSinusoid`. The spline basis functions h_j over the considered time interval are in the left panel. Multiplying the spline functions with the cosine and sine of a given period then yields the modulated cosines and sines in the center and right panel, respectively.

A.2.4 Linearly Independent Spline Basis for Time-varying Sinusoids:

`AmpPhModulatedSinusoid`

For study areas where the amplitude of seasonal oscillations varies significantly, DISSTANS offers varying-amplitude sinusoidal modeling. Such modeling can also be used to improve the fitting of shortterm transient processes, as potentially periodic parts can then be accommodated by the seasonal model.

The simple `Sinusoid` class models a seasonal signal, given a certain frequency ω , as the linear combination of a sine and cosine combination, allowing one to estimate both phase ϕ and amplitude A as a linear problem:

$$g_{\text{Sinusoid}}(t) = A \cos(\omega t - \phi) = a \cos(\omega t) + b \sin(\omega t). \quad (\text{A.10})$$

If we want to allow the overall amplitude A to change over time, we can extend the definition of a (and similarly, b):

$$a \rightarrow a(t) = \bar{a} + \Delta a(t) \quad (\text{A.11a})$$

$$b \rightarrow b(t) = \bar{b} + \Delta b(t). \quad (\text{A.11b})$$

To keep the problem linear, we can use a spline representation for $\Delta a(t)$, $\Delta b(t)$:

$$a(t) = \bar{a} + \sum_j a_j h_j(t) \quad (\text{A.12a})$$

$$b(t) = \bar{b} + \sum_j b_j h_j(t) \quad (\text{A.12b})$$

where the a_j (and respectively, b_j) are the parameters m_j to estimate, and h_j are the spline basis functions (more on h_j below). Expanding $g_{\text{Sinusoid}}(t)$ with the extended

definition leads to a natural separation of components:

$$\begin{aligned}
& a(t) \cos(\omega t) + b(t) \sin(\omega t) \\
&= (\bar{a} + \Delta a(t)) \cos(\omega t) \\
&\quad + (\bar{b} + \Delta b(t)) \sin(\omega t) \\
&= (\bar{a} \cos(\omega t) + \bar{b} \sin(\omega t)) \\
&\quad + (\Delta a(t) \cos(\omega t) + \Delta b(t) \sin(\omega t)) \\
&= (\bar{a} \cos(\omega t) + \bar{b} \sin(\omega t)) \\
&\quad + \sum_j (a_j h_j(t) \cos(\omega t) + b_j h_j(t) \sin(\omega t)) \\
&= g_{\text{Sinusoid}}(t) + g_{\text{AmpPhModulatedSinusoid}}(t).
\end{aligned} \tag{A.13}$$

Here, the first term represents the *nominal* component, and the second term the *deviation* component. In DISSTANS, the terms correspond to the `Sinusoid` and `AmpPhModulatedSinusoid`, respectively.

Note that the h_j are not the same as for the dictionary of splines defined above. The dictionary is comprised of a single (cardinal) spline, that is of a defined length scale (i.e., period), and centered at specified timestamps. Here, for `AmpPhModulatedSinusoid`, we do not need the spline to be the same one shifted and scaled, instead we can default to the more general notion of B-Splines — a complete basis for polynomials of a given degree on a given interval. This relaxation allows us to use SciPy’s basis function implementation directly (Virtanen et al., 2020). Fig. A.3 shows an example set of spline basis functions h_j , as well as the resulting modulated cosine and sine terms used as the spanning functions for `AmpPhModulatedSinusoid`.

Although it is not strictly necessary to include \bar{a} and \bar{b} explicitly in $a(t)$ and $b(t)$ (splines can also represent any constant function), the separation is useful because it allows a regularized solver to not penalize the nominal component.

A.3 Solver Functions

The provided solver functions are least-squares (therefore parametric) solvers, with varying degrees of added complexity. They each

1. Build the mapping and observation matrices for a given `Timeseries` object of observations and `ModelCollection` object (\mathbf{G} and \mathbf{d} , respectively, see Appendix A.2),

2. Divide the solution process into independent sub-problems if there is no data component covariance (decreasing the computational burden),
3. Call a lower-level solver to minimize the cost function $\|\mathbf{G}\mathbf{m} - \mathbf{d}\|_2^2$ (potentially subject to regularization),
4. Optionally calculate the formal model parameter covariance matrix \mathbf{C}_m , and
5. Return a `Solution` object (containing the best-fitting \mathbf{m}).

To prevent convergence or numerical issues, the solvers and the `Solution` class keep track of model parameters that cannot be estimated (because they are not observable given the timespan of the observations) or should not be estimated (useful, for example, if some splines in a `SplineSet` are assumed to be zero). The regularized solvers additionally keep track of which model's parameters should be regularized, allowing for a flexible regularization approach.

The first, most basic solver is `linear_regression()`, which provides the above-mentioned features as a wrapper to the least-squares routine in SciPy (Virtanen et al., 2020). It can therefore be regarded as a minimal code example for new, user-defined solvers. The cost function to be minimized is:

$$\|\mathbf{G}\mathbf{m} - \mathbf{d}\|_2^2 \tag{A.14}$$

and the posterior covariance matrix \mathbf{C}_m given the data covariance matrix \mathbf{C}_d is

$$\mathbf{C}_m = (\mathbf{G}^T \mathbf{C}_d^{-1} \mathbf{G})^{-1} \tag{A.15}$$

where $^{-1}$ is the generalized pseudo-inverse for matrices.

The second provided solver, `ridge_regression()`, adds L2 regularization, and also relies on the least-squares routine in SciPy. It minimizes the cost function

$$\|\mathbf{G}\mathbf{m} - \mathbf{d}\|_2^2 + \lambda \|\mathbf{m}_{\text{reg}}\|_2^2 \tag{A.16}$$

where λ is a chosen regularization penalty hyperparameter, and \mathbf{m}_{reg} is the subset of \mathbf{m} that should be regularized. Furthermore, λ can vary between data components to account for different noise levels. The posterior covariance matrix takes the regularization into account:

$$\mathbf{C}_m = (\mathbf{G}^T \mathbf{C}_d^{-1} \mathbf{G} + \lambda \mathbf{I}_{\text{reg}})^{-1}. \tag{A.17}$$

The third solver, `lasso_regression()`, uses CVXPY (Diamond and Boyd, 2016; Agrawal et al., 2018) to provide L1 and, by means of weighted iterations, station-specific L0 regularization (Candès et al., 2008). In its basic form, the solver minimizes

$$\|\mathbf{G}\mathbf{m} - \mathbf{d}\|_2^2 + \lambda \|\mathbf{m}_{\text{reg}}\|_1. \quad (\text{A.18})$$

By defining a reweighting function and iterating on the L1-regularized solution, the `lasso_regression()` solver approximates the solution for the L0-regularized³ least-squares problem, minimizing

$$\|\mathbf{G}\mathbf{m} - \mathbf{d}\|_2^2 + \|\mathbf{m}_{\text{reg}}\|_0. \quad (\text{A.19})$$

Because the result of an L0-regularized solution is approximately the same as if an unregularized problem was solved with only a subset of model parameters to be estimated, the posterior covariance matrix for `lasso_regression()` is the same as for `linear_regression()`, but setting to zero the covariances which were not estimated.

Note that the reweighting function does not explicitly appear in the L0 cost function, although it does mimic the role of the regularization penalty λ from the L2 and L1 cost functions. Specifically, during the iteration process, the reweighting function returns penalties that anticorrelate with parameter amplitude: small parameters will be penalized heavily, and large parameters will receive very little penalty. With this approach, the solver converges to the L0-regularized solution where parameters either have near-zero penalty (if the parameters are deemed significant by the solver), or near-zero amplitude (because of their high penalty). Consequently, the final cost function does not contain an explicit penalty hyperparameter λ , although care needs to be taken when specifying the reweighting function such that it is able to distinguish insignificant from significant parameters. Thresholds for this distinction are not hard cut-offs; they are defined within the context of `ReweightingFunction` objects, and usually correspond to the location of an L-shape bend in the reweighting function's shape. While the appropriate choice of functions and scales will vary between applications, a good (empirical) starting point are functions whose penalties close to an input value of zero are of a similar order of magnitude of the data being fitted. For more details about the implementation of the L0-regularized solver, including examples of reweighting functions, see Candès et al., 2008.

The `Network.spatialfit()` method extends the possibilities of station-specific L0 regularization to also take into account the weights of a given model at nearby stations. The approach implemented here follows Riel et al. (2014) closely, with the

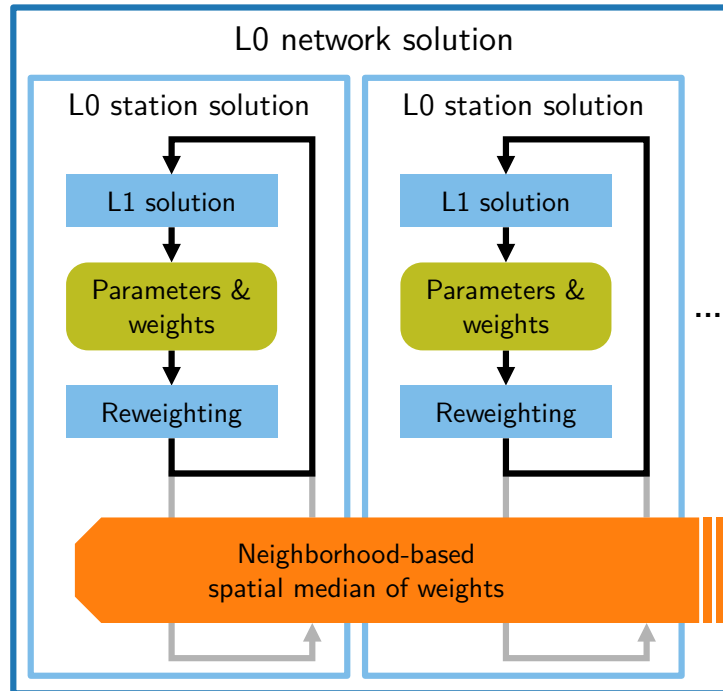


Figure A.4: Flowchart of the spatiotemporal L0-regularized solver as described in Riel et al. (2014). Symbols and colors from Fig. A.5. At each station, an L1-regularized least-squares fit is computed, where each parameter has an associated weight. The weight is inversely correlated to the parameter magnitude. Parameters close to zero are iteratively penalized, whereas significant parameters have their penalty gradually reduced to zero. Iterated L1 regularization effectively approximates an L0-regularized solution (see Candès et al., 2008). By combining the weights between stations with a median in an intermediate step, parameters that are significant at other nearby stations as well are promoted, and parameters that are insignificant are demoted.

