

LA-UR-20-22489 (Accepted Manuscript)

An exponential build-up in seismic energy suggests a months-long nucleation of slow slip in Cascadia

Hulbert, Claudia L.
Rouet-Leduc, Bertrand Philippe Gerard
Jolivet, Romain
Johnson, Paul Allan

Provided by the author(s) and the Los Alamos National Laboratory (2021-07-06).

To be published in: Nature Communications

DOI to publisher's version: 10.1038/s41467-020-17754-9

Permalink to record: <http://permalink.lanl.gov/object/view?what=info:lanl-repo/lareport/LA-UR-20-22489>

Disclaimer:

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

ARTICLE



<https://doi.org/10.1038/s41467-020-17754-9>

OPEN

An exponential build-up in seismic energy suggests a months-long nucleation of slow slip in Cascadia

Claudia Hulbert^{1,2}✉, Bertrand Rouet-Leduc¹  ², Romain Jolivet^{1,3}  & Paul A. Johnson²

Slow slip events result from the spontaneous weakening of the subduction megathrust and bear strong resemblance to earthquakes, only slower. This resemblance allows us to study fundamental aspects of nucleation that remain elusive for classic, fast earthquakes. We rely on machine learning algorithms to infer slow slip timing from statistics of seismic waveforms. We find that patterns in seismic power follow the 14-month slow slip cycle in Cascadia, arguing in favor of the predictability of slow slip rupture. Here, we show that seismic power exponentially increases as the slowly slipping portion of the subduction zone approaches failure, a behavior that shares a striking similarity with the increase in acoustic power observed prior to laboratory slow slip events. Our results suggest that the nucleation phase of Cascadia slow slip events may last from several weeks up to several months.

¹ Laboratoire de Géologie, Département de Géosciences, École Normale Supérieure, PSL Université, CNRS UMR 8538, Paris, France. ² Los Alamos National Laboratory, Geophysics Group, Los Alamos, NM, USA. ³ Institut Universitaire de France, 1 rue Descartes, 75005 Paris, France. ✉email: claudia.hulbert@ens.fr

Since their discovery in Japan at the turn of the millennium^{1,2}, slow slip events and associated tectonic tremor and low-frequency earthquakes (LFEs) have been identified in most subduction zones as well as other tectonic environments^{2–9}. Slow slip events release energy over much longer durations than classic earthquakes, from a few days to months or even years². In subduction settings, slow slip occurs deep along the subduction interface (i.e. roughly 50 km depth), down-dip from the nucleation zone of damaging earthquakes, at the transition from brittle to ductile deformation¹⁰. At such depths, slow slip and tremor are thought to take place where temperatures drive dehydration of subducting material that increases pore pressure, inhibiting dynamic failure in the brittle/ductile transition regime^{11,12}. The slowly slipping region is considered to mark the transition from unstable (seismogenic) to stable (creeping) sliding and therefore may define the depth limit of megathrust ruptures¹².

A growing body of literature suggests that slow, aseismic slip and rapid, seismic slip bear strong resemblance^{9,13,14}. In particular, recent studies find that they follow comparable scaling relationships in terms of duration and magnitude^{14–16}. Slow slip events may therefore provide an opportunity to study fundamental rupture physics, as they take place over long periods of time without radiating large amplitude seismic waves. Considering the lack of observational evidence of earthquake nucleation mechanisms, we propose to explore the period leading up to a slow slip event as a window into a better understanding of the nucleation of a slip instability in nature. Here, we transpose a methodology developed on laboratory experiments to the occurrence of slow slip events in Cascadia.

Laboratory studies of slow slip¹³ from a bi-axial shear device^{17–19} suggest that the amplitude of acoustic noise coming from a fault follows characteristic patterns throughout the slip cycle. Such patterns allow us to estimate key properties of the laboratory fault, including friction on the fault as well as fault displacement rate. In a first effort to generalize these results to a natural fault system, the analysis of slow slip in Cascadia²⁰ revealed that statistical characteristics of continuous seismic signals can be used to estimate the displacement rate of GPS stations at the surface. These characteristics are related to seismic power, which is analogous to the acoustic power measured in laboratory experiments (we define seismic power as the average of seismic energy per unit of time, i.e. the squared measured ground velocity per unit of time). This proportionality between seismic power and surface displacement enables a quantitative characterization of slow slip events from seismic data.

Machine learning (ML) analysis of seismic data is an expanding field, with recent studies focusing on event detection²¹, phase identification²², phase association^{23,24}, or patterns in seismicity²⁵. In the following, we investigate whether seismic signatures can be found in the period leading up to any known manifestation of major slow slip occurrence anywhere in the Cascadia region. We find that for most episodic slip and tremor events, features within tectonic tremor frequency bands increase a few months before any detection of cataloged tremor or any geodetic signature of fault slip is made in the Cascades. We interpret this growth in the seismic power of the subduction zone as the signature of a nucleation phase that can be detected long before being observed in tremor catalogs or GPS data.

Results

Seismic power analysis and the occurrence of slow slip in Cascadia. We analyze seismic data on Vancouver Island, Canada, where the Juan de Fuca oceanic plate subducts beneath the North American plate (Fig. 1a). The quasi-periodic occurrence of slow

slip events (approximately every 14 months)^{6,26–30} is manifested by the North American plate lurching southwesterly over the Juan de Fuca plate, generating bursts of tectonic tremor over the area (Fig. 1b). Smaller slow slip events occurring between these large periodic events have been identified recently, pointing towards a large variability in the size and timing of slow earthquakes in the area³¹. The regular occurrence of slow slip events in Cascadia, especially in the vicinity of Vancouver Island²⁶, compares well to the afore-mentioned laboratory experiments. Furthermore, continuous seismic recordings in this region are available for well over a decade. Supervised ML used in the work described below requires robust training and testing sets including many slip events. The long history of recurrent slow slip observed in this region makes it an ideal case to (i) apply a methodology that has been developed in the laboratory and (ii) determine if there is information carried in the seismic signal characteristic of nucleation and upcoming failure.

We rely on the Pacific Northwest Seismic Network (PNSN) Tremor Logs²⁹ to identify slow slip failures. The Tremor Logs report ongoing slow slip on the basis of tremor activity detected anywhere in Cascadia by the PNSN. In what follows, the Tremor Logs will be used to identify slow slip timings. These timings are represented by vertical gray bars in Fig. 2b. Smoothed tremor rates in Vancouver Island from the PNSN tremor catalog²⁹ are also plotted for comparison.

Tremor logs take into account the entire Cascadia region and are not geographically limited to Vancouver Island. This is important for our analysis, as it precludes contamination of seismic data by an ongoing slow slip event. Indeed, slow slip events in Cascadia take place over a large section of the west coast of the USA and Canada, and the beginning and migration of events do not follow systematic patterns. Preceding any local manifestation of slow slip activity, seismic data might include information from an event that has already started elsewhere and is migrating towards our region of interest.

For this reason, using GPS displacement or only local tremor to identify failure times would not be robust for our analysis. GPS displacement is a local measurement, and so is the occurrence of local tremor, whereas seismic sensors may capture signatures of slipping segments located farther away. Furthermore, identifying slow slip with GPS requires temporal smoothing and consequently does not allow one to determine precisely when the rupture begins, which may introduce errors of days to weeks and ultimately arbitrarily improve the performance of our analysis. This explains why we favor the PNSN tremor logs to determine slip timing (note that tests using local peaks in tremor rates or local GPS displacement as a proxy for failure times provides similar or better results—see Supplementary).

We rely on nine seismic stations from the Plate Boundary Observatory³² (Fig. 1b). We find that borehole stations are much more robust than surface ones because of contamination by seasonal signals. A continuous, clipped and de-sampled seismic waveform for one seismic station is shown in Fig. 2a. We first process the seismic data by correcting for the instrument gain, for each day during the period analyzed (2005–2018). The daily data intervals are band-passed between 8 and 13 Hz (within successive bands of 1 Hz) and clipped, to limit the contribution of microseismic noise and earthquakes, respectively, and to focus on low-amplitude signals. These particular frequency bands have been identified in our previous work as the most informative of the behavior of the system²⁰. Once the data are pre-processed, for each day we compute a number of statistical features linked to signal energy. Building on our previous work, the features correspond to inter-quantile ranges of seismic data within tremor frequency bands (8–13 Hz, by 1 Hz increments). These features are representative of seismic energy in tremor frequency bands,

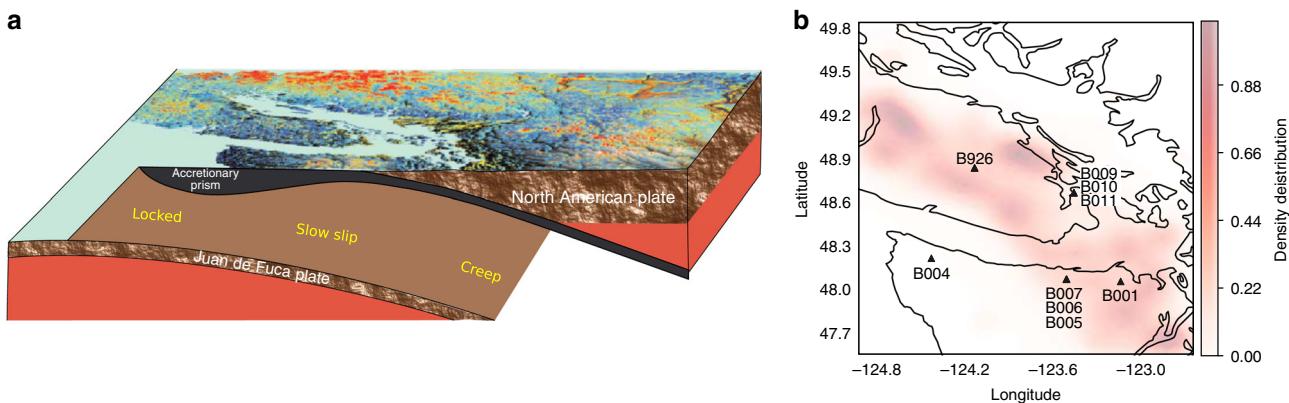


Fig. 1 Schematic of the subduction zone and seismic array analyzed. **a** Sketch of the subduction zone underneath Vancouver Island, Canada. **b** Map of the region analyzed. The seismic stations used in this study are shown with black triangles. Red shaded areas correspond to the density distribution of tremors from the PNSN tremor catalog from 2009 to late 2018.