goal to identify signals close to the noise floor, suppress local noise, and promote sparse models in both time *and* space. A visual summary of the method is given in Fig. A.4. DISSTANS is able to perform the station-specific fits in parallel, resulting in a large runtime improvement. Lastly, `Network.spatialfit()` can also minimize the jointly L1- and L0-regularized problem:

$$\|\mathbf{G}\mathbf{m} - \mathbf{d}\|_2^2 + \lambda \|\mathbf{m}_{\text{reg,L1}}\|_1 + \|\mathbf{m}_{\text{reg,L0}}\|_0. \quad (\text{A.20})$$

A.4 Data Formats

All timeseries datasets are stored as objects of `Timeseries` subclasses. The `Timeseries` parent class defines an internal data structure that all further processing done by DISSTANS methods of all levels rely on. DISSTANS also implements properties such as the calculation of a timeseries length or reliability, the possibility to use Python’s in-built mathematical operators to create new timeseries, and conve-

nience functions such as cutting the timeseries or building covariance matrices at a particular timestep.

Subclasses, in turn, define how any particular input file gets loaded to match the common structure. The two provided subclasses are `GipsyTimeseries` (for JPL's GipsyX .tseries files) and `UNRTimeseries` (for UNR's .tenv3 files). User-defined classes can easily be created by adhering to the format of the two existing subclasses, and checking the documentation of `Timeseries`.

A.5 Synthetic Data

The creation of synthetic data is another feature directly integrated into DISSTANS. Each `Model` and `ModelCollection` object has the two methods `read_parameters()` and `evaluate()`, which integrate into existing Python workflow by accepting and returning (respectively) NumPy arrays. A typical workflow to generate datasets therefore is to instantiate `Model` objects (e.g., a polynomial of a certain order), define and read in the parameters of the model, and finally evaluate the individual models (or a `ModelCollection` containing the individual ones). If the data is then to be used within DISSTANS, a simple `Timeseries` constructor exists for NumPy arrays, otherwise one can use the regular NumPy methods for exporting the data.

A.6 Step Detector

DISSTANS includes the `StepDetector` class to perform statistics-based assessments on whether step models should be added to a timeseries to estimate offsets due to physical (e.g., earthquakes) or non-physical (e.g., maintenance events) processes. Since there is no fully-automated algorithm that approaches the performance of manual inspection (e.g., Gazeaux et al., 2013), the focus here lies on providing a semi-automated method that is to be used in conjunction with end-user interaction. The method implemented in DISSTANS is based on the Akaike Information Criterion (AIC, e.g., Burnham and Anderson, 2002). For a given window size, the algorithm will evaluate the residuals of fitting the timeseries with two different models: one with a simple linear slope and offset, and the other with a linear slope, offset, and an additional offset in the center of the window. Then, for each timestep, the relative probability of the step versus the no-step model being true is calculated. In the last step, the maxima of the step probabilities are calculated and thresholded. The user can then examine the steps (alongside their respective variance reductions) and determine whether to add the offsets as steps to be estimated.

Our step detector approach is similar to the one in GipsyX (Bertiger et al., 2020),

but where GipsyX considers multiple window sizes, DISSTANS only uses a single one. Of course, **StepDetector** can be run multiple times with different window sizes, such that their combined results can provide a more robust step probability estimate.

A.7 Example Workflow

Even though DISSTANS is modular and therefore highly flexible, we propose the workflow presented in Fig. A.5 as a general starting point for timeseries decomposition with DISSTANS.

The first step is the acquisition and preparation of the raw input datasets, i.e., the GNSS network station displacement timeseries (and, if available, associated maintenance and seismic catalogs). Applying quality metrics such as requiring a minimum number of observations or station reliability (through their respective **Timeseries** attributes `num_observations` and `reliability`) ensures that the fitting process is not hindered by bad data.

We view the second step as a “preprocessing” one, where we identify and remove statistical outliers and the common mode errors (CME, see Dong et al., 2006; Huang et al., 2012) from the observations (see Fig. A.6 for more detail). The relevant functions are `median()`, `clean()` and `common_mode()`, which are called on the entire network (respecting parallelization) through the **Network** methods `call_func_ts_return()` and `call_func_no_return()`. The CME is systematic for the entire network, reflects noise in the estimation of the reference frame and manifests itself as a high-frequency noise realization that should be estimated independently of model fits (which could create additional systematic errors). To estimate the CME (e.g., using PCA/ICA), we first remove empirically the potentially interesting, low-frequency signal using a low-pass running median. A median filter is robust when handling large steps in the data (which may be present before any step removal is performed). Outlier removal is performed on the residual between low-pass and input signal, based on the residual’s variance.

Offsets (or steps) in the data are the big obstacles for model fitting. Left unaccounted for, they will influence every other model component (e.g., the secular plate velocity). While big jumps in the data can easily be spotted by comparing a measurement with the variance around the mean of previous observations, smaller offsets that are either below or similar to the data variance, and/or are accompanied by transient motion, are more challenging to detect. Ideally, all occurring offsets are known in advance based on ancillary catalogs, and could be categorized into equipment

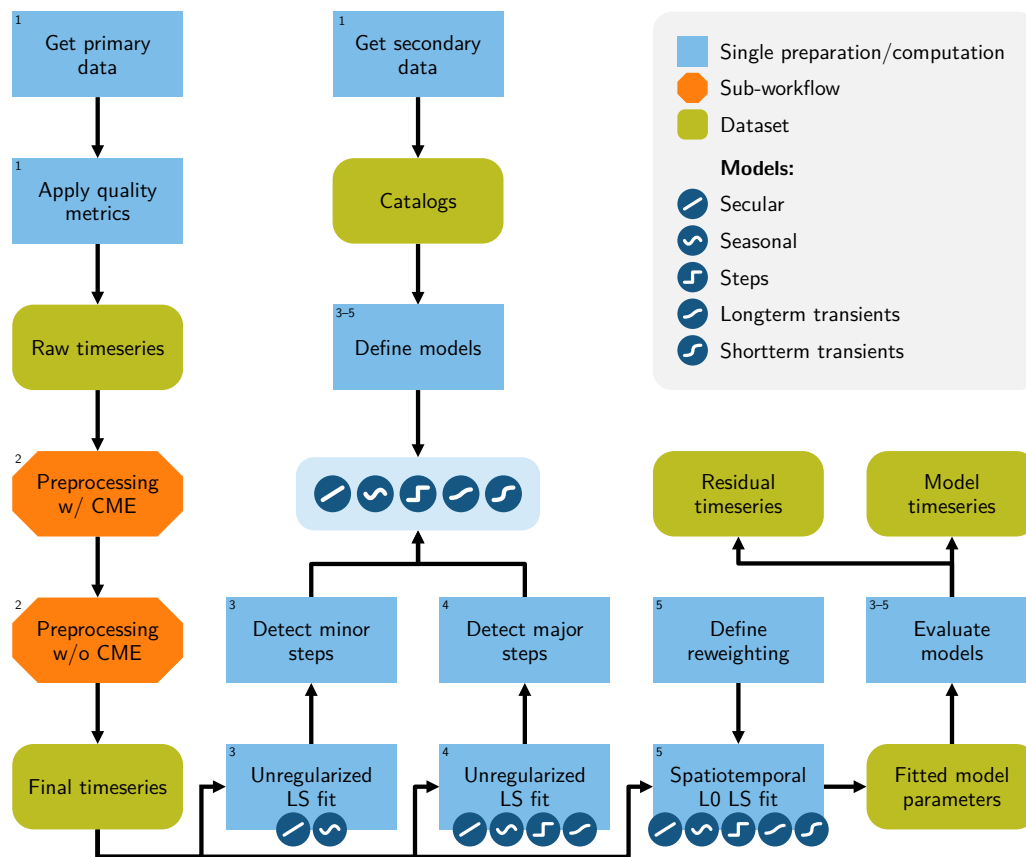


Figure A.5: Example workflow for using DISSTANS, explained in detail in Appendix A.7. Blue rectangles represent single computational steps, orange rectangles with cut corners represent sub-workflows discussed in more detail elsewhere, and green, rounded rectangles represent datasets at their different stages of processing. The numbered steps in the text correspond to the numbering in the top left corners of the rectangles.

changes and physical processes.

Maintenance events (e.g., antenna replacements, software changes, receiver upgrades) usually are well-recorded and accessible. Functions like `parse_maintenance_table()` and `parse_unr_steps()` are useful for these purposes. However, not all maintenance events automatically have a visible effect in the data, and therefore there are “grey zones” where the addition of a modeled step may be more harmful than beneficial. In these cases, we can perform an iterative process between fitting larger signals, and then checking again for evidence of smaller offsets.