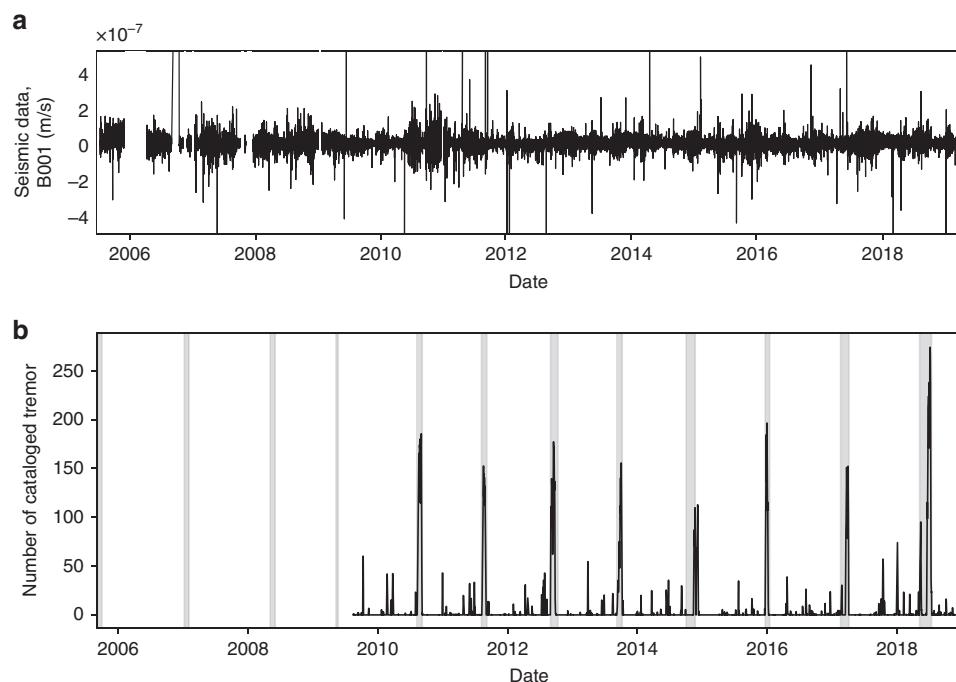


Fig. 2 Signals analyzed and timings of slow slip events. **a** Down-sampled (solely for visualization purposes), clipped (at 5E–7 m/s) continuous seismic waves for one station analyzed (B001). Our goal is to rely exclusively on continuous seismic waves to identify signatures preceding an upcoming slow slip event. **b** Slow slip timing from PNSN tremor logs (gray shaded areas), and smoothed (over 10 days) catalogued tremor rates in Southern Vancouver Island from the PNSN tremor catalog for comparison.

but with outlier values removed, which makes them more robust to signals not of interest for our analysis (such as earthquakes) and to potential sources of noise. A more extensive description of the features used can be found in the “Methods” section. These daily features are then averaged within a time window. Anomalous data points are detected within each window and removed before averaging. The results shown in Fig. 3 use a time window of 3 months (i.e. features are averaged over 90 days), but our methodology is robust to changes in the window size (see Supplementary). Each window is indexed by its latest day: the value of the features over the 3 months considered is associated with the last day of the window, to ensure that the analysis is made using only past data. Two successive time windows are offset by one day, and therefore contain 89 days in common. The averaged features over these time windows are used as input to

the ML algorithm. In the following, ‘seismic features’ will refer to those features averaged over a time window.

We apply a supervised ML approach to assess whether continuous seismic waves carry the potential signature of an upcoming slow slip failure. We assess whether a given time window of the continuous seismic data can be used to find signatures of impending failure for the next slow slip event. In the training phase, the algorithm takes as input the seismic features calculated from the first (contiguous) 50% of the seismic data (training set), and attempts to find the best model that maps these features to the time remaining before the next slow slip event (label or target). Details on how we build the model can be found in the “Methods” section.

Once a model is trained, it is evaluated on data it has never seen—the remaining (contiguous) 50% of the data (testing set). It