A similar case can be made for the presence of coseismic displacements. Large, nearby earthquakes produce offsets that can be predicted from seismic catalogs and simple forward modeling of the expected displacement at any given station. (The earthquakes module provides this functionality in DISSTANS.) However, smaller

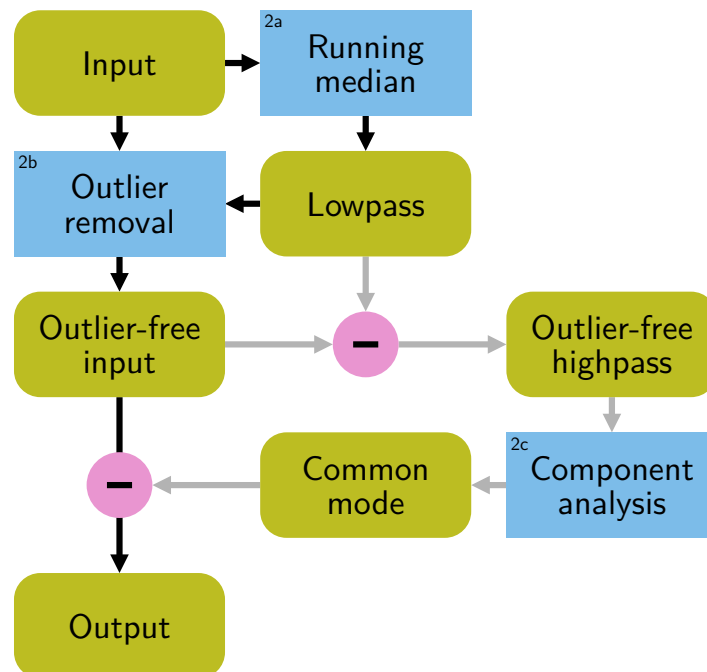


Figure A.6: Preprocessing sub-workflow, following the same symbolic and coloring as Fig. A.5 (step 2), with rose circles representing mathematical operations. First, a running median of the input is calculated, which results in a lowpass filtered timeseries. The variance of the input around the lowpass timeseries is used to detect outliers. Removing them from the input yields the outlier-free input. Without common mode estimation, this is also the final output. To remove the common mode, the difference between the lowpassed input and the outlier-free input is calculated, which yields an outlier-free, highpassed input. The dominant component of this timeseries is the best estimate of the common mode error. Removing this from the outlier-free input yields the outlier-free, common-mode-removed output.

events might not necessarily warrant an additional modeled step, and very fast transients would be better fit by transient models. Therefore, we recommend an iterative approach here as well.

The next steps in the proposed workflow are iterations of step-detection and model-fitting. In the third step, an unregularized least-squares fit with only a polynomial and some sinusoidal models is performed at each station individually. Using the `StepDetector` class, extremely prominent offsets in the data are well resolved, and are added to a list of offsets to be fit (with the `Step` model class).

In the fourth step, using the initial simple models, the defined list of offsets, and a `Sp1ineSet` dictionary of longterm transient splines, another unregularized least-squares fit is computed. Together with external maintenance and seismic catalogs, a second run of the `StepDetector` then aims to identify smaller steps that are to be estimated.

For the fifth step of the proposed workflow, the `Network.spatialfit()` method is used to perform a network-wide fit using the aforementioned polynomial, sinusoidal, and step models, as well as an expanded spline dictionary that includes also shorter-term transients. Only the spline parameters are subject to the spatiotemporal L0 regularization, although the regularization can be extended to all models. Defining an appropriate **ReweightingFunction** ensures a sparse, yet well-fitting solution. When seasonal effects are found to be strongly varying in time, allowing the seasonal signal to vary in amplitude over time (using **AmpPhModulatedSinusoid** models), can also improve the fit.

The results at each step are a set of model parameters for each data component, together with a complete parameter covariance matrix. They can be evaluated at all stages to yield the overall model-predicted timeseries (including its predicted uncertainty), as well as the individual constituent contributions. The residuals can be computed using the `Network.math()` methods and analyzed using the `Network.gui()` method to assure no systematic misfit is present. Of course, there are many variations to this example workflow.

A.8 Visualization

Because the raw data contained within **Timeseries** objects are standard pandas DataFrames, they can be plotted using standard Matplotlib code using their `Timeseries.time` and `Timeseries.data` attributes. Utilizing commonly-used Python object formats enables easy inspections of a particular station, timeseries, or fit; and allows for non-standard user-desired plotting. Model parameter values and covariances (accessed through their `Model.parameters` and `Model.covariances` NumPy array attributes) are also directly plottable with Matplotlib.

There are high-level visualization routines already included in DISSTANS. The core functionality is contained within the `Network.gui()` method, which provides a clickable map of the network (optionally with satellite imagery background), and a separate figure with all the timeseries contained by a station. If a timeseries contains fitted models, the overall model prediction is plotted, and optionally, can be split up into the different model components, and if there are **SplineSet** models present at a station, a scalogram can be shown. All figures can also be saved directly to files.

Furthermore, to visualize station motion in a map view, the `Network.wormplot()` method can produce still maps and animated videos of station displacements (or individual model constituents of them). Lastly, the `Network.graphical_cme()` method

performs common mode estimation (see Appendix A.7) and presents the temporal and spatial components separately for validation purposes.

Appendix B

DISSTANS SUPPLEMENTAL MATERIAL

B.1 Synthetic Network




Figure B.1: Convergence of the iterative, spatiotemporal L0-regularized solver (line) for 8416 total spline coefficients. For reference, values for the local L1-regularized (circles) and local L0-regularized (diamonds) solutions are also plotted on the axes. The results are shown both for the individual components (blue and orange, right vertical axis), as well as the overall solution (black, left vertical axis). The total number of non-zero parameters (i.e., the sum over all stations and components of the number of non-zero coefficients) as well as the number of *unique* non-zero parameters (i.e., the number of all splines that are non-zero at least at one station, per component) converges monotonically onto their final values. The latter number specifically demonstrates the effect of spatial sparsity.

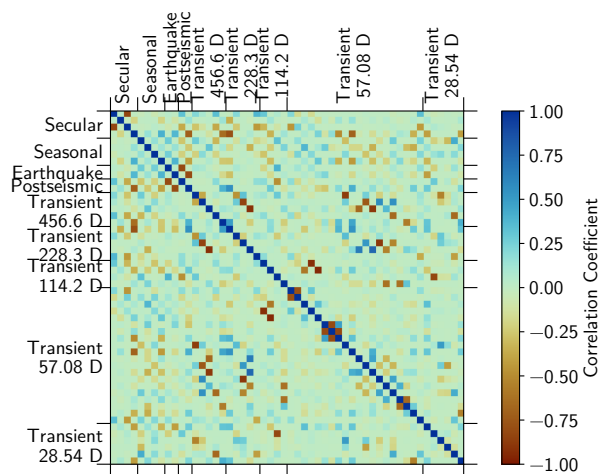


Figure B.2: Model parameter correlation matrix for all models and both east-north components. The covariances for spline parameters that are estimated to be close to zero are set to zero as well, and not shown. Tradeoffs between models and within splines are clearly identifiable.

	East [mm/a]	North [mm/a]
Spatial L0	0.108107	0.227728
Local L0	0.456424	0.512325
Linear + Steps	0.623014	0.605845
Linear	1.275072	1.272709
MIDAS	0.854414	0.871200

Table B.1: Root-Mean-Squared-Error between the secular velocity estimates and the true secular velocity, averaged across the entire network, for the methods presented in Fig. B.5, and both data components individually. The spatial L0 solution significantly outperforms the other solutions (including the local L0 solution).

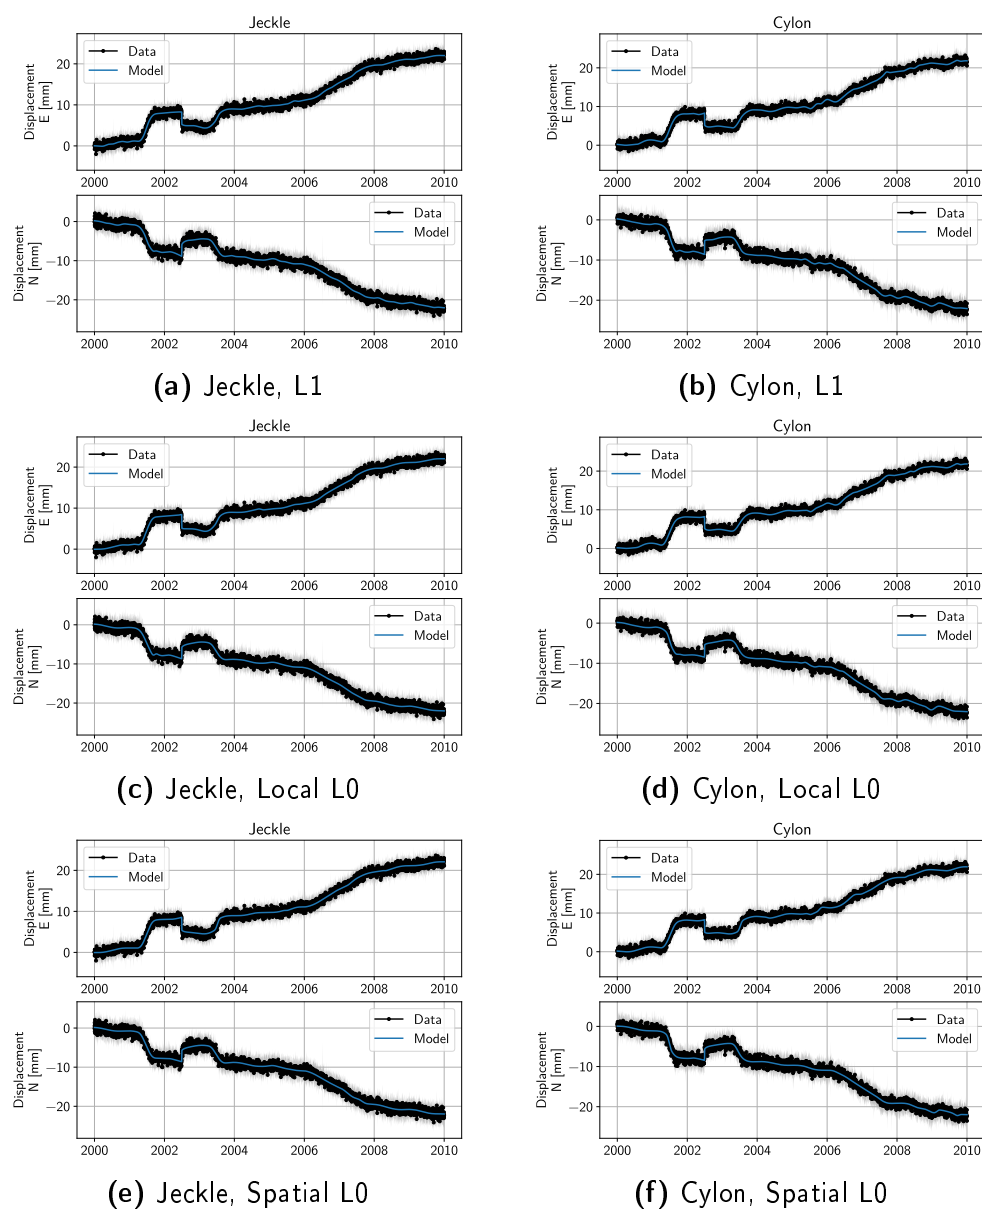





Figure B.3: Timeseries comparison for the two neighboring stations Jeckle and Cylon and three regularization schemes in the two east and north components. Black dots are the synthetic observations, with the grey shading corresponding to three standard deviations of simulated observation uncertainty. The blue line represents the final overall model fit. The fitted models are virtually indistinguishable between regularization schemes, except for the overfitting of some colored noise at station Cylon, which is reduced with the spatial L0 regularization.