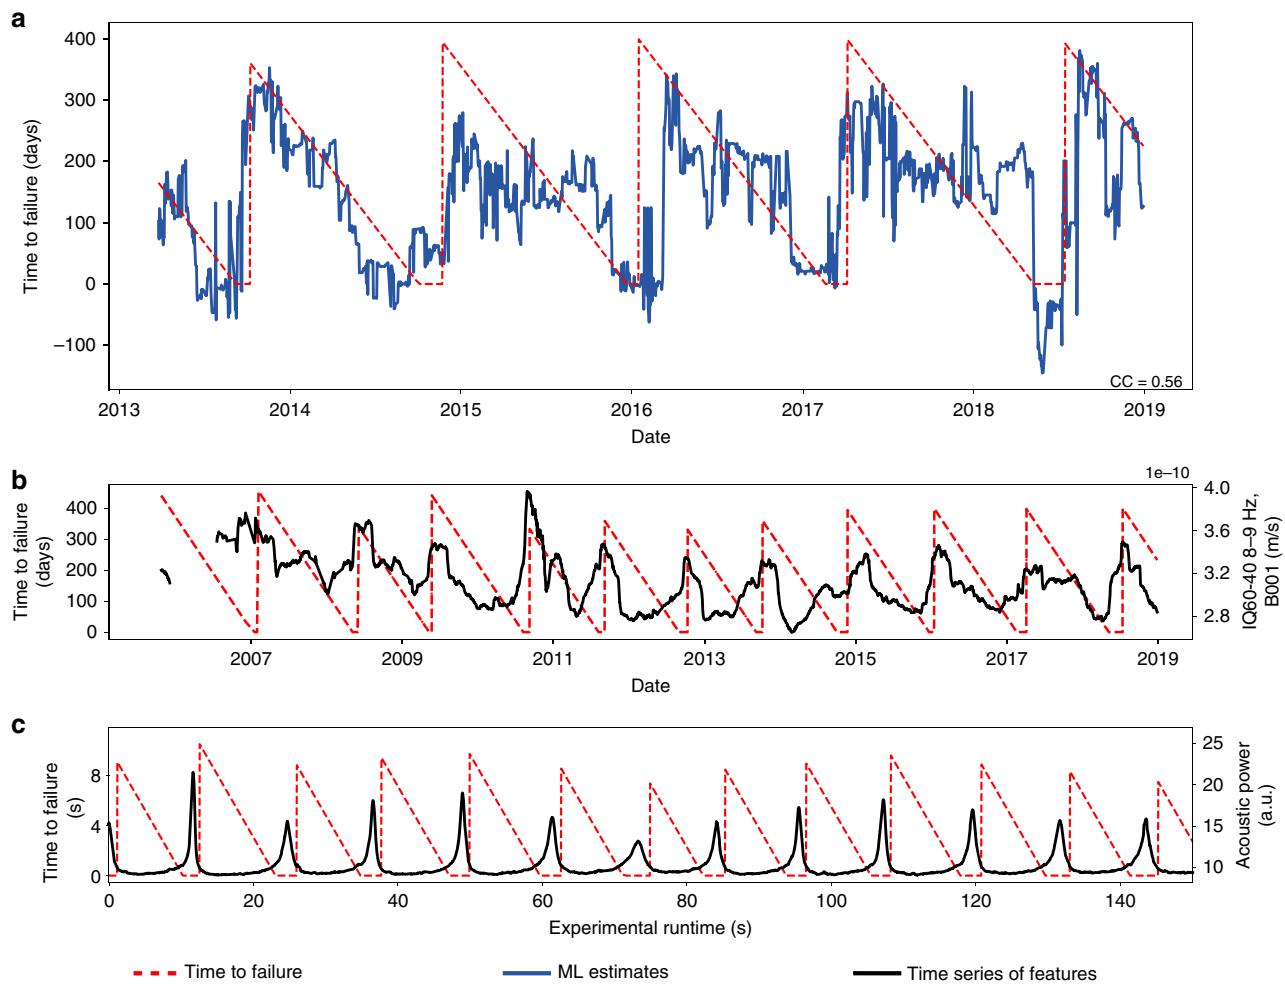


Fig. 3 Slip timing estimations, patterns identified by our model, and comparison to shear experiments in the laboratory. **a** ML estimates in testing of time to failure (in blue), and ground truth time to failure (in red) from the PNSN tremor logs. For most events, estimates are close to zero for several weeks preceding the rupture. **b** The most important feature identified by our model plotted against time, for the best stations (B001), for time window intervals of 3 months (black curve). This energy-based feature shows clear patterns with respect to the time remaining before the next slow slip event (left axis). **c** Shows the best statistical feature found in laboratory slow slip experiments for comparison (acoustic power, right-hand vertical axis). The best features in the laboratory **c** and Cascadia **b** are related to the energy of the seismic waves, and appear to follow similar patterns, with a progressive increase in amplitude that peaks towards the end of each slow slip event.

is important to note that in the testing phase, the model only has access to the seismic features calculated from the continuous seismic data, and has no information related to slow slip timing (the label). In the testing phase, the label is used exclusively to measure the quality of the model's estimates, i.e. how close these estimates are compared to the true label values obtained from PNSN Tremor Logs. If the model is able to estimate an imminent failure from seismic data it has never seen before, then it means that one or several features carry the signature of an impending slow slip event. We use Pearson's correlation coefficient (CC) as the evaluation metric, to compare the output of our model to the true test label values.

We rely on gradient-boosted trees algorithms that are relatively transparent in their analysis in contrast to many other methods³³. These algorithms can be probed to identify which features are important in the model predictions, and why. Identifying the important statistical features allows us to make comparisons with laboratory experiments, and gain insight into the underlying physics.

Estimating slip failure times from continuous seismic data. Estimations of the time remaining before the next slow slip event

on the testing set are shown in Fig. 3. Figure 3a shows the ML time to failure estimations using data from station B001 (in blue), and the measured time remaining before the next slow slip event (ground truth, dashed red line). This ground truth can be understood as a countdown to the next slip, and is equal to zero whenever the PNSN Tremor Logs reported an ongoing slow slip. Each point of the model estimations of slip timing (blue curve) is made from seismic features from the three preceding months.

The seismic data long before failure and during failure appears very different to the trained model, that easily distinguishes between these two extreme cases in all the examples in our testing set. Estimations far from failure (for large values of time to failure) are noisy and not very accurate, but are of lesser interest for our analysis. They are often characterized by a long plateau during which the estimations remain relatively constant, suggesting that none of our seismic features follow patterns in the early stages of the cycle that are informative of the upcoming rupture. Interestingly, a similar but shorter plateau can be observed in laboratory slow slip events as well¹³, especially following events of larger magnitudes. These observations suggest that slow slip events may be followed by a fundamentally unpredictable phase—maybe due to the local re-arrangement of

the system following the preceding failure. Other possible explanations could be that early signals are too small to be perceived amongst seismic noise, or that our selected set of features does not allow the algorithm to capture the early evolution of the system.

When close to failure (in the weeks preceding the rupture, when the time to failure approaches zero), the ML estimations are close to zero for most slow slip events—with the exception of the 2018 slow slip. Each point on the blue estimation curve in Fig. 3a is derived from a single time window of seismic data (built from the previous 90 days), that does not include any information from the rupture itself. Thus, the results suggest that at times closer to the end of the slow slip cycle, a snapshot of continuous seismic waves is imprinted with fundamental information regarding the upcoming failure of the system. The fact that estimations of failure times tend to drop abruptly also suggests that the evolution towards failure is less smooth than in the laboratory, and may be characterized by a sudden transition instead.

For most events, the signatures of impending failure can be observed weeks to months before the rupture. Furthermore, they are sufficiently similar from one cycle to the next for a model trained on past data to recognize them several years (cycles) afterwards. Whether the system remains stable enough for this exercise to hold over long periods of time is an open question—if more data was available, it might be the case that re-training on more recent data would be required after a while.

Note that because the events are most often separated by 13 or 14 months, the model cannot rely on seasonal signals to make its estimations. To prove this is the case, we show that a seasonal sinusoidal cannot lead to good estimations (i.e. simply knowing the time of the year, the ML model cannot make good predictions—see Supplementary). Therefore, the seismic data analyzed contains identifiable seismic patterns that match a slow slip cycle of 13–14 months and are independent of the season.

Because we choose to rely on transparent ML algorithms, we can identify the most important features used by our model to make its estimations of failure times, and therefore make comparison with laboratory experiments. We find that the best features identified in Cascadia follow very similar patterns compared to those identified in the laboratory. In the case of laboratory slow slip events, the most important feature by far for forecasting failure time is the seismic power^{13,34}, shown in Fig. 3c. In the case of the Cascadia subduction zone we rely on interquartile ranges, closely related to the root of the seismic power but with outlier values removed (these outliers are more likely to come from anthropogenic noise and/or local and distant earthquakes). Figure 3b illustrates that seismic energy follows

similar patterns in the laboratory and Cascadia: (i) a progressive increase, especially when failure approaches, with peaks in energy reached toward the end of the slow slip (as the slipping phase of the cycle emits larger seismic amplitudes than the loading phase, due to intense episodes of tremor within the frequency bands considered); and (ii) an often abrupt decrease within each slip cycle, towards the end of an event. Cycles in Cascadia are clearly apparent, but much noisier. This is likely due to background noise and to the fact that slow slip events occur over a very large region and might not be at the same stage everywhere.

In the laboratory, these signals carry fundamental information regarding the frictional state of the system. The fact that a similar behavior can be observed in Cascadia, with energy patterns matching the 14-month slow slip cycle, may also provide indirect information regarding the evolution of friction and slip rate at much larger scale. In simulated data, the origin of the seismic signal is due to the evolution of the force network in the granular gouge^{35,36}. In the Cascadia subduction zone, we can posit that part of this energy is emitted from a large number of asperities located on the fault interface. This strong resemblance suggests that some of the frictional physics may scale from the laboratory to subduction in Cascadia, bringing additional evidence for the self-similarity of slow slip nucleation and rupture, in the laboratory and in the field.

An exponential build-up in seismic energy preceding failure suggests a long nucleation phase for slow slip events in Cascadia. Stacking the evolution of our main feature for all cycles shows that seismic power within these tremor frequency bands increases exponentially as the fault approaches failure. Figure 4 shows the density distribution of the datapoints for all cycles considered (including events from both the training and the testing set). In the laboratory, we observe a clear exponential growth of the acoustic power, starting about 5–7 s before failure, for an interevent time of ca. 10 s. We interpret this increase in acoustic power as the signature of a growing slip instability. In the case of the Cascadia slow slip events, stacked data highlight a comparable increase in seismic power emerging ~100 days before failure.

This exponential build up observed preceding failure is consistent with laboratory experiments focusing on earthquake nucleation³⁷. During the nucleation of a slip instability, slip rate and rupture size both grow exponentially until the rupture has reached a size allowing for dynamic rupture propagation³⁸. The observed exponential build-up in seismic energy in tremor frequency bands can be explained by an exponential growth of

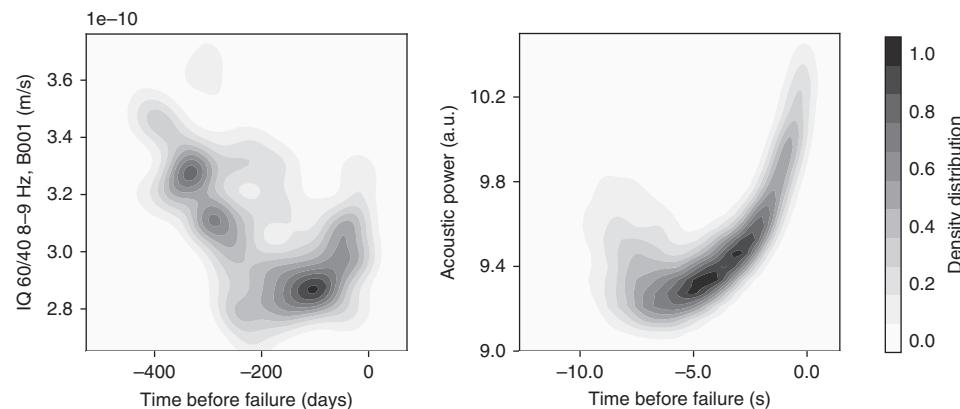


Fig. 4 Density distribution. Distribution of features with respect to time before failure, in between slip events—energy increases exponentially when failure becomes close. Right: laboratory slow slips. Left: Cascadia slow earthquakes.