(a) Jeckle, L1




(b) Cylon, L1




(c) Jeckle, Local L0



(d) Cylon, Local L0



(e) Jeckle, Spatial L0



(f) Cylon, Spatial L0

Figure B.4: See next page.

Figure B.4: Scalograms of the transient model for the stations and regularization in Fig. B.3. The horizontal and vertical axes correspond to time and the discrete periods of the splines, respectively. Patches (colored by the spline coefficient's value) in this time-period-space represent a single spline in the dictionary, with their extent in time defined as the active period of the spline (i.e., having non-zero gradient), and their height defined by the relative magnitude of the particular spline compared to all splines active at that time. Using the L1 solver, the transients (two shortterm, one longterm) are sparsely fitted in time, but not in space (i.e., each station's timeseries is fit using different splines). The local L0 regularization does not change this general behavior. Spatial L0 regularization leads to the transients being sparsely fitted in time *and* space (i.e., every station's timeseries is fit with a similar set of splines). Modeling the transients with coefficients sparse in time, space and period is beneficial in the context of identifying signals close to the noise floor that are appearing at multiple stations, since the respective coefficients will be penalized less, allowing for a more physically-consistent decomposition. Conversely, the penalization of coefficients that are only seen at isolated stations makes it easier to identify local shortterm noise processes.




Figure B.5: Comparison of the secular velocity estimates (East component) presented in Section 2.4.1, Fig. 2.2 with other methods. The dark gray dots correspond to the sum of true secular velocity and the generated noise; the light gray dots additionally include the transient constituents. The true secular velocity (orange line) and secular velocity estimated by the spatial L0 solver (blue line) correspond to the same lines in Fig. 2.2. The secular velocity estimated by the local L0 solver (compare Fig. 2.3) is shown with the green line (essentially overlapping the spatial L0 result). The three other lines correspond to other commonly used methods: local, unregularized least squares with (red) and without (purple) steps at the transient center times (assumed to be known); and the MIDAS solution (brown). In our synthetic example, the presence of the transient signal throughout most of the considered timeseries is significant enough to heavily deteriorate the estimated produced by the simple least squares as well as the MIDAS methods. For this station, and in the east component, adding spatial awareness only provides a minimal benefit, although Table B.1 shows that across components and stations, there are significant benefits.

B.2 Influence of Number of Stations

The code for this analysis, the synthetic model parameters, as well as the exploration of additional explored hyperparameters, can be found in Tutorial 5 of the online documentation.

In this section, we use a synthetic network of $N = 20$ stations, distributed randomly, that is only affected by a single transient process and white noise, to explore the dependence of the model error on the number of stations used. The noise level relative to the maximum amplitude of the transient signal, σ , is one of the hyperparameters we vary. The other variable is the number of stations $2 \leq n \leq N$ used by the spatial L0-regularized solver. For each test case, we therefore subsample the original network to create a subnetwork of smaller size n , comprised of randomly selected stations. (We also calculate the result of using a local L0-regularized solver for comparison, where by construction $n = 1$). The number of samples m , for each n to test, is given by the maximum of either the amount of possible permutations, or a defined maximum value M based on computational considerations ($M = 50$ in our case).

For each σ and each n , we therefore have m samples to test. The metric we choose to compare is the root-mean-squared true model error (RMSE), calculated from the final fit of each sampled subnetwork (ensuring the solvers iterate long enough to converge). For each n , we therefore compute the double mean of the RMSE, ϵ , first across the subnetwork, and then across samples. We also compute the standard deviation σ_ϵ of the samples of the subnetwork-wide mean RMSEs.

Fig. B.6 shows the results of our experiment. For all of the cases, the mean RMSE ϵ decreases with increasing number of stations used in the fitting process (approximately by $1/\sqrt{n}$). Furthermore, the variance of the errors decreases as well. Importantly, for the case of $\sigma = 3$ (i.e., the white noise standard deviation is three times the maximum magnitude of the transient signal), the local L0-regularized solution has a high error variance centered close to the maximum allowable error (defined as not fitting a transient at all). Including multiple stations in the estimation process, however, decreases the mean error and error variance significantly — with 20 stations, as low as the mean error for the local L0-regularized solution for $\sigma = 1$. In the highest noise case presented here, $\sigma = 10$, most local L0-regularized solutions actually overfit the data. Incorporating spatial awareness prevents the solver to do so. Overall, as shown by the reduction of error, error variance, and susceptibility to overfitting, the importance of using spatial awareness for transient model fitting




Figure B.6: Sample mean (ϵ , colored lines) and sample standard deviation (σ_{ϵ} , vertical error bars) for all the sampled subnetworks as a function of the number of stations used in the solution process (n , horizontal axis), and noise level ratio (σ , different colors). The dashed, horizontal grey line corresponds to the maximum allowable error if no transient signal is fitted at all. The dotted grey line is a reference line parallel to $1/\sqrt{n}$.

becomes clear.

B.3 Long Valley Caldera: Secular Velocity Comparison

The code for this analysis can be found in Example 2 of the online documentation.

B.3.1 Method

Qualitatively, our results of transient and seasonal constituents are comparable to, e.g., Ji et al. (2013), Montgomery-Brown et al. (2015), and Silverii et al. (2020). To quantitatively validate the decomposition of the input timeseries from the Long Valley Caldera Region (LVCR) into its different constituents, we would need published, already-decomposed timeseries for the same study area. However, we are not aware of such products, and reproducing decompositions based on individual studies is beyond the scope of the paper.

A different way to still be able to perform a quantitative validation of our method is to recognize the fact that if our method is successful at distinguishing motion due to transient processes from longterm, secular motion (while still taking into account seasonal, seismic, and maintenance signals), then such estimates of secular motion should be free of physical influences other than longterm plate motion and deformation. Specifically for the Long Valley Caldera, we would expect that on top of a “background” field of motion, we would not see any influence from the magmatic caldera inflation. (Of course, if the caldera intrusion has a steady-state component, we would have no way of inferring this from only GNSS data, and are therefore neglecting this possibility.)

Geodetically, we can fit an average field of motion by assuming our study area (a circle of 100 km radius around the caldera center) is, to first order, approximated by a rigid body moving on a sphere. We can then estimate a best-fit rotation matrix (or equivalently, an Euler pole) using standard weighted least-squares (e.g., Goudarzi et al., 2014).

In this section, we compare the results obtained using DISSTANS and its spatiotemporal L0 regularization approach with published MIDAS-derived secular velocities (Blewitt et al., 2016) and the Geodesy Advancing Geosciences and EarthScope (GAGE) facility’s secular velocities (Herring et al., 2016). We first build a `Network` object that contains all three different velocity models. Then, the `Network.euler_rot_field()` method calculates the predicted velocity due to best-fit motion on a spherical Earth. Lastly, we remove the best-fit “background” secular velocities from the previously estimated, “model” secular velocities to produce “residual” secular velocities. The smaller these residuals, the better can our study area be approximated

	DISSTANS	MIDAS	GAGE
All	2.692	3.065	3.883
Outside LVCR	2.880	2.788	2.789
Inside LVCR	2.490	3.318	4.730

Table B.2: Root-Mean-Squared (RMS) residual magnitudes (in mm/a) between the modeled and background horizontal secular velocities for this study (DISSTANS) as well as the published velocities from MIDAS (Blewitt et al., 2016) and GAGE (Herring et al., 2016). The rows correspond to different subsets of the data over which the RMS is calculated. “Inside LVCR” corresponds to the stations shown in the lower panels of Figs. B.7–B.9, “All” to the ones shown in the upper panels of Figs. B.7–B.9, and “Outside LVCR” to the ones that are in the latter but not in the former. For stations outside the LVCR, the models produce approximately the same residual RMS (approx. $\pm 3\%$), but within the LVCR, DISSTANS reduces the residuals by approx. 25–47%, respectively.