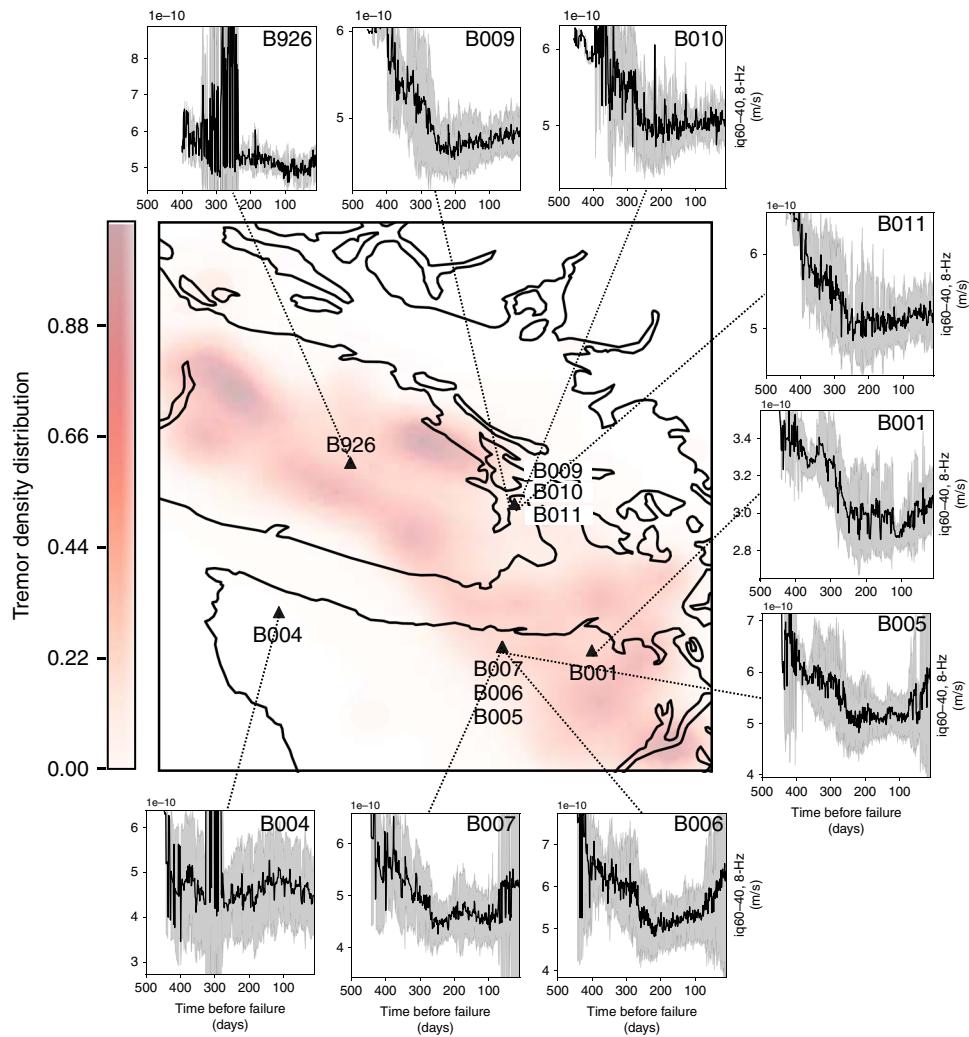


Fig. 5 Behavior on all stations. Mean (black curve) and standard deviation (gray shade) of our feature (seismic power in a tremor frequency band) for all stations considered, in-between slow slip events. For most stations, an increase in seismic power can be observed as failure approaches. This increase is particularly marked for the stations further inside the slowly slipping region. The only station outside of the slow slip area (B004) is also the only station to not exhibit such an increase in seismic power.

either the size or the slip rate (or both) of the nucleation phase of a slow slip event. In particular, we know from recent work that the same seismic features map accurately to the slip rate in both the laboratory^{13,39} and in Cascadia²⁰. Therefore, our results suggest that slow slip often begins with an exponential acceleration on the fault, that can be small enough to not be captured in catalogued tremor.

The evolution of stacked seismic power early in the slip cycle is characterized by a decreasing trend in both Cascadia and the laboratory. Part of this trend might come from the rearrangement of the system after failure, although the underlying drivers remain unclear. In Cascadia, the observed signals appear to be mostly driven by tectonics (with clear 14-month patterns), but superposed with low-amplitude seasonal fluctuations.

The increase in seismic power preceding failure can also be observed on other seismic stations (Fig. 5), in particular for the stations further inside the slowly slipping region (B001, B005, B006, B007, and B926 to a lesser extent as is appears noisier). Among these stations, B001 is characterized by the smallest standard deviation close to failure, which explains why it is the most useful station for our ML analysis (see Supplementary). Stations closer to the border of the area that experiences slow slip still show an increase, but not as pronounced (B009, B010, B011).

Interestingly, the only station analyzed that is located outside of the slow slip area (B004) is also the only station that does not show any increase preceding the rupture.

Discussion

We interpret the observed build-up in seismic noise within tremor frequency bands preceding failure as the signature of an increase in the number and intensity of low-amplitude tremors, too small to appear on several stations and be catalogued with array-based techniques. In our recent work⁴⁰, we showed that a neural network can detect many more tremors compared to the catalog it has been trained on. The number of tremors detected on surface stations in the same area by neural network seems to accelerate around 100 days before the rupture as well (see Supplementary), although the behavior preceding failure is not as clear as in the continuous seismic noise (from borehole stations). This is a strong argument in favor of our signal being driven by underlying tremor.

Because the signatures of upcoming failure can be recognized by our trained model many years afterwards, our results argue in favor of the predictability of slow slip rupture. If indeed slow slip events share a strong resemblance with earthquakes, long-term signatures of growing slip instabilities might be measurable preceding

destructive earthquakes as well. This hypothesis is in line with studies of the nucleation of large interplate earthquakes⁴¹, and with recent work on the detection of small-amplitude foreshock activity in California^{42,43} showing that foreshock activity may be observable preceding a significant fraction (i.e. 30% to 50% according to the analysis) of earthquakes with magnitude >4. Finer detection of small tremors and foreshocks may ultimately illuminate a long nucleation phase for large earthquakes, and pave the way for better constraints on rupture timing.

Methods

Seismic data and features used to build the models. We rely on seismic data from the the Plate Boundary Observatory³², from mid-2005 to late 2018. We build on our previous work related to the estimation of the GPS displacement from seismic data in Cascadia to create our features. From this previous work, we know that the best features that map to GPS displacement are intermediate inter-quartile ranges within the 8–13 Hz band⁴⁴.

The data processing is as follows:

- (i) We correct for the instrument gain, for all waveforms.
- (ii) We clip the signal at 5E–7 m/s, to limit the impact of earthquakes and anthropogenic noise on the analysis, and focus on low-amplitude signals.
- (iii) For every station and every day we then compute the 40–60 and 25–75 inter-quartile ranges, for the following frequency bands: 8–9, 9–10, 10–11, 11–12, 12–13 Hz. Because these features discard the values of the waveforms below and above a certain percentile, they are robust to outlier values.
- (iv) Once these features are built, anomalous feature datapoints are detected for each day and removed using Isolation Forests⁴⁵, with the automatic contamination threshold.