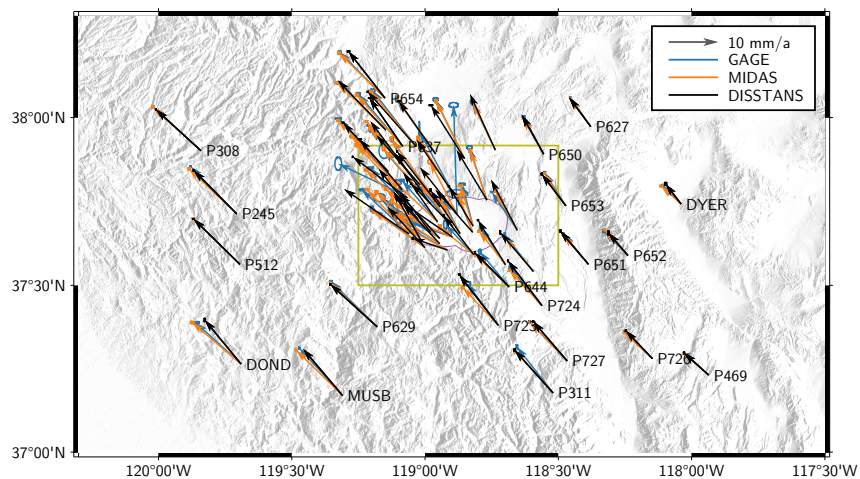
by a rigid plate.

In our comparison, we do not include methods that rely on a priori removal of a secular velocity in order to extract the transient. These approaches currently represent the majority of transient extraction methods; e.g., in combination with filtering (e.g., Silverii et al., 2020), with vbICA (e.g., if the timeseries is strongly correlated, Gualandi et al., 2016), with Singular Spectrum Analysis (SSA, e.g., Walwer et al., 2016), or with stacking (e.g., Kano and Kato, 2020). We omit these methods on the basis that they either do not claim to capture the longterm secular velocity in the first place (e.g., because the analyzed timespans are short, Kano and Kato, 2020) or that they are heavily reliant on assumptions (e.g., assuming a certain timespan represents steady-state velocity, Silverii et al., 2020).

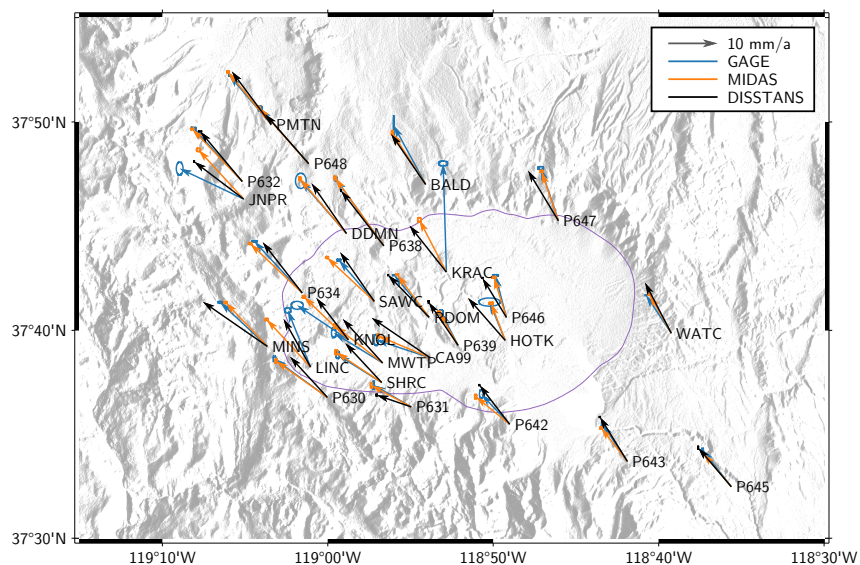
B.3.2 Results

We first want to note that neither the MIDAS nor the GAGE solution explicitly aim to model transient processes, with the exception of postseismic, decaying transient motion. While MIDAS does aim to be robust against shortterm transients, in general, our comparison is therefore not a “fair” one — both MIDAS and GAGE velocity fields are estimated in a fully-automated fashion and for the majority of global GNSS stations, they provide high-quality, reliable secular velocity estimates. The goal of this subsection is simply to highlight the differences in model results owing to our explicit modeling of transient processes.

Fig. B.7 shows the modeled horizontal secular velocities from the DISSTANS, MIDAS, and GAGE solutions. While over the entire study region, the velocities mostly match each other, differences are obvious when zooming into the Long

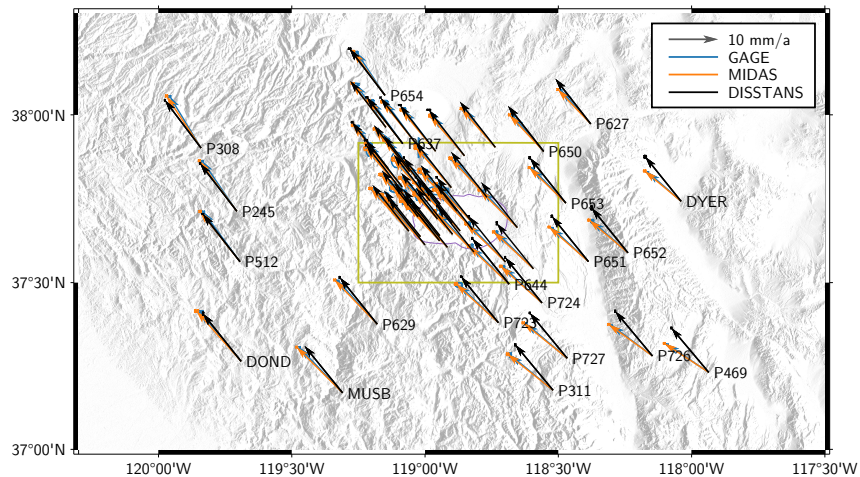


(a) Model (Study Region)

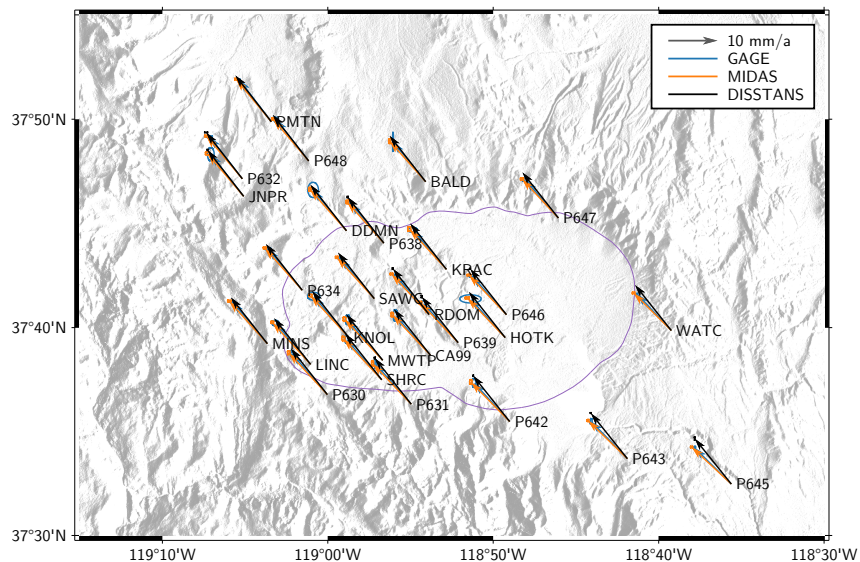


(b) Model (Long Valley Caldera Region)

Figure B.7: Upper panel: Modeled horizontal secular velocities in the study region for the three different solutions DISSTANS (Section 2.4.2), MIDAS (Blewitt et al., 2016), and GAGE (Herring et al., 2016) (in GAGE's North America-fixed reference frame). The Caldera ring fault (USGS Quaternary Fault Database, Bailey, 1989) is shown in purple. Uncertainties are shown as one-standard-deviation ellipses for each solution. The green rectangle shows the extent greater Long Valley Caldera Region (LVCR). Lower panel: Same as upper, zoomed into the LVCR. While the DISSTANS-derived velocities mostly match the published velocities outside the LVCR, they are significantly different within the LVCR.

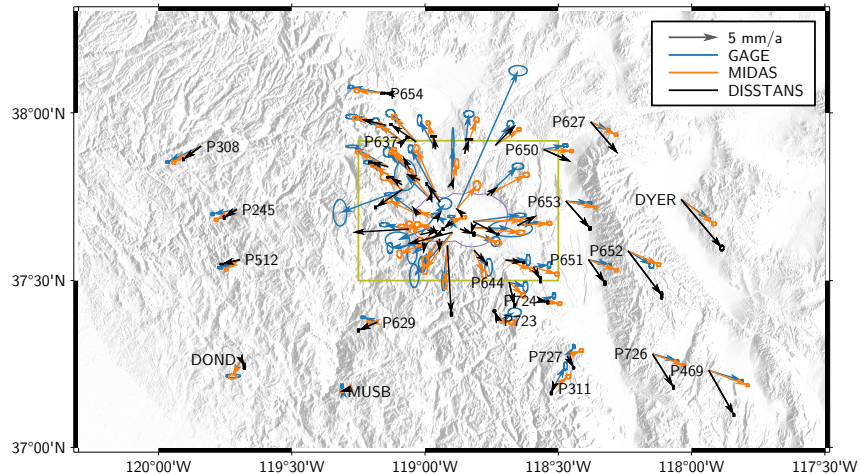


(a) Background (Study Region)

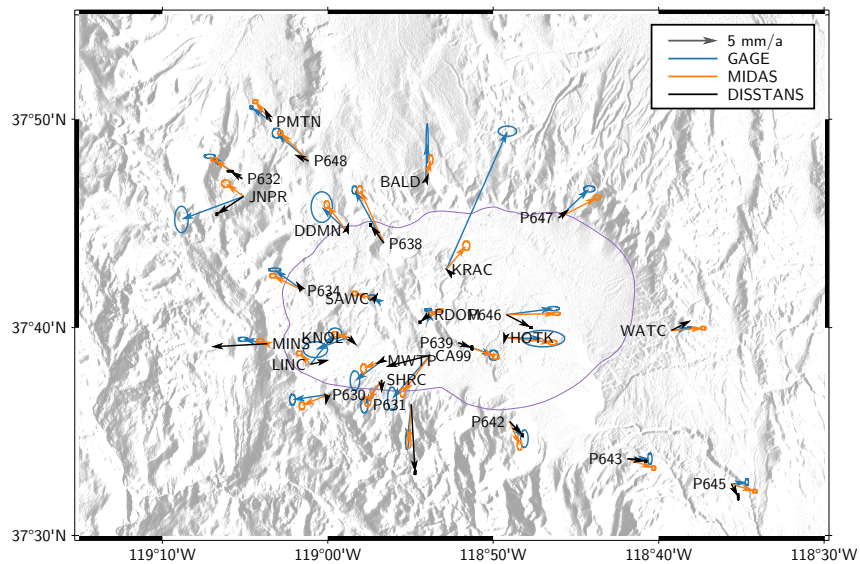


(b) Background (Long Valley Caldera Region)

Figure B.8: Background velocity fields as calculated by the best-fit Euler pole for the entire study area and the Long Valley Caldera Region in the upper and lower panels, respectively. Uncertainties, fault outlines, and colors are the same as in Fig. B.7. The DISSTANS-derived background velocity field slightly differs from the MIDAS- and GAGE-derived fields, but exhibit the same overall pattern.