The features are then averaged over a time window and used as input to the ML models (the paper shows results with 3-month windows). The difference between the features' values at the beginning and the end of a given time window is also given as input to the algorithm. This leads to a total of $2^*5^*2 = 20$ features for each station. We build different models for each of the stations; combining stations or stacking station data does not lead to improved performance. Models built for each seismic station can be found in the Supplementary.

Model construction and model hyperparameters. We rely on the XGBoost library⁴⁶ for the gradient boosted trees' regression, shown in Fig. 2 of the paper (and for results presented below). The problem is posed in a regression setting. Model hyperparameters are set by five-fold cross-validation, using Bayesian optimization (skopt library⁴⁷). The associated hyper-parameters for the regression in Fig. 3 of the main text (B001 station, 3-month windows) are the following: colsample_bytree = 0.9000000000000002, learning_rate = 0.20419758706505553, max_depth = 8, min_child_weight = 5, n_estimators = 947, reg_alpha = 100, reg_lambda = 100, subsample = 0.9899999999999999.

Model evaluation. We use Pearson's CC as evaluation metric. The Pearson coefficient is a measure of the strength of the linear relationship between the predicted and ground truth values. For a perfect positive linear relationship, the Pearson coefficient is 1.

Data availability

The seismic data used was obtained from the Plate Boundary Observatory³², and is publicly available in the IRIS DMC repository (<https://ds.iris.edu/ds/nodes/dmc/earthscope/pbo/>). The data can be accessed through the Obspy open source library, or other seismic query open-source softwares.

Code availability

The code is based on two open-source libraries, XGBoost⁴⁶ and Scikit-Optimize⁴⁷, that are available online. We are currently building an open-source implementation of the code, that will be published online in the months to come.

Received: 31 March 2020; Accepted: 6 July 2020;

Published online: 18 August 2020

References

1. Obara, K. Nonvolcanic deep tremor associated with subduction in southwest japan. *Science* **296**, 1679–1681 (2002).
2. Obara, K. & Kato, A. Connecting slow earthquakes to huge earthquakes. *Science* **353**, 253–257 (2016).
3. Shelly, D. R., Beroza, G. C. & Ide, S. Non-volcanic tremor and low-frequency earthquake swarms. *Nature* **446**, 305–307 (2007).
4. Beroza, G. C. & Ide, S. Slow earthquakes and nonvolcanic tremor. *Annu. Rev. Earth Planet. Sci.* **39**, 271–296 (2011).
5. Peng, Z. & Gomberg, J. An integrated perspective of the continuum between earthquakes and slow-slip phenomena. *Nat. Geosci.* **3**, 599 (2010).
6. Gomberg, J. C. & Group, B. W. Slow-slip phenomena in Cascadia from 2007 and beyond: a review. *GSA Bull.* **122**, 963–978 (2010).
7. Rubinstein, J., Shelly, D. R. & Ellsworth, W. L. *Non-volcanic Tremor: A Window into the Roots of Fault Zones* (Springer Science+Business Media B.V., 2010).
8. Rousset, B., Bürgmann, R. & Campillo, M. Slow slip events in the roots of the San Andreas fault. *Sci. Adv.* **5**, eaav3274 (2019).
9. Jolivet, R. & Frank, W. B. The transient and intermittent nature of slow slip. *AGU Adv.* **1**, e2019AV000126 (2020).
10. Scholz, C. H. *The Mechanics of Earthquakes and Faulting*, i–i 3rd edn (Cambridge University Press, 2019).
11. Ougier-Simonin, A. & Zhu, W. Effect of pore pressure buildup on slowness of rupture propagation. *J. Geophys. Res.* **120**, 7966–7985 (2015).
12. Gao, X. & Wang, K. Rheological separation of the megathrust seismogenic zone and episodic tremor and slip. *Nature* **543**, 416 (2017).
13. Hulbert, C. et al. Similarity of fast and slow earthquakes illuminated by machine learning. *Nat. Geosci.* **12**, 69 (2019).
14. Michel, S., Gualandi, A. & Avouac, J.-P. Similar scaling laws for earthquakes and Cascadia slow-slip events. *Nature* **574**, 522–526 (2019).
15. Jolivet, R. et al. The burst-like behavior of aseismic slip on a rough fault: the creeping section of the Haiyuan Fault, China. *Bull. Seismol. Soc. Am.* **105**, 480–488 (2014).
16. Frank, W. B. & Brodsky, E. E. Daily measurement of slow slip from low-frequency earthquakes is consistent with ordinary earthquake scaling. *Sci. Adv.* **5**, eaaw9386 (2019).
17. Marone, C. Laboratory-derived friction laws and their application to seismic faulting. *Annu. Rev. Earth Planet. Sci.* **26**, 643–696 (1998).
18. Kaproth, B. M. & Marone, C. Slow earthquakes, preseismic velocity changes, and the origin of slow frictional stick-slip. *Science* **341**, 1229–1232 (2013).
19. Scuderi, M., Marone, C., Tinti, E., Di Stefano, G. & Collettini, C. Precursory changes in seismic velocity for the spectrum of earthquake failure modes. *Nat. Geosci.* **9**, 695 (2016).
20. Rouet-Leduc, B., Hulbert, C. & Johnson, P. A. Continuous chatter of the Cascadia subduction zone revealed by machine learning. *Nat. Geosci.* **12**, 75–79 (2019).
21. Perol, T., Gharbi, M. & Denolle, M. Convolutional neural network for earthquake detection and location. *Sci. Adv.* **4**, e1700578 (2018).
22. Ross, Z. E., Meier, M.-A. & Hauksson, E. P wave arrival picking and first-motion polarity determination with deep learning. *J. Geophys. Res.* **123**, 5120–5129 (2018).
23. McBrearty, I. W., Delaney, A. A. & Johnson, P. A. Pairwise association of seismic arrivals with convolutional neural networks. *Seismol. Res. Lett.* **90**, 503–509 (2019).
24. Ross, Z. E., Yue, Y., Meier, M.-A., Hauksson, E. & Heaton, T. H. Phaselink: a deep learning approach to seismic phase association. *J. Geophys. Res.* **124**, 856–869 (2019).
25. Holtzman, B. K., Paté, A., Paisley, J., Waldhauser, F. & Repetto, D. Machine learning reveals cyclic changes in seismic source spectra in geysers geothermal field. *Sci. Adv.* **4**, eaao2929 (2018).
26. Rogers, G. & Dragert, H. Episodic tremor and slip on the Cascadia subduction zone: the chatter of silent slip. *Science* **300**, 1942–1943 (2003).
27. Kao, H. et al. A wide depth distribution of seismic tremors along the northern Cascadia margin. *Nature* **436**, 841–844 (2005).
28. Wech, A. G., Creager, K. C. & Melbourne, T. I. Seismic and geodetic constraints on Cascadia slow slip. *J. Geophys. Res.* **114**, B10316 (2009).
29. Wech, A. G. & Bartlow, N. M. Slip rate and tremor genesis in Cascadia. *Geophys. Res. Lett.* **41**, 392–398 (2014).
30. Hawthorne, J. C., Bostock, M. G., Royer, A. A. & Thomas, A. M. Variations in slow slip moment rate associated with rapid tremor reversals in Cascadia. *Geochim. Geophys. Geosyst.* **17**, 4899–4919 (2016).
31. Frank, W. B. Slow slip hidden in the noise: the intermittence of tectonic release. *Geophys. Res. Lett.* **43**, 10,125–10,133 (2016).
32. Silver, P. G. et al. A plate boundary observatory. *Iris Newsl.* **16**, 3 (1999).
33. Friedman, J. H. Greedy function approximation: a gradient boosting machine. *Ann. Stat.* **29**, 1189–1232 (2001).
34. Rouet-Leduc, B. et al. Machine learning predicts laboratory earthquakes. *Geophys. Res. Lett.* **44**, 9276–9282 (2017).
35. Gao, K., Guyer, R. A., Rougier, E., Ren, C. X. & Johnson, P. A. From force chains to acoustic emission. *Phys. Rev. Lett.* **123**, 048003 (2019).
36. Ren, C. X. et al. Machine learning reveals the state of intermittent frictional dynamics in a sheared granular fault. *Geophys. Res. Lett.* **46**, 7395–7403 (2019).