(a) Residuals (Study Region)



(b) Residuals (Long Valley Caldera Region)

Figure B.9: Residual secular velocities for the entire study area and the Long Valley Caldera Region (LVCR) in the upper and lower panels, respectively. Uncertainties, fault outlines, and colors are the same as in Fig. B.7. The DISSTANS solution has smaller residual velocities in the LVCR than the MIDAS and GAGE solutions. (Note the different vector scale length.)

Valley Caldera Region (LVCR) itself. Calculating the best-fit velocities assuming a rigid-body motion for the three different solutions (independently) yields similar results (Fig. B.8) (differences around or below millimeter/year level). Excluding stations inside the LVCR in the estimation process also does not affect the resulting velocities significantly (differences on the order of half a millimeter/year).

Fig. B.9 shows the residual velocities (difference between modeled horizontal and best-fit background velocities). Qualitatively, the residuals from the DISSTANS solution are visibly reduced inside the LVCR compared to the MIDAS and GAGE solutions. Specifically, the MIDAS and GAGE solutions show a clear expansion component for stations in or near the caldera itself; this expansion pattern is much less prominent in the DISSTANS solution. Outside of the LVCR, all residuals show coherent patterns of motion, indicative of the imperfection of the assumption on the background velocity field (see below). Table B.2 quantifies the differences between model residuals using the Root-Mean-Square (RMS) of the residual magnitudes (i.e., the length of the residual vectors). Crucially, in the LVCR, where we expected the residuals to decrease by modeling the transients explicitly, we find that they are reduced by 25–47% (depending on the model). This reduction implies that the original modeled secular velocity field produced by DISSTANS more closely approximates the homogenous background velocity field. The reduction of residuals inside the LVCR is accompanied by only a small increase in residuals of about 3% outside the LVCR. Interestingly, the residual RMS of our velocity model is more similar between stations outside and inside the LVCR (range of 0.4 mm/a), whereas the residuals of the MIDAS and GAGE solution show a larger variance (ranges of 0.5 and 1.9 mm/a).

Overall, we interpret the reduction of residuals in the Long Valley Caldera Region for the DISSTANS solution to demonstrate the benefit of spatial awareness and explicit modeling of transient processes. By separating transient from longterm motion in a spatially-aware framework, the resulting secular velocity field is more homogenous than the MIDAS and GAGE solutions, and diminishes significantly the effect of magmatic inflation periods on the secular velocity estimate in the vicinity of the caldera.

B.3.3 Note on the assumed background velocity field

By comparing the secular velocities instead of the displacement timeseries during transient episodes, we are able to show quantitative differences between our solution

and two other published secular velocity fields — the MIDAS (Blewitt et al., 2016) and GAGE (Herring et al., 2016) models. To demonstrate the effect of explicitly modeling transient processes on the resulting estimated secular station velocities, we have furthermore estimated and removed a “background” field of motion from the secular velocities, and shown the residual velocities. We take the background field of motion to be the best-fit velocity field for our small study area (100 km radius around the Long Valley Caldera center), assuming rigid body motion on a sphere. Note that we do not assume the background velocity field to represent the true underlying longterm velocity (e.g., we expect distributed shearing in our study area because of the remote North America-Pacific plate boundary), only that such a velocity field should be able to reproduce the modeled secular velocities to first order (which is the case). The estimation of the background velocities is performed using simple, unregularized, weighted least-squares. The results obtained from the three different input fields (DISSTANS, MIDAS, and GAGE) are similar, supporting our assumption that the background field is able to capture most of the secular velocity signal. Therefore, using the residual velocity fields for our comparison simplifies the highlighting of the differences between the processing strategies of the three secular velocity models.

B.4 Long Valley Caldera: Seasonal Signals

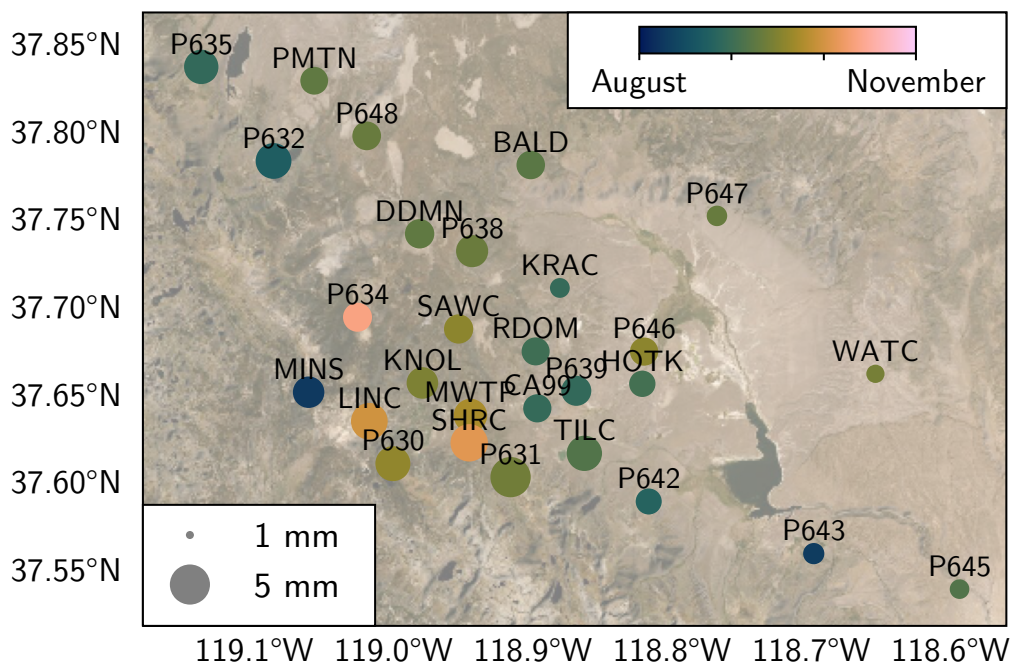


Figure B.10: Unregularized (average) component of the vertical seasonal fitted annual model. The marker size corresponds to the amplitude, and the color corresponds to the time of the sinusoid's maximum during the year. Background satellite imagery by Earthstar Geographics & Esri.




Figure B.11: Full vertical seasonal model (unregularized and regularized, annual and semi-annual) for the selected stations from Fig. 2.5, in the same order. Black dots are the overall model's residuals, centered on the seasonal model.

*Appendix C***EXPLORATION OF THE NORTHERN JAPAN 3D MESH****C.1 Finite Fault Models Defining Asperities**

1. Ibaraki

- 2011: Simons et al. (2011)

2. Miyagi

- 1978: Yamanaka and Kikuchi (2004)
- 2005: U. S. Geological Survey (2017, Event ID p000dx2) and Miura et al. (2006)

3. Sanriku

- 1968: Nagai et al. (2001)
- 1994: U. S. Geological Survey (2017, Event ID usp0006qh3), Nakayama and Takeo (1997), Nagai et al. (2001), and Hayes (2017)

4. Tokachi

- 2003: U. S. Geological Survey (2017, Event ID p000c8kv), Yamanaka and Kikuchi (2003), Koketsu et al. (2004), Tanioka et al. (2004), Yagi (2004), Hayes (2017), and Kobayashi et al. (2021)

5. Nemuro

- 1973: Nishimura (2009)

6. Tohoku

- 2011: U. S. Geological Survey (2017, Event ID p000hvnu), Ammon et al. (2011), Fujii et al. (2011), Hayes (2011), Ide et al. (2011), Lay et al. (2011), Shao et al. (2011), Yagi and Fukahata (2011), Yamazaki et al. (2011), Gusman et al. (2012), Satake et al. (2013), and Yue and Lay (2013)

7. Hokkaido

- 1600s: Nanayama et al. (2003)

A majority of the finite fault models were downloaded using SRCMOD (Mai and Thingbaijam, 2014).

C.2 Differences Between Mesh Discretizations

Our Bayesian inversion approach (Section 3.5), based on the Markov-chain Monte Carlo methods presented in Minson et al. (2013), Duputel et al. (2014), Jolivet et al. (2014), Duputel et al. (2015), and Jolivet et al. (2015), requires thousands to millions of evaluations of the forward model (Section 3.4) to approximate the posterior probability distribution of the rheological parameters to be estimated. Therefore, to make full-cycle probabilistic inversions feasible, the forward model has to run within minutes or less. One significant factor for the runtime is the number of patches in the discretized mesh representing the fault interface, with fewer and larger patches leading to faster cycle integrations than more and smaller ones.

In Section 5.2.1, we present two versions of the discretized plate interface between the Pacific and Eurasian plates, and have to rely on the coarser of the two for all inversions. In this section, we assess whether the coarse mesh leads to similar results than the fine mesh, or not. We cannot run the same inversion with the fine mesh for the computational reasons alluded to above, which would be the most natural way to determine the equivalence of the obtained results. We can, however, sample rheological parameters from the posterior distribution obtained with the coarse mesh, and run forward models on the fine mesh using these parameter samples. This is straightforward since the spatial variability of the rheological parameters is defined independent of the mesh (see Section 5.2.3).

Figs. C.1 and C.2 show the difference between the model predictions of the coarse and fine mesh in terms of preseismic velocity and postseismic cumulative displacement, respectively. In both cases, the differences are, on average, significantly smaller than the residuals between the predictions of the coarse mesh and the observations (compare Figs. 5.11 and 5.17). The only exception to this is the predicted postseismic cumulative displacement close to the Miyagi asperity. Here, the modeled displacement is extremely sensitive to the shape of the asperity and the values of the tapered slip outside the asperity. Fig. C.4 shows the fault slip rate over time for the fine mesh (compare to the results of the coarse mesh in Fig. 5.8), and highlights that the tapered slip is, because of the smaller patch sizes, much more localized around the asperities. Generally, however, the fault slip rate on the fine

mesh exhibits the same first order shape and decay as the slip rates from the coarse mesh.

We can visualize the overall correlation between the modeled surface displacement for the two meshes by plotting the coarse predictions versus the fine ones. Fig. C.3 shows the strong correlation ($R^2 = 0.98$) between the outputs of both forward models for each component separately. Moreover, it shows that the largest differences between the coarse and fine predictions are located close to the Miyagi asperity, as expected from Figs. C.2 and C.4. Fig. C.5 shows that the excess moment release on the fault also closely tracks the predictions from the coarse mesh (compare Fig. 5.21). Finally, we calculate the average displacement error from the fine forward model as $\delta_d = 0.43$. This value is closer to the $\delta_d = 0.40$ error from the coarse forward model than the errors from the other test inversions summarized in Table 4.2, except for cases (2) and (6), which only add more samples or remove the preseismic information, respectively.

Jointly, the presented qualitative and quantitative evidence suggests that the forward models using the fine mesh discretization shown in Fig. 5.2 produce results similar to forward models using the coarse mesh from Fig. 5.3. Since the rheological parameters have not been optimized for the fine mesh, however, we choose to show the results from the coarse mesh throughout Chapter 5.




Figure C.1: Interseismic velocity differences for Northern Honshu between forward models using the fine and coarse mesh. The arrows show the horizontal difference, colored by the vertical difference. The asperities included in the inversion framework are shown by outlined and filled polygons for the fine and coarse meshes, respectively (same colors as in Fig. 5.1). Other map content as in Fig. 5.11.




Figure C.2: Postseismic cumulative displacement difference for the timeframe 2011–2024 in Northern Honshu between forward models using the fine and coarse mesh. The arrows show the horizontal difference, colored by the vertical difference. The asperities included in the inversion framework are shown by outlined and filled polygons for the fine and coarse meshes, respectively (same colors as in Fig. 5.1). Other map content as in Fig. 5.11.




Figure C.3: Correlation between the predicted surface displacements at all timesteps for forward models using the coarse and fine mesh, split up into the three data components, and offset for visual clarity. Blue, red, and green dots refer to the East, North, and Up components, and are shaded by distance to the Miyagi asperity, with smaller distances marked by darker colors. The joint best-fit linear correlation slope for all components is shown by the dashed gray line; the solid gray line shows a 1:1 correlation.

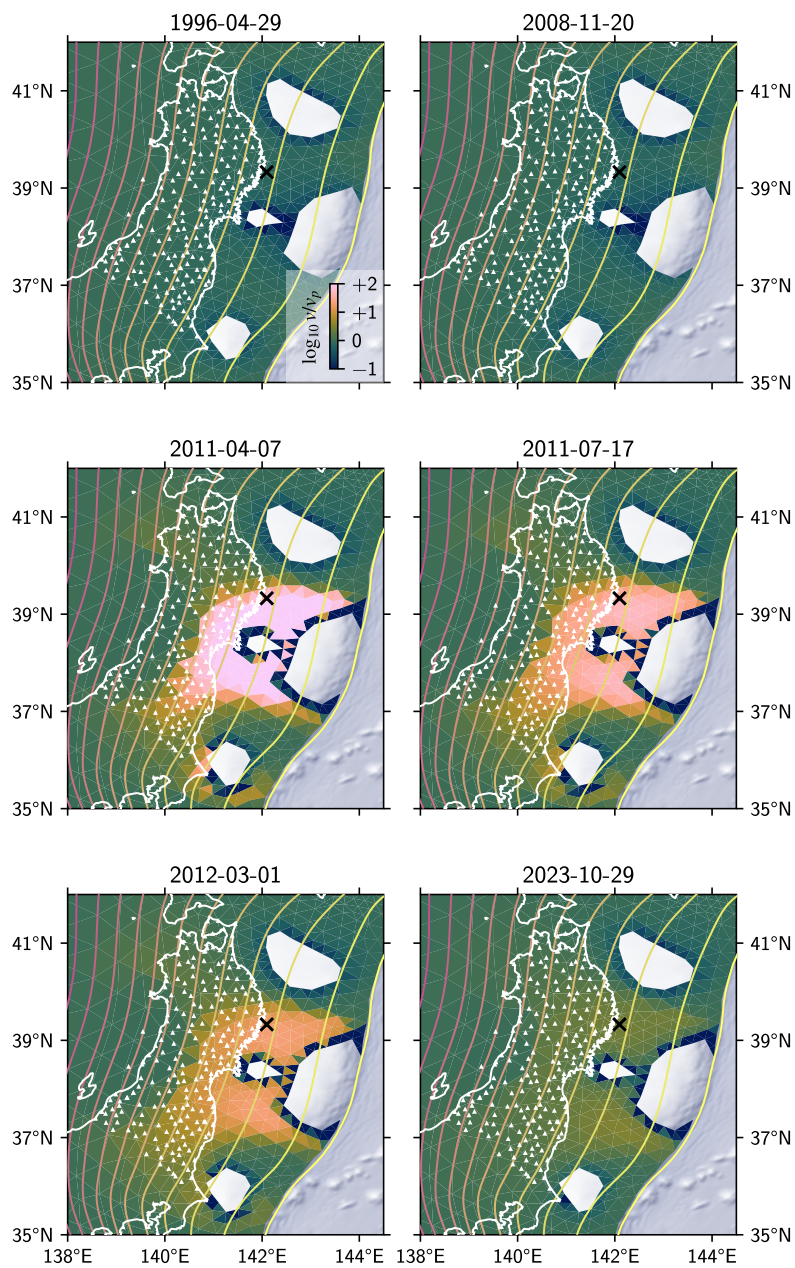


Figure C.4: Mean fault slip rate on the Northern Japanese megathrust for the forward models using the fine mesh at six different times (two pre- and four postseismic). The discretized interface is colored according to the base-10-logarithm of the ratio between inferred slip rate and the assumed plate convergence rate. Additional map elements as in Fig. 5.8.




Figure C.5: Mean excess moment magnitude over the entire creeping fault interface for the interval between the observation timestamps for the forward models using the coarse and fine mesh. Does not include contributions of the plate convergence.

C.3 Maximum Fit Quality of the Tohoku-oki Reference Model using Kinematic Regression

To assess to which degree the displacement misfits of the earthquake cycle inversion framework (Chapter 5) is due to the forced evolution of slip according to a rheological model, we also perform static least-squares regression for the fault slip rate over time, using the same mesh and observations. To take into account GNSS stations starting or ending service, the data we fit is the average velocity (or equivalently, cumulative displacement) between two timestamps (ignoring the coseismic period).

We first perform the inversion without any regularization or constraints. In this case, the fit to the data is expected to be as perfect as possible given the fault interface discretization and the choice of the displacement kernels (Green's functions). We omit any figures for the unregularized fit since the residuals are close to zero. To obtain a more realistic, albeit degraded, fit to the data, we then force the slip rate to always be within 45° of the rake angle. We furthermore regularize the inversion in a mechanically-meaningful fashion by adding the stress kernel to the least-squares cost function (similar to Lindsey et al., 2021), and optimize the scalar penalty factor of the stress kernel using the RMS of the residuals. This furthermore allows the regularization in the preseismic timeframe to be applied to the deviation from the plate rate, rather than absolute velocity.

Fig. C.6 shows the best-fit fault slip rates for a globally optimal penalty factor. Pre-Tohoku-oki, the recovered average slip rate is close to the plate rate almost everywhere, as expected. Around the asperities, which are fully-locked just as in the cycle inversion, the slip rates are slightly below plate rate, showing the effect of the stress-based regularization. Immediately after the Tohoku-oki earthquake, the average slip rates are about one order of magnitude above the plate convergence rate almost everywhere below the data coverage, as well as around the rupture area. Over the course of the years after 2011, the slip rates start to decay smoothly everywhere, hinting at a below-plate-rate zone down-dip of the Miyagi asperity. We skip an interpretation of the very deep fault sections (below the west coast of Honshu) as well as the areas far north and south of the observing GNSS stations.