37. Latour, S., Schubnel, A., Nielsen, S., Madariaga, R. & Vinciguerra, S. Characterization of nucleation during laboratory earthquakes. *Geophys. Res. Lett.* **40**, 5064–5069 (2013).
38. Ohnaka, M. A constitutive scaling law and a unified comprehension for frictional slip failure, shear fracture of intact rock, and earthquake rupture. *J. Geophys. Res.* **108**, 2080 (2003).
39. Rouet-Leduc, B. et al. Estimating fault friction from seismic signals in the laboratory. *Geophys. Res. Lett.* **45**, 1321–1329 (2018).
40. Rouet-Leduc, B., Hulbert, C., McBrearty, I. W. & Johnson, P. A. Probing slow earthquakes with deep learning. *Geophys. Res. Lett.* **47**, e2019GL085870 (2020).
41. Bouchon, M., Durand, V., Marsan, D., Karabulut, H. & Schmittbuhl, J. The long precursory phase of most large interplate earthquakes. *Nat. Geosci.* **6**, 299–302 (2013).
42. Trugman, D. T. & Ross, Z. E. Pervasive foreshock activity across southern California. *Geophys. Res. Lett.* **46**, 8772–8781 (2019).
43. van den Ende, M. P. A. & Ampuero, J.-P. On the statistical significance of foreshock sequences in southern California. *Geophys. Res. Lett.* **47**, e2019GL086224 (2020).
44. Rouet-Leduc, B. et al. Estimating fault friction from seismic signals in the laboratory. *Geophys. Res. Lett.* **45**, 1321–1329 (2018).
45. Pedregosa, F. et al. Scikit-learn: machine learning in python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
46. Chen, T. & Guestrin, C. Xgboost: a scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794 <https://doi.org/10.1145/2939672.2939785> (ACM, 2016).
47. Head, T. et al. *Scikit-optimize: Sequential Model-based Optimization in Python* (Zenodo, 2017).

Acknowledgements

We thank C. Duverger and A. Schubnel for useful discussions and comments. We thank C. Marone, C. Bolton, and J. Rivière for the laboratory data. C.H. and B.R.-L. were supported by Institutional Support (LDRD) at Los Alamos. C.H.'s work was also supported by a joint research laboratory effort in the framework of the CEA-ENS Yves Rocard LRC (France), and both C.H. and R.J. were funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Geo-4D project, grant agreement 758210). P.A.J. and CH were supported by DOE Office of Science (Geoscience Program, grant 89233218CNA000001).

Author contributions

C.H. and B.R.-L. conducted the analysis. R.J. helped analyzing the results and provided expertise on earthquake nucleation. P.A.J. supervised the project. All authors contributed to writing the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41467-020-17754-9>.

Correspondence and requests for materials should be addressed to C.H.

Peer review information *Nature Communications* thanks Benjamin Holtzman, Quentin Bletery and the other anonymous reviewer for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

This is a U.S. government work and not under copyright protection in the U.S.; foreign copyright protection may apply 2020