The resulting timeseries at representative stations landwards of the Tohoku-oki rupture area are shown in Fig. C.7, with an average displacement error of $\delta_d^{\text{kin}} = 2.4$ cm. They show an excellent agreement with the observations for most components and stations. Figs. C.8 and C.9 show that the misfit is mostly systematic. In the preseismic period, there seems to be an unmodeled shortening of the west coast of Honshu,

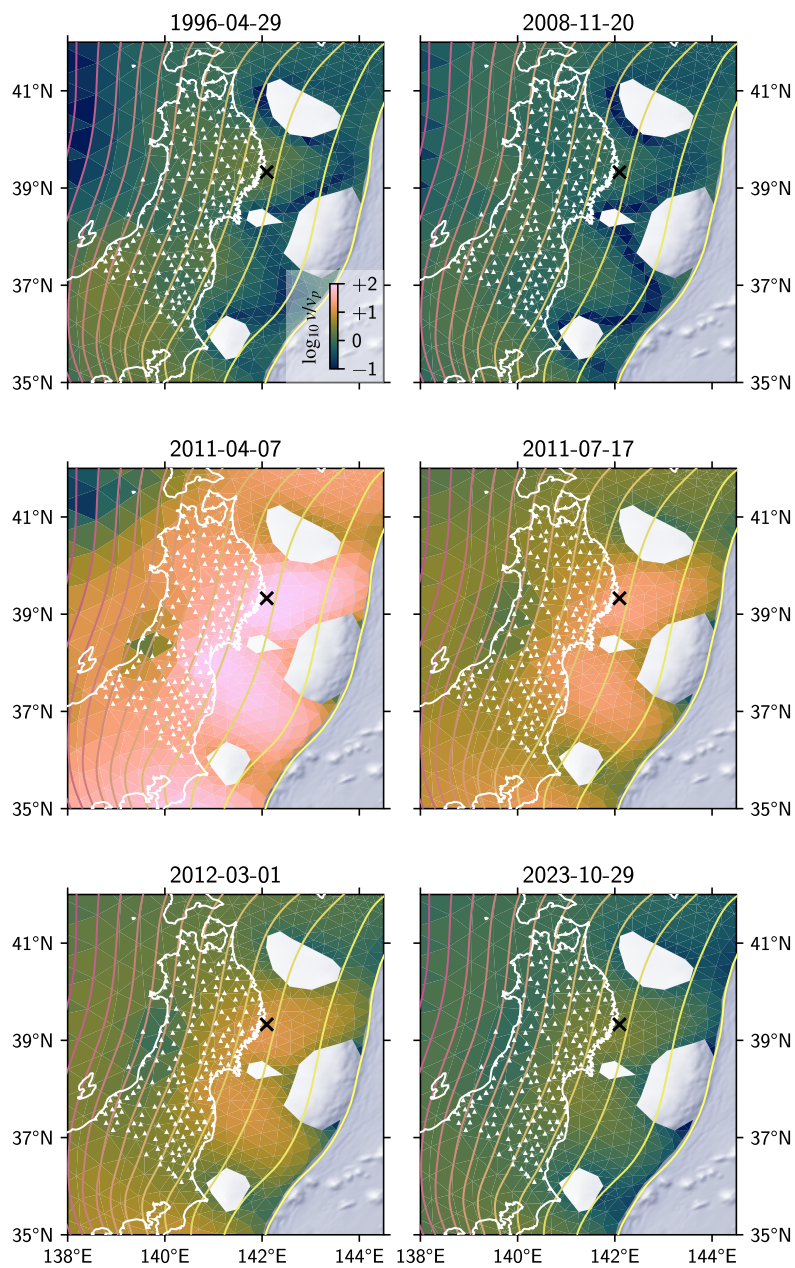


Figure C.6: Fault slip rate on the Northern Japanese megathrust as inferred by the kinematic inversion at six different times (two pre- and four postseismic). The discretized interface is colored according to the base-10-logarithm of the ratio between inferred slip rate and the assumed plate convergence rate. Trench, slab depths, and GNSS station locations as in Fig. 5.18.

as well as some edge effects in the far north and south. For the postseismic time-frame, the trenchward residual of the model fit suggests that more afterslip inside or close to the Miyagi asperity would be necessary to improve the misfit. When compared to the total signal of postseismic transient after the 2011 earthquake, however, the residuals appear small. Finally, Fig. C.10 shows the excess moment release over time on the fault interface. The overall trend of the moment release is similar to the trend of the earthquake cycle models, indicating that the modeled fault slip rates in our proposed framework are, to first order, a good approximation of the real kinematics (even though the absolute value of the release is about 0.2 orders of magnitude higher).

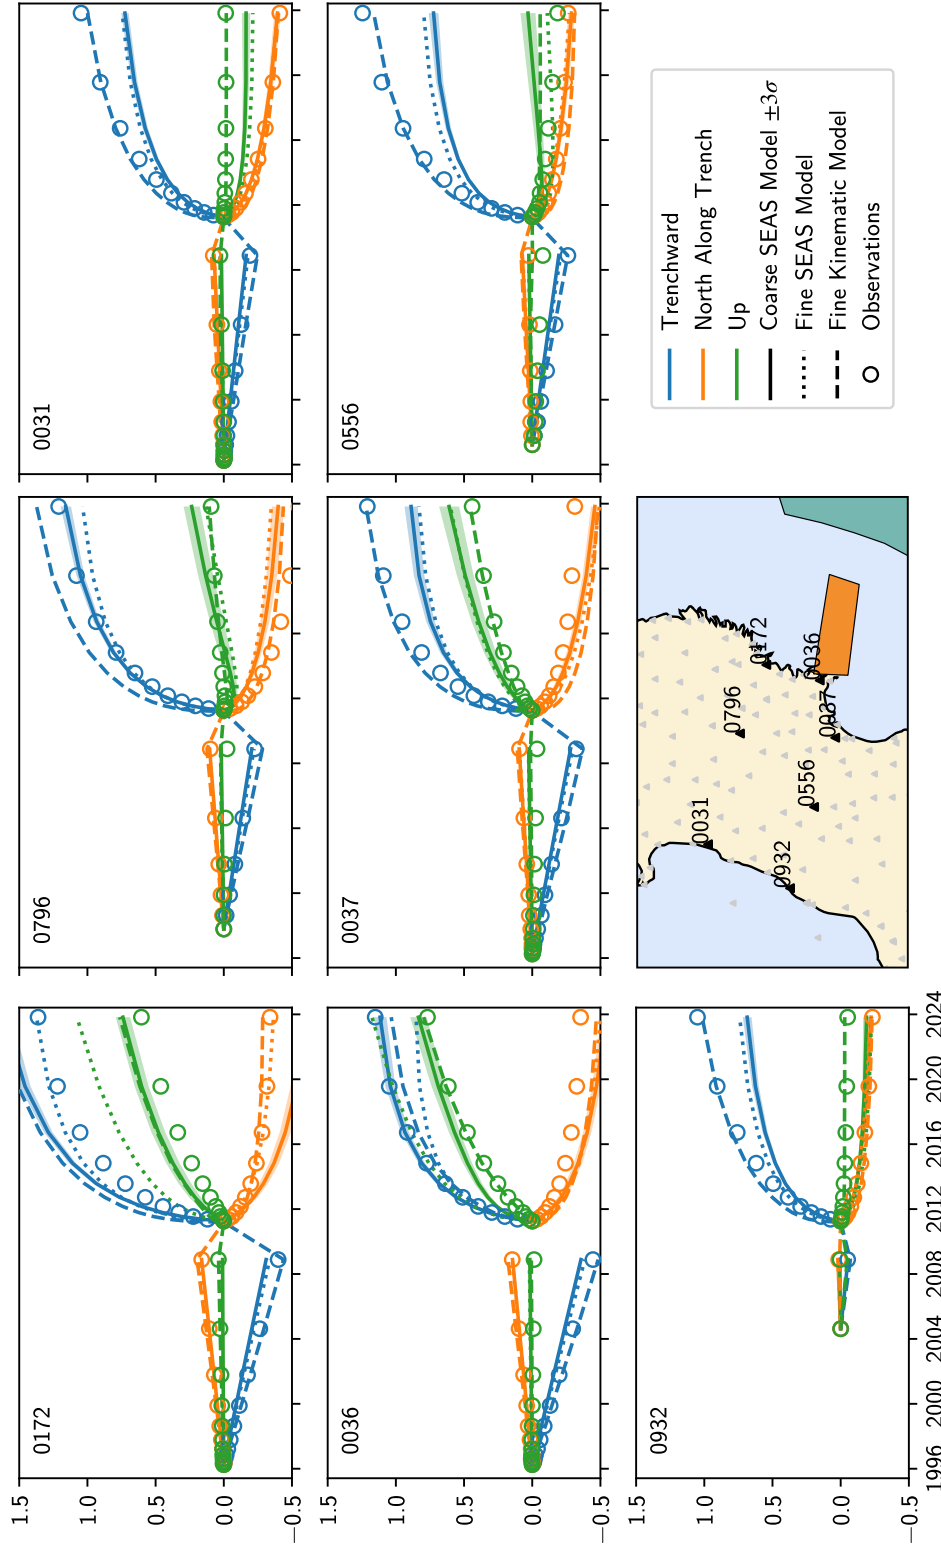


Figure C.7: (Same as Fig. 5.10 but including the kinematic inversion results.) Observed and modeled timeseries for the reference model and kinematic inversion at select stations. The vertical axis is displacement [m] and the horizontal axis is time. Trenchward, along-strike, and vertical data components are shown in blue, orange, and green, respectively. Dashed areas show the sample mean and uncertainty of the cycle model output based on the 2011 Tohoku-oki earthquake. Continuous lines and shaded areas show the best-fit displacement timeseries from the kinematic inversion. The center bottom panel shows a zoomed-in map view of the station network, where the stations shown in this figure are labeled and marked with black triangles. The Miyagi and Tohoku asperities are shown as orange and green polygons.




Figure C.8: Interseismic velocity residuals from the kinematic inversion for Northern Honshu Island, Japan. The residuals are defined as the difference between the observed velocities and the best-fit velocity from the kinematic inversion. The arrows show the horizontal residual, colored by the vertical residual. The asperities included in the inversion framework are shown by polygons (same colors as in Fig. 5.1). Trench and slab as in Fig. 5.18.




Figure C.9: Postseismic cumulative displacement residuals from the kinematic inversion for the timeframe 2011–2024 in Northern Honshu. The residuals are defined as the difference between the observed cumulative displacement in 2024 and the best-fit displacement from the kinematic inversion. The arrows show the horizontal residual, colored by the vertical residual. Trench, slab, and asperities as in Fig. C.8




Figure C.10: Excess moment magnitude over the entire creeping fault interface for the interval between the observation timestamps for the kinematic and cycle models. Does not include contributions of the plate convergence.

Footnotes for Section 2 and Appendices A and B:

¹Automatic detection only by GipsyX.

²Only GAMIT/GLOBK.

³Here, we use the definition of Candès et al. (2008): $\|\mathbf{x}\|_0 = |\{i : x_i \neq 0\}|$, i.e., the number of non-zero elements